Risk Management for Central Banks and Other Public Investors

EDITED BY
Ulrich Bindseil, Fernando González
and Evangelos Tabakis

CAMBRIDGE

CAMBRIDGE www.cambridge.org/9780521518567

This page intentionally left blank

Risk Management for Central Banks and Other Public Investors

Domestic and foreign financial assets of all central banks and public wealth funds worldwide are estimated to have reached more than USD 12 trillion in 2007. How do these institutions manage such unprecedented growth in their financial assets and how have they responded to the 'revolution' of risk management techniques during the last fifteen years? This book surveys the fundamental issues and techniques associated with risk management and shows how central banks and other public investors can create better risk management frameworks. Each chapter looks at a specific area of risk management, first presenting general problems and then showing how these materialize in the special case of public institutions. Written by a team of risk management experts from the European Central Bank, this much-needed survey is an ideal resource for those concerned with the increasingly important task of managing risk in central banks and other public institutions.

Ulrich Bindseil is Head of the Risk Management Division at the European Central Bank.

Fernando González is Principal Economist at the European Central Bank.

Evangelos Tabakis is Deputy Head of the Risk Management Division at the European Central Bank.

Risk Management for Central Banks and Other Public Investors

Edited by

Ulrich Bindseil, Fernando González and

Evangelos Tabakis

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK
Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9780521518567

© Cambridge University Press 2009

This publication is in copyright. Subject to statutory exception and to the provision of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press.

First published in print format 2009

ISBN-13 978-0-511-47916-8 eBook (EBL)

ISBN-13 978-0-521-51856-7 hardback

Cambridge University Press has no responsibility for the persistence or accuracy of urls for external or third-party internet websites referred to in this publication, and does not guarantee that any content on such websites is, or will remain, accurate or appropriate.

Contents

List of figures *page* x
List of tables xii
List of boxes xv
Foreword xvii
José-Manuel González-Páramo
Introduction xx
Ulrich Bindseil, Fernando González and Evangelos Tabakis

Part I Investment operations 1

1 Central banks and other public institutions as financial investors 3
Ulrich Bindseil
1 Introduction 3
2 Public institutions' specificities as investors 4
3 How policy tasks have made central banks large-scale investors 10
4 Optimal degree of diversification of public institutions' financial assets 17
5 How actively should public institutions manage their financial assets? 23
6 Policy-related risk factors 29
7 The role of central bank capital – a simple model 34
8 Integrated risk management for public investors 41
9 Conclusions 48

2 Strategic asset allocation for fixed-income investors 49
Matti Koivu, Fernando Monar Lora, and Ken Nyholm
1 Introduction 49
2 A primer on strategic asset allocation 50
3 Components of the ECB investment process 68

4 Forward-looking modelling of the stochastic factors 75
5 Optimization models for SAA under a shortfall approach 89
6 The ECB case: an application 99

3 Credit risk modelling for public institutions' investment portfolios 117
Han van der Hoorn
1 Introduction 117
2 Credit risk in central bank and other public investors' portfolios 118
3 The ECB's approach towards credit risk modelling: issues and parameter choices 122
4 Simulation results 143
5 Conclusions 155

4 Risk control, compliance monitoring and reporting 157
Andres Manzanares and Henrik Schwartzlose
1 Introduction 157
2 Overview of the distribution of portfolio management tasks within the Eurosystem 159
3 Limits 161
4 Portfolio management oversight tasks 179
5 Reporting on risk and performance 189
6 IT and risk management 196

5 Performance measurement 207
Hervé Bourquin and Roman Marton
1 Introduction 207
2 Rules for return calculation 208
3 Two-dimensional analysis: risk-adjusted performance measures 213
4 Performance measurement at the ECB 219

6 Performance attribution 222
Roman Marton and Hervé Bourquin
1 Introduction 222
2 Multi-factor return decomposition models 224
3 Fixed-income portfolios: risk factor derivation 228
4 Performance attribution models 241
5 The ECB approach to performance attribution 257
6 Conclusions 267

Part II: Policy operations 269

7 Risk management and market impact of central bank credit operations 271
Ulrich Bindseil and Francesco Papadia
1 Introduction 271
2 The collateral framework and efficient risk mitigation 274
3 A cost–benefit analysis of a central bank collateral framework 284
4 Conclusions 300

8 Risk mitigation measures and credit risk assessment in central bank policy operations 303
Fernando González and Phillipe Molitor
1 Introduction 303
2 Assessment of collateral credit quality 307
3 Collateral valuation: marking to market 315
4 Haircut determination methods 318
5 Limits as a risk mitigation tool 337
6 Conclusions 338

9 Collateral and risk mitigation frameworks of central bank policy operations – a comparison across central banks 340
Evangelos Tabakis and Benedict Weller
1 Introduction 340
2 General comparison of the three collateral frameworks 342
3 Eligibility criteria 348
4 Credit risk assessment and risk control framework 353
5 Conclusions 357

10 Risk measurement for a repo portfolio – an application to the Eurosystem's collateralized lending operations 359
Elke Heinle and Matti Koivu
1 Introduction 359
2 Simulating credit risk 360
3 Simulating liquidity-related risks 366
4 Issues related to concentration risks 368
5 Risk measures: Credit Value-at-Risk and Expected Shortfall 376
6 An efficient Monte Carlo approach for credit risk estimation 379

7 Residual risk estimation for the Eurosystem's credit operations 387
8 Conclusions 393

11 Central bank financial crisis management from a risk management perspective 394
Ulrich Bindseil
1 Introduction 394
2 Typology of financial crisis management measures 396
3 Review of some key results of the literature 399
4 Financial stability role of central bank operational framework 416
5 The inertia principle of central bank risk management in crisis situations 418
6 Equal access FCM measures 422
7 FCM measures addressed to individual banks (ELA) 434
8 Conclusions 437

Part III: Organizational issues and operational risk 441

12 Organizational issues in the risk management function of central banks 443
Evangelos Tabakis
1 Introduction 443
2 Relevance of the risk management function in a central bank 444
3 Risk management best practices for financial institutions 445
4 Six principles in the organization of risk management in central banks 448
5 Conclusions 459

13 Operational risk management in central banks 460
Jean-Charles Sevet
1 Introduction 460
2 Central bank specific ORM challenges 463
3 Definition of operational risk 465
4 ORM as overarching framework 468
5 Taxonomy of operational risk 469
6 The ORM lifecycle 471
7 Operational risk tolerance policy 472
8 Top-down self-assessments 476

9 Bottom-up self-assessments 479
10 ORM governance 483
11 KRIs and ORM reporting 484
12 Conclusions 488

References 490
Index 507

Figures

2.1	Evolution of Strategic Asset Allocation	*page* 53
2.2	The efficient frontier	59
2.3	Adapted efficient frontier and VaR constraint	65
2.4	Efficient frontier in E[r]–VaR space	66
2.5	Components of an investment process	69
2.6	The overall policy structure of the investment process	73
2.7	Modular structure of SAA tools	76
2.8	Generic yield curves	103
2.9	Normal macroeconomic evolution: (a) GDP YoY % Growth; (b) CPI YoY % Growth	105
2.10	Projected average evolution of the US Government yield curve in a normal example	106
2.11	Projected distribution of yields in a normal example: (a) US Gov 0–1Y; (b) US Gov 7–10Y	107
2.12	Distribution of returns in a normal example: (a) US Gov 0–1Y; (b) US Gov 7–10Y	109
2.13	Inflationary macroeconomic evolution: (a) GDP YoY % Growth; (b) CPI YoY % Growth	112
2.14	Projected average evolution of the US Government yield curve in a non-normal example	113
2.15	Projected distribution of yields in a non-normal example: (a) US Gov 0–1Y; (b) US Gov 7–10Y	113
2.16	Distribution of returns in a non-normal example: (a) US Gov 0–1Y; (b) US Gov 7–10Y	115
3.1	Asset value and migration (probabilities not according to scale)	130
3.2	Impact of asset correlation on portfolio risk (hypothetical portfolio with 100 issuers rated AAA–A, confidence level 99.95%).	142
3.3	Comparison of portfolios by rating and by industry	144
3.4	Simulation results for Portfolio I	146
3.5	Comparison of simulation results for Portfolios I and II	152

3.6	Lorenz curves for Portfolios I and II	153
3.7	Sensitivity analysis for Portfolio I	155
7.1	Marginal costs and benefits for banks of posting collateral with the central bank	288
7.2	One-week moving average spread between non-EEA and EEA issuers in 2005	293
7.3	Spread between the three-month EURIBOR and three-month EUREPO rates since the introduction of the EUREPO in March 2002 – until end 2007	294
7.4	Evolution of MRO weighted average, 1 Week repo, and 1 Week unsecured interbank rates in 2007	298
7.5	Evolution of LTRO weighted average, 3M repo, and 3M unsecured interbank rates in 2007	298
8.1	Risks involved in central bank repurchase transactions	305
8.2	Basic determinants of haircut calculations	319
8.3	Holding period	320
8.4	Relationship between position size and liquidation value	324
8.5	Yield-curve differentials	328
8.6	Value-at-Risk due to credit risk for a single exposure	334
10.1	Important types of concentrations in the Eurosystem collateral framework	369
10.2	Lorenz curve for counterparties with respect to amount of collateral submitted	370
10.3	Lorenz curve for collateral issuers with respect to amount of collateral submitted	372
10.4	Herfindahl–Hirschmann Indices (HHI) of individual counterparties with respect to their collateral submitted	375
10.5	Variance reduction factors, for varying values of $\hat{\theta}$ and asset correlations	386
10.6	The effect on Expected Shortfall of changed liquidation time assumptions. *Source:* ECB's own calculations	390
10.7	The effect on Expected Shortfall of changed credit quality assumptions. *Source:* ECB's own calculations	391
10.8	The effect on Expected Shortfall of changed assumptions on issuer-counterparty correlations	391
11.1	Liquidity shocks and associated marginal costs to a specific bank	424
13.1	Taxonomy of operational risk	470
13.2	Drivers of the risk impact-grading scale of the ECB	474
13.3	Operational risk tolerance: illustrative principles	475

Tables

1.1 Foreign reserves (and domestic financial asset of G3 central banks) in December 2007 *page* 13
1.2 Different reasons for holding foreign exchange reserves – importance attributed by reserve managers according to a JPMorgan survey 15
1.3 Risk quantification and economic capital, in billions of EUR, as at end 2005 16
1.4 Modified duration of fixed-income market portfolios 19
1.5 Asset classes used by central banks in their foreign reserves management 21
1.6 Asset classes currently allowed or planned to be allowed according to a JPMorgan survey 22
1.7 Derivatives currently allowed or planned to be allowed according to a JPMorgan survey 23
1.8 Trading styles of central bank reserves managers according to a JPMorgan survey 28
2.1 Example of the eligible investment universe for a USD portfolio 100
2.2 Classification scheme 102
2.3 Transition matrices 102
2.4 Intercepts of the Nelson–Siegel state equation 102
2.5 Autoregressive coefficients of the Nelson–Siegel state equation 103
2.6 Returns in a normal example: average and standard deviation 108
2.7 Optimal portfolio composition in a normal example 110
2.8 Summary information for the optimal portfolio in a normal example 110
2.9 Returns in a non-normal example: average and standard deviation 114
2.10 Optimal portfolio composition in a non-normal example 116
2.11 Summary information for the optimal portfolio in a non-normal example 116
3.1 Migration probabilities and standard normal boundaries for bond with initial rating A 129

3.2	Risk-weighting of Standardized Approach under Basel II	135
3.3	Original and augmented migration probabilities for bond with initial rating A	140
3.4	Common migration matrix (one-year migration probabilities)	146
3.5	Parameters for Nelson–Siegel curves	146
3.6	Simulation results for Portfolio I	147
3.7	Decomposition of simulation results into default and migration	148
3.8	Simulation results for Portfolio II, including decomposition	151
3.9	Sensitivity analysis for Portfolio I	154
4.1	Rating scales, numerical equivalents of ratings and correction factors for counterparty limits	174
7.1	Shares of different types of collateral received by 113 institutions responding to the 2006 ISDA margin survey	279
7.2	Comparison of the key recommendations of ISDA Guideline for Collateral Practitioners with the Eurosystem collateralization framework	281
7.3	Bid–ask spreads as an indicator of liquidity for selected assets (2005 data)	283
7.4	Example of parameters underlying a cost–benefit analysis of collateral eligibility	287
7.5	Social welfare under different sets of eligible collateral and refinancing needs of the banking system	288
7.6	Information on the set of bonds used for the analysis	292
7.7	Spreads containing information on the GC and Eurosystem collateral eligibility premia – before and during the 2007 turmoil	299
8.1	Summary of ECAF by credit assessment source in the context of the Single List	314
8.2	Liquidity score card	331
8.3	Eurosystem liquidity categories for marketable assets	332
8.4	Eurosystem levels of valuation haircuts applied to eligible marketable assets in relation to fixed coupon and zero-coupon instruments	333
8.5	The distribution of bond values of an A rated bond	335
8.6	'Through-the-cycle' credit migration matrix	336
8.7	'Point-in-time' credit migration matrix	336
8.8	99 per cent credit risk haircut for a five-year fixed coupon bond	337
9.1	Differentiation of collateral policy depending on type of operation	343
9.2	Comparison of sizes of credit operations (averages for 2006, in EUR billions)	346

9.3	Comparison of eligibility criteria	350
9.4	Comparison of haircuts applied to government bonds	355
9.5	Comparison of haircuts of assets with a residual maturity of five years	355
10.1	Default probabilities for different rating grades	363
10.2	Liquidation time assumptions used for the different asset classes	364
10.3	Comparison of various variance reduction techniques with 0.24 asset correlation	387
10.4	Comparison of various variance reduction techniques with 0.5 asset correlation	387
10.5	Breakdown of residual risks in the base case scenario	389
10.6	Composition of submitted collateral over time and composition of residual financial risks over time	392
11.1	FCM typology and illustration from August–December 2007	400

Boxes

2.1	The VAR macro model	*page* 78
2.2	Transformation of yields and relative slope	83
3.1	Credit spreads and the limitations of diversification	121
4.1	Modified duration versus VaR	163
4.2	Calculation of rate reasonability tolerance bands at the ECB	184
4.3	ECB Risk Management – Regular reports	193
4.4	The systems used by the ECB Risk Management Division (RMA)	201
8.1	Historical background in the creation of in-house credit assessment systems in four Eurosystem central banks	310
8.2	In-house credit assessments by the Bank of Japan	311
8.3	The Qualified Loan Review programme of the Federal Reserve	312
9.1	Survey of credit and market risk mitigation in a collateral management in central banks	356

Foreword

The reader familiar with central bank parlance will have certainly noticed that our vocabulary is full of references to risks. It seems that no speech of ours can avoid raising awareness of risks to price stability or evade the subject of risks to the smooth functioning of the financial system. Indeed, one way to describe our core responsibility is to say that the central bank acts as a risk manager for the economy using monetary policy to hedge against inflationary risks. However, we tend to be less willing to share information on the ways we manage financial risks in our own institutions. It is thus not surprising that a book that sheds light on risk management in the world of central banks and other public investors in a systematic and comprehensive way has not been published so far. And I am very happy that the initiative to prepare such a book has been taken by staff of the European Central Bank.

Central banks' own appetite for financial risks is not always easy to understand. Our institutions have historically been conservative investors, placing their foreign reserves mostly in government securities and taking very little, if any, credit risk. Progressively, the accumulation of reserves in some countries, either as a result of their abundant natural resources or of foreign exchange policies, has led their central banks to expand their investment universe and, with it, the financial risks they face. More recently, the landscape of public investors has been enriched by sovereign wealth funds, state-backed investors from emerging economies that made their presence more than noticeable in international capital markets and have occasionally created controversy with their investment strategies.

While managing investment portfolios is one area where risk management expertise is needed, central banks have other core concerns. They are in charge of monetary policy in their jurisdiction. They are also expected to intervene when the stability of the financial system is at stake. In order to steer the system out of a crisis, they are prepared, if needed, to take those

risks which other market participants rush to shed. They are prepared to provide additional liquidity to the system as a whole or lend to specific banks on special conditions. Such behaviour, which may seem to put risk management considerations on hold, at least temporarily, further complicates the effort of an outsider to understand the role of risk management in the central bank.

Being responsible for risk management in a public institution, like a central bank, does not simply rely on technical risk management expertise. Although the requirement for a high degree of fluency in quantitative techniques is not less important than in private financial institutions, it must be combined with a deep understanding of the role of the public institution and its core functions. In our institutions, financial decisions are not taken based only on risk and return considerations but also take into account broader social welfare aspects.

Central bank risk managers provide decision makers with assessments of financial risks in the whole range of central banks' operations, whether these are linked to policy objectives or are related to the management of investment portfolios. They should be able to deliver such assessments not only under normal market conditions but, even more so, under conditions of market stress. Decision makers also seek their advice to understand and draw the right conclusions from the use of the latest instruments of risk transfer in the markets and the implementation of risk management strategies by financial institutions in our jurisdictions.

The European Central Bank placed, from the very beginning, particular attention to risk management. As a new member of the central bank community, it had the ambition of fulfilling the highest governance standards in organizing its risk management function within the institution and applying state-of-the-art tools. No less than that would be expected from a new central bank that would determine monetary policy and oversee financial stability for an ever-increasing number of European citizens, playing the lead role in a system of cooperating central banks.

Central banks and other public investors have been entrusted with the management of public funds and are expected to do so in a transparent way that is well understood by the public. This book systematically explains how central banks have addressed financial risks in their operations. It discusses issues of principle but also provides concrete practical information. It explains how risk management techniques, developed in the private sector, apply to central banks and where idiosyncrasies of our institutions merit special approaches. The blend of analysis and information provided in the

next pages makes me confident that this book will find an eager readership among both risk managers and central bankers.

José Manuel González-Páramo
Member of the Executive Board of
the European Central Bank

Introduction

Ulrich Bindseil, Fernando González and Evangelos Tabakis

Domestic and foreign financial assets of central banks and public wealth funds worldwide are estimated to have reached in 2007 more than USD 12 trillion, which is more than 15 per cent of world GDP, and more than 10 per cent of the global market capitalization of equity and fixed-income securities markets. Reflecting unprecedented growth of their financial assets, and the revolution of risk management techniques and best practices during the last fifteen years, the investment and risk management policies and procedures of central banks and other public investors have undergone a profound transformation. The purpose of this book is to provide a comprehensive and structured overview of issues and techniques in the area of public institutions' risk management. On each of the main areas of risk management, the book aims at first presenting the general problems as they also would occur in private financial institutions, then to discuss how these materialize in the special case of public institutions, and finally to illustrate this general discussion by describing the European Central Bank's (ECB) specific approach. Due consideration is given to the specificities of public institutions in general and central banks in particular. On the one side, their public character relates to certain policy tasks, which will also impact on their investment policies, in particular with regard to assets which are directly considered policy assets (e.g. monetary policy assets, foreign reserves to stand ready for intervention purposes). Secondly, the public character of these institutions has certain implications regardless of policy tasks, such as particular duties of transparency and accountability, less flexibility in terms of human resource policies and contracting, being outside the radar of regulators, etc. These characteristics will also influence optimal investment policies and risk management techniques of public institution.

The book targets portfolio managers, risk managers, monetary policy implementation experts of central banks and public wealth funds, and staff in supranational financial institutions working on similar issues. Moreover, staff from the financial industry who provide services to central banks would also have an interest in this book. Similarly, treasury and liquidity managers of banks will find the risk management perspective of central banks' liquidity

providing operations useful in understanding central bank policies. Around a half of the chapters also provide methodological discussions which are not really specific to central banks or other public investors, but which are equally relevant for any other institutional investors. Finally, students in both finance and central banking will find the book important as bridging theory and practice and as providing insights in a key area of central banking other than monetary policy on which very little has traditionally been made public.

The authors of this book all work or worked in the ECB's Risk Management Division (except two, who work in the ECB's Directorate General Market Operations), and the topics covered reflect the area of expertise of the respective authors. Thus, the book obviously reflects the experience of the ECB and the specific challenges it has had to address. Nevertheless, the book aims at working out the generic specificities and issues relating to all public institutions' risk management functions.

There are two types of books with which the present one can be compared. First, there are a number of books on central bank investment policies and risk management, like Bernadell *et al.* (2004), Danmarks Nationalbank (2004), Pringle and Carver (2007, but also previous editions), Johnson-Calari and Rietveld (2007) or Bakker and van Herpt (2007). These books however do not aim at being comprehensive and conceptually structured, nor do they go really into depth. In contrast, the present book is intended to be a comprehensive reference book, structured along the main areas of central bank investment and risk management, reviewing systematically the existing literature, going into depth, and using state-of-the art methods. Second, there are at least two recent books by teams from the institutional investor/asset allocation area of major investment banks, namely Litterman (2003) and Dynkin *et al.* (2006). These books are similar in authorship as they are produced by a team of experts from one institution and cover topics in the broader area of financial management, including risk management. However the two books have a different perspective, namely that of investment management, and do not cover the risk control and risk mitigation aspects of risk management.

Structure of the book: Investment *vs.* policy operations; different risk types

The book is structured into three main parts: the first deals with the risk management for *investment operations* of public institutions. Investment

operations are defined broadly as financial operations of public institutions which are not or only limitedly constrained by the policy mandates of the public institution. Still, the public character of the institution should influence its investment and risk management policies, relative to a non-public institutional investor. The second part deals with *policy operations* of central banks, whereby the focus is on collateralized lending operations, as such monetary policy operations are standard today for central banks to control short-term interest rates. Most issues arising in this context are, however, also relevant for collateralized lending programmes that a financial institution would establish, and techniques discussed are therefore relevant for the financial industry. Finally, a short third part deals with organizational issues and operational risk management in public financial institutions.

While the segregation of risk management approaches into those relating to investment and those relating to policy operations may seem straightforward for central bankers, its compatibility with the idea of integrated financial risk management may be questioned. Why wouldn't all risks be mapped eventually into one risk framework? It appears a standard problem of any bank that risks from different business lines seem at a first look difficult to aggregate, but that these problems need to be overcome because segregated risk management is inferior. In contradiction to this, in many central banks, the organizational units for risk management are segregated: one would be responsible for investment operations, and the other for policy operations. In the case of the ECB, both risk management functions are assigned to one division, not to aggregate risk across the two 'business lines', but for achieving intellectual economies of scale and scope. A probably valid explanation in the case of the ECB for not integrating the two business lines in terms of risk management is that monetary policy operations are in the books of the national central banks (NCBs) of the Eurosystem, and not in the books of the ECB. Therefore, also, losses would arise with NCBs. The responsibility of the ECB's risk management for defining the risk framework for policy operations is based on the fact that losses relating to monetary policy operations are *shared* across NCBs. In contrast, the ECB's investment operations are genuinely in the books of the ECB, and directly affect its P&L. Therefore, integrating the two types of operations would mean ignoring that the associated P&Ls are not for the same institutions, and thus should be part of different risk budgets, etc. While the ECB has thus a valid excuse for keeping the two issues separated, which affects the structure of the current book, other central banks should

probably not follow this avenue, as all types of operations end up affecting their P&L.

The structure of this book from the risk *type* perspective may appear less clear than for a typical risk management textbook. While Chapter 3 is purely on the credit risk side, Chapters 2, 5 and 6 are about market risk management. Chapters 7–10 are mainly on the credit risk side; however, potential losses in reverse repo operations are also driven by liquidity and market risk when it comes to liquidating collateral in the case of a counterparty default. Chapter 4 addresses risk control tasks aiming at both credit and market risk. Operational risk management as discussed in Chapter 13 is a rather different animal, but as operational risk contributes in Basel II a third component to capital requirements, it is thought that a book on public institutions' risk management would be incomplete if not also discussing, at least in one chapter, issues relating to operational risk in public institutions. In the ECB, the more limited interaction between operational and financial risk management is reflected by having separate entities being responsible for each.

Part I: Investment operations

Part I of the book, on investment operations, begins with a chapter (**Central banks and other public institutions as financial investors**) discussing the 'nature' of central banks and other public institutions as investors. The chapter aims at providing tentative answers to questions like: What are the special characteristics of such investors implied by their policy mandates? What are the basic risk–return properties of their balance sheets? What capital do they need and what are their long-run financial perspectives? In which sense should they be 'active' investors and how diversified should they be? Are they unique in terms of aversion against reputation risk? The chapter suggests that while on one side, many financial industry risk management techniques (like VaR, limit setting, reporting, performance attribution) are directly applicable to public institutions, the foundations of integrated risk management (e.g. risk budgeting, economic capital calculation, desired credit rating) are very special for public institutions, and in fact are more difficult to derive than in the case of a private financial institution.

Chapter 2 (**Strategic asset allocation for central banks and public investors**) contains a general introduction to strategic asset allocation and a

review of the key issues relating to it. It also provides a review of central bank practice in this area (also on the basis of available surveys), and a detailed technical presentation of the ECB's approach to strategic asset allocation. The importance of strategic asset allocation in public institutions can hardly be overestimated, since it typically drives more than 90 per cent of the risks and returns of public institution's investments. This also reflects the need for transparency of public investments, which can be fulfilled in principle by a strategic asset allocation approach, but less by 'active management' investment strategies.

Chapter 3 discusses **Credit risk modelling for public institutions' investment portfolios**. Portfolio credit risk modelling in general has emerged in practice only over the least ten years, and in public institutions only very recently. Its relevance for central banks, for example, is on the one hand obvious in view of the size of the portfolios in questions, and their increasing share of non-government bonds. On the other hand, public investors tend to hold credit portfolios of very high average credit quality, still concentrated in a limited number of issuers, which poses specific challenges for estimating sensible credit risk measures.

Chapter 4 on **Risk control, compliance monitoring and reporting** turns to the core regular risk control tasks that any institutional financial investor should undertake. There is typically little systematic literature on these topics which are so relevant and also often challenging in practice.

Chapter 5 on **Performance measurement** again deals in more depth with one core risk control subject of interest to all institutional investors. While in principle being a very practical issue, it often raises numerous technical implementation issues. Chapter 6, on **Performance attribution** complements Chapter 5. While performance attribution is a topic which can fill a book in its own right, this chapter includes a discussion of the most fundamental principles and considerations when applying performance attribution in the typical central bank setting. In addition, the fixed-income attribution framework currently applied by the European Central Bank is introduced.

Part II: Policy operations

Chapters 7 to 11 cover central bank **policy operations** conducted as reverse repo operations. Chapter 7 on **Risk management and market impact of central bank credit** reviews the role and effects of the collateral framework

which central banks, for example, use in conducting temporary monetary policy operations. First, the chapter explains the design of such a framework from the perspective of risk mitigation. It is argued that by means of appropriate risk mitigation measures, the residual risk on any potentially eligible asset can be equalized and brought down to the level consistent with the risk tolerance of the central bank. Once this result has been achieved, eligibility decisions should be based on an economic cost–benefit analysis. The chapter looks at the effects of the collateral framework on financial markets, and in particular on spreads between eligible and ineligible assets.

Chapter 8 goes in more depth with regard to methodological issues of **risk mitigation measures and credit risk assessments in central bank policy operations**. It motivates in more detail the different risk mitigation measures, and how they are applied in the Eurosystem. In particular, valuation issues and haircut setting are explained. To ensure that accepted collateral fulfils sufficient credit quality standards, central banks tend to rely on external or internal credit quality assessments. While many central banks today rely exclusively on ratings by rating agencies, others still rely on internal credit quality assessment systems.

Chapter 9 provides **a comparison of risk mitigation measures and credit risk assessment in central bank policy operations** across in particular three major central banks, namely the Federal Reserve, the Bank of Japan and the Eurosystem.

Chapter 10 (**Risk measurement for a repo portfolio**) presents a state-of-the art approach to estimating tail risk measures for a portfolio of collateralized lending operations, as it is relevant for any investor with a large repo portfolio, and as it has been implemented for the first time by a central bank in 2006 by the ECB.

Chapter 11 turns to **central bank financial crisis management from a risk management perspective**. Financial crisis management is a key central bank policy task and unsurprisingly financial transactions in such an environment will imply particular risk taking, which needs to be well justified and well controlled. The second half of 2007 provided multiple illustrations for this chapter.

Part III: Organizational issues and operational risk

Part three of the book consists of Chapters 12 and 13. Chapter 12 is on **Organizational issues in the risk management function of central banks**,

and covers organizational issues of relevance for any institutional investor, such as segregation of duties; Chinese walls; policy *vs.* investment operations, optimal boundaries of responsibilities vis-à-vis other business areas etc. The final Chapter 13 treats **Operational risk management in central banks** and presents in some detail the ECB's approach to this.

Part I

Investment operations

1 Central banks and other public institutions as financial investors

Ulrich Bindseil

1. Introduction

Domestic and foreign financial assets of all central banks and public wealth funds worldwide are estimated to have reached in 2007 more than USD 12 billion. Public investors, hence, are important players in global financial markets, and their investment decisions will both matter substantially for their (and hence for the governments') income and for relative financial asset prices. If public institutional investors face such large-scale investment issues, some normative theory of their investment behaviour is obviously of interest. How far would such a theory deviate from a normative theory of investment for typical private large-scale institutional investors, such as pension funds, endowment funds, insurance companies, or mutual funds? Can we rationalize with such a theory what we observe today as central bank investment behaviour? Or would we end concluding like Summers (2007), who compares central bank investment performance with the typical investment performance of pension and endowment funds, that central banks waste considerable public money with an overly restrictive investment approach?

In practice, central bank risk management is extensively using, as it should, risk management methodologies and tools developed and applied by the *private* financial industry. Those tools will be described in more detail in the following chapters of the book. While public institutions are in this respect not fundamentally different from other institutional investors, important specificities remain, due to public institutions' policy mandate, organizational structure or financial asset types held. This is what justifies discussing all these tasks in detail in this book on central bank and other public institutions' risk management, instead of simply referring to general risk management literature. The present chapter focuses more on the main *idiosyncratic* features of public institutions in the area of investment and

risk management, which do not relate so much to the set of risk management tools to be applied, but more on how to integrate them into one consistent framework reflecting the overall constraints and preferences of, for example, central banks, and how to correspondingly set the basic key parameters of the public institution's risk management and investment frameworks.

The rest of this chapter is organized as follows: Section 2 reviews in more detail the specificities of public investors in general, which are likely to be relevant for their optimal risk management and investment policies. Section 3 turns to the specific case of central banks, being by far the largest type of public investors. It explains how the different central bank policy tasks on the one side have made such large investors out of central banks, and on the other side may constrain the central bank in its investment decisions. Sections 4 and 5 look each at one specific key question faced by public investors: first, how much should public investors diversify their assets, and second, how actively should they manage them. Sections 6 and 7 are devoted again more specifically to central banks, namely by looking more closely at what non-alienable risk factors are present in central bank balance sheets, and at the role of central bank capital, respectively. Section 6, as Section 3, reviews one by one the key central bank policy tasks, but in this case to analyse their role as major non-alienable risk factors for integrated central bank risk management. Also on the basis of Sections 6 and 7, Section 8 turns to *integrated* financial risk management of public institutions, which is as much the holy grail of risk management for them as it is for private financial institutions. Section 9 draws conclusions.

2. Public institutions' specificities as investors

Public institutions are specific as financial investors as they operate under unique policy mandates and are subject to constraints which do not exist for private institutional investors. These specificities will have implication for optimal investment behaviour. The following specificities 1) to 5) are relevant for all public investors, while 6) to 10) only affect central banks.

1) Public institutions may appear to be, relative to some private institutional investors (like an insurance, or an endowment fund), subject to some specific constraints: (i) Less organizational flexibility, including more complex and therefore more costly decision-making procedures. This may argue against 'decision-intensive' investment styles; (ii) Decision

makers less specialized on investment. For instance central bank board members are often macroeconomists or lawyers, and come more rarely from the investment or risk management side; (iii) Higher accountability and transparency requirements, possibly arguing against investment approaches that are by nature less transparent, such as active portfolio management; (iv) Less leeway in the selection and compensation of portfolio managers due to rules governing the employment of public servants. This may argue against giving leeway to public investors' portfolio managers, as compared to less constrained institutional investors. There are certainly good reasons for these organizational specificities of public institutions. They could in general imply, everything else being equal, a certain competitive disadvantage of central banks in active portfolio management or in diversification into less standard asset classes, relative to private players.

2) **Public institutions being part of the consolidated state sector**. It could be argued that when investing into domestic financial assets, public institutions should have a preference for Government securities as they are part of the state sector, and as the state sector should not lengthen unnecessarily its consolidated balance sheet (i.e. the consolidated state balance sheet should be 'lean'). A lean state sector may be defended on the basis of the general argument that the state should concentrate on its core business, and avoid anything else, since it is likely to be uncompetitive relative to private players (which are 'fitter' as they survive free market competition). The Fed may be viewed as a central bank following the 'lean consolidated state sector' approach most closely; as more than 90 per cent of its assets are domestic Government bonds held outright (see Federal Reserve Bank of New York 2007, 11). Thus, if one consolidates the US federal Government and the Federal Reserve System, a large part of the Fed balance sheet can be netted off.

3) **Public institutions have a very special owner: the Government, and therefore, indirectly, the people (or 'the taxpayer')**. When discussing how a specific institutional investor should invest, it is natural to first look at who 'owns' the institutional investor or, more generally, who owns the returns on the assets that are managed. One tends to describe (or to explain) the preferences of investors with (i) an investment horizon, (ii) relative risk–return preferences, expressed in some functional form, (iii) possibly some non-alienable assets or liabilities (for individuals, this would for instance be human capital), which exhibit specific correlations with financial assets, and thereby determine the optimal asset allocation. If one would

view the central bank in its role as investor as a pure agent of the Government or of the people, one needs to look in more detail to these three characteristics of its owner. The opposite approach is to view a public institution as a subject on its own, and to see payments to its owners (to which it is obliged through its statutes) as 'lost' money from its perspective. Under this approach, the three dimensions (i)–(iii) of preferences above need to be derived taking directly the perspective of the public institution.

4) Public institutions do not have the task to maximize their income. Instead, for instance the ECB has, beyond its primary task to conduct monetary policy, the aim to contribute *to an efficient allocation of resources*, i.e. it should have social welfare in mind. According to article 2 of the ESCB/ECB Statute: 'The ESCB shall act in accordance with the principle of an open market economy with free competition, favouring an efficient allocation of resources. . .'. The question thus arises in how far certain investment approaches, such as e.g. active portfolio management, are socially efficient. As Hirshleifer (1971) had demonstrated, there is no general insurance that private and social returns are equal in the case of information producing activities. Especially in the case of what he calls 'foreknowledge', it seems likely *that private returns of information producing activities tend to exceed social returns*, such that at the margin, investment into such information would tend to be detrimental to social welfare (i.e. to an efficient allocation of resources). In his words:

> The key factor. . .is the distributive significance of foreknowledge. When private information fails to lead to improved productive alignments (as must necessarily be the case in a world of pure exchange. . .), it is evident that the individual's source of gain can only be at the expense of his fellows. But even where information is disseminated and does lead to improved productive commitments, the distributive transfer gain will surely be far greater than the relatively minor productive gain the individual might reap from the redirection of his own real investment commitments. (Hirshleifer 1971, 567)

One could thus argue that it is questionable that an institution, which according to its statute should care about social welfare, engages into active portfolio management. On the other side, it could be felt that this argument applies to a lesser extent to foreign reserves, since a central bank should probably always care more about the welfare of *its own* country than about the one of others, such that egoistic profit maximization in the case of foreign reserves would be legitimate. Also beyond the issue of active management, the question is to be raised whether what is rational from the

perspective of a private, selfish investor would be economically (or 'socially') efficient if applied by the central bank. Unless one has concrete indications of the contrary, public institutions should probably assume that this is the case, i.e. that by adopting state-of-the-art investment and risk management techniques from the financial industry; they also contribute to the social efficiency of their investments.

5) Public institutions and reputation risk. Reputation risk may be defined as risks arising from negative public opinion for the P&L of an institution or more generally for the ability to conduct relevant tasks. This risk may be related to the risks of litigation and loss of independence. It is also called sometimes 'headline' risk as events damaging the reputation of a public institution are typically taken up by the media. Reputation risk is often linked to financial losses (i.e. in case of losses due to the failure of a counterparty), but not necessarily. For instance, it may be deemed a 'scandal' in itself that a central bank invests into some issuer, be it public or private, which is judged not to adhere to ethical standards. Or it could be considered that the central bank should not invest into some 'speculative' derivatives, although these derivatives are in fact used for hedging, what the press, the government or the public however may not understand. All investors may be subject to reputation risk, but clearly to a varying degree. Central banks' rather-developed sensitivity for reputation risk may stem from the following three factors:

(i) Their need for credibility for achieving their policy tasks, such as maintaining price stability. Credibility is not supported by being perceived as unethical or amateurish.

(ii) Central banks tend to 'preach' to the rest of the world what is right and wrong. For instance, they often criticize the spending behaviour and lack of reform policies of Governments. Or, as banking supervisors, they impose high governance standards on banks, and search for weaknesses of banks to intervene against them. Again, such roles do not appear compatible with own weaknesses, which again is a credibility issue.

(iii) Central banks worry about preserving their independence. Independence is a privileged status, and it is obviously endangered if the central bank shows weaknesses which could help the adversaries of central bank independence (and those which were criticized or lectured by it) to argue that 'these guys need to be controlled more closely by democratically elected bodies'.

A classical example for central bank headline risk is the attention the small exposure of Banca d'Italia to LTCM got in 1998, including a need for the

Governor to justify the Bank in the Italian Parliament. Reputation risk may depend first of all on whether a task is implied by the statutes of a public investor. If for instance holding foreign reserves is a duty of a central bank, then associated financial risks should imply little reputation risk. The more remote an activity is to the core tasks assigned to the public investor, the higher the danger of getting questions like: 'How could you lose public money in this activity and why did you at all undertake it as you have not been asked to do so?' If taking market or credit risk for the sake of increasing income is not an explicit mandate of a public institution, then market or credit risk will have a natural correlation to reputation risk.

Reputation risk is obviously closely linked to transparency, and maybe transparency is the best way to reduce reputation risk. What has been made public and explained truthfully to the public can less be reproached to the central bank in case of non-favourable outcomes – in particular if no criticism was voiced *ex ante*. Central banks have gone a long way in terms of transparency over the last decades, not only in terms of monetary policy (e.g. transparency on their methodology and decision making), but also in the area of central bank investments. For instance the ECB has published in April 2006 an article in its Monthly bulletin revealing a series of key parameters of its investment approach (ECB 2006a, 75–86). Principles of central bank transparency in foreign reserves management are discussed in section 2 of IMF (2004).

6) **Central banks are normally equipped with large implicit economic capital through their franchise to issue banknotes**. This could be seen to imply that they can take considerable risks in their investments, and harvest the associated higher expected returns. At least for a majority of central banks, the implicit capital is indeed considerable, which is discussed in more detail in Section 7. Still, for some other central banks, financial buffers may be less extensive. For instance, central banks which are asked to purchase substantial amounts of foreign reserves to avoid revaluation of their currency may be in a potentially loss-making situation, in particular if, in addition: (i) the demand for banknotes in the country is relatively limited; (ii) domestic interest rates are higher than foreign rates; (iii) their own currency is under revaluation pressure, which would imply accounting losses.

7) **Central bank independence (relevant mainly for *domestic* financial assets)**. The need for central bank independence may be viewed to be relevant in this context as implying that the central bank should stay out from investing into securities or other assets issued by its own countries' Government. In particular World War I taught a lesson in this respect to

e.g. the US, the UK, and more than to anyone else, to Germany. Under Government pressure, the central banks purchased during the war massive amounts of Government paper and kept interest rates artificially low. It has been an established doctrine for a long time that the excessive purchase of Government paper by the central bank is a sign of, or leads to, a lack of central bank independence. For instance article 21.1 of the ECB/ESCB Statutes reflects this doctrine by prohibiting the direct purchase of public debt instruments by the ECB or by NCBs.

8) Central banks have insider information on the evolution of short-term rates, at least in their own currency, and thus on the yield curve in general. One may argue that insider information should not be used for ethical or for other reasons, and that therefore certain types of investment positions (in particular yield curve and duration positions in domestic fixed-income assets) should not be taken by central bank portfolio managers. As a possible alternative, 'Chinese walls' or other devices can be established around active managers of domestic portfolios in the central bank. For foreign exchange assets, the argument holds to a lesser extent.

9) Central banks may have special reasons to develop market intelligence, since they need to implement monetary policy in an efficient way, and need to stand ready to operate as lender of last resort. Especially the latter requires an in-depth knowledge of financial markets and of all financial instruments. While some forms of market intelligence may be developed in the context of basic risk-free debt instruments, a more advanced and broader understanding of financial markets may depend on diversifying into more exotic asset classes (e.g. MBSs, ABSs, CDOs, equity, hedge funds) or on using derivatives (like futures, swaps, options, or CDSs). Also active portfolio management may be perceived as a way to understand best the logic of the marketplace, as it might be argued that only with active management do portfolio managers have strong incentives to understand all details of financial markets. For instance the Reserve Bank of New Zealand has stated this doctrine, motivating active portfolio management openly (taken from the IMF 2005, statement 773 – see also the statement by the Bank of Israel, IMF 2005, statement 663):

773. The Bank actively manages foreign reserves. It does so because it believes that active management: generates positive returns (in excess of compensation for risk and of active management overheads) and so reduce the costs of holding reserves; and encourages the dealers to actively participate in a wider range of instruments and markets than would otherwise be the case and so improves the Bank's market intelligence and contacts, knowledge of market practices, and foreign exchange intervention and risk management skills. The skills and experience gained

from reserves management have been of value to the Bank in the context of its other roles too. For instance, foreign reserves dealers were able to provide valuable input when the Bank, in the context of its financial system oversight responsibilities, was managing the sale of a derivatives portfolio of a failed financial institution. It is not possible to be precise about how much added-value is obtained from active management but, in time of crises, extensive market knowledge, contacts and experience become invaluable.

10) At least some central banks tend to be amongst the exceptionally big investors. The most striking examples are the Asian central banks and in particular China and Japan with reserves, mostly in USD, at or beyond 1 trillion USD. The status as *big* investor has two important consequences. First, such central banks should probably go further than others in **diversifying** their investment portfolio. In the CAPM (Capital Asset Pricing Model), all investors should hold a widely diversified market portfolio, but in reality, transactions and information costs of many kinds are making such full diversification inefficient. Participation in a diversified fund can reduce these costs, but will not eliminate them. The easiest way to model these costs preventing full diversification is to assume fixed set-up costs per asset type, which may be viewed as the costs for the front, back and middle office to understand the asset type sufficiently and to prepare for the integration and handling of associated transactions. These fixed set-up costs will be lower for some and higher for other asset types. Under such assumptions, it is clear why smaller investors will end up being less diversified. Set-up costs can be economized to some extent through *outsourcing* or through purchasing investment vehicles like funds. Also, some important forms of diversification, like e.g. into an equity indices, may require relatively low set-up costs, and hesitations of central banks (large or small) with their regard may be due to other reasons. Second, large central banks with a substantial weight in some markets (e.g. US Treasuries) may influence relative prices in these markets, in particular when doing large transactions. This may potentially worsen their returns, and implies the need to smooth transactions over time, and, again, to diversify. Also it increases liquidity risks, i.e. the risks that the quick liquidation of relevant positions is only possible at a discount.

3. How policy tasks have made central banks large-scale investors

The starting point in analysing the specificities of central banks as investors is clearly the question why central banks are at all facing 'investment' issues.

Return-oriented investment is clearly not amongst the standard policy tasks of central banks. To be a bit more specific on the core central bank tasks and how they have made large-scale investors out of central banks, consider the tasks specifically prescribed for the ECB, which are quite standard for central banks with statutes defined in the last decade or so. According to the Treaty establishing the European Community (article 105.2), the basic tasks of the ECB are: (i) the definition and implementation of monetary policy for the euro area; (ii) the conduct of foreign exchange operations; (iii) the holding and management of the official foreign reserves of the euro area countries (portfolio management); (iv) the promotion of the smooth operation of payment systems. Further tasks according to the ECB/ESCB Statutes relevant in this context are: (v) banknotes – the ECB has the exclusive right to authorize the issuance of banknotes within the euro area (article 16 of the ECB/ESCB Statute); (vi) financial stability and supervision – the Eurosystem contributes to the smooth conduct of policies pursued by the authorities in charge related to the prudential supervision of credit institutions and the stability of the financial system (see Article 25 of the ECB/ESCB Statutes). Finally, article 2 ('Objectives') of the ECB/ESCB Statutes also prescribes that: (vii) 'The ESCB shall act in accordance with the principle of an open market economy with free competition, favouring an efficient allocation of resources. The rest of this section explains how such tasks made "investors" out of today's central banks.'

3.1 Banknotes issuing and payment systems

As long as there is demand for banknotes, and the central bank has an issuance monopoly, banknotes will constitute an important and unique unremunerated liability in the balance sheets of central banks. As any other liability, banknotes need to be counterbalanced by some assets. Unless there are specific constraints on the asset composition derived from other policy tasks, the existence of banknotes in itself thus creates the need for investment decisions by the central bank. Although academic and even central bank visionaries have forecast the end of banknotes for a long time, the trend growth of banknotes of the large currencies (in particular USD and EUR) has been even above the nominal growth rate of the respective economies. For example, at end June 2007, euro banknotes in circulation stood at EUR 633 billion, more or less the same as USD banknotes. In the past, there were examples of central bank *payment systems* creating unremunerated liabilities of the central bank of considerable size and of almost

similar magnitude as banknotes (e.g. the Reichsbank in 1900, see table 2.2 in Bindseil (2004, 52)). Today, however, payment systems tend to be so efficient as to create very little unremunerated central bank liabilities. In so far, they have a negligible impact on central bank balance sheets.

3.2 Monetary policy implementation

In short, central bank monetary policy consists in setting short-term money interest rates such that price stability is maintained over time. Short-term interest rates are controlled by steering the scarcity of deposits of banks with the central bank, as the short-term money interest rate is essentially anchored in the interbank lease rate for such deposits. The *supply* of deposits can be influenced by the central bank by injecting or absorbing deposits, namely by purchasing or selling securities, or by lending or absorbing funds through so-called 'reverse operations'. The *demand* for deposits can also be influenced by the central bank, notably by imposing reserve requirements on banks. What matters in the present context is that the steering of short-term rates eventually consists in manipulating the scarcity of deposits 'at the margin', i.e. the lease price of deposits is set, as any other price, by the *marginal* demand and supply. This means that the central bank is constrained in its asset management decisions from the perspective of monetary policy only *marginally* in the sense that the assets that are used to steer the scarcity of deposits at the margin need to be suitable to do so, *while the choices regarding the entire rest of the assets remain unaffected* (for a survey of monetary policy implementation explaining the impact on the asset side of the balance sheet, see e.g. Bindseil 2004, chapters 2 and 3). It is noteworthy that in the more distant past, there had been the view that the entire asset composition of central banks (so not only the one at the margin) does matter for its ability to control inflation: this was the case under the famous real bills doctrine, according to which 'real' (in contrast to 'financial') trade bills would be good, i.e. non-inflationary assets (see Bindseil 2004, 107 and the literature mentioned there). Recognizing that monetary policy is implemented only 'at the margin', the central bank can segregate a large part of its domestic financial assets from the monetary policy operations to consider them domestic financial investment portfolios. Monetary policy portfolios typically consist of short-term reverse repo operations which the central bank conducts through tender procedures in the market. Financial risks of those tend to be very limited: credit risk is mitigated through collateralization (see Part II of the book), while market

Table 1.1 Foreign reserves (and domestic financial asset of G3 central banks) in December 2007

	USD billion
Top five reserve holders:	
• China	1,528
• Japan	954
• Russia	446
• Taiwan	270
• India	268
Euro Area	235
Total foreign reserves of central banks	6,459
Sovereign wealth funds	≈3,000
Total official reserves	≈9,500
Memo: Domestic financial assets of three currency areas (approximate)	
US Fed (end 2006)	800
Eurosystem (euro area central banks, end 2007)	1,150
Bank of Japan (end 2007)	920
Grand total of financial assets in this table	*12,400*

Sources: IMF; JPMorgan 'New trends in reserve management – Central bank survey', February 2008; for domestic financial assets: central bank websites.

risks appear negligible in the sense that the operations tend to have short-term maturity (mainly up to three months).

3.3 Foreign exchange policies and reserves

The tasks to implement foreign exchange rate policies and to hold foreign reserves have over the last decade led to an unprecedented accumulation of foreign reserves of which the likelihood of use in foreign exchange interventions to support the domestic currency is negligible. The growth in reserves reflects essentially increased oil and other raw material revenues that the relevant countries capitalized in foreign currency, and the trade surpluses of Asian economies that the relevant countries did not want to be reflected in an appreciation of their respective currencies. According to JPMorgan, reserve accumulation would have reached a new record in 2007 with an annual increase of 26 per cent. Table 1.1 shows an overview of official reserves figures as of December 2007.

Accordingly, global *foreign* reserves stood at end 2007 at USD 6.5 trillion, of which around USD 4 trillion were owned by Asian central banks.

Sovereign wealth funds, which are typically split-ups of excess central bank reserves, constituted around EUR 3 trillion. Defining 'excess reserves' as foreign reserves which would not be needed to cover all foreign debt coming due within one year, Summers (2007) notes that excess reserves of 121 developing countries sum up to USD 2 trillion, or 19 per cent of their combined GDP. China's excess reserves would be 32 per cent of its GDP and for instance for Malaysia this figure even stands at 49 per cent and at 125 per cent for Libya. Excess reserves could be regarded as those reserves for which *central banks only face an investment problem, and have no policy constraint* (except, maybe, the foreign currency denomination). In fact, three cases of central banks may have to be differentiated with regard to the origin and implied policy constraints of foreign reserves. First, the case of a large area being the '$n+1$' country not caring so much about foreign exchange rates and thus not needing foreign reserves. The US (and maybe to a lesser extent the euro area) falls into this category, and the Fed will therefore hold very little or no foreign reserves for policy reasons. In this case, still, the central bank may hold foreign reserves for pure investment reasons. However, this typically adds substantial market risk, without allowing to improve expected returns, which would therefore rarely be done by such a central bank. Second, central banks may want to hold foreign reserves as ammunition for foreign reserve intervention in case of devaluation pressures on their own currency in foreign exchange markets. This has obviously consequences on the currency denomination of assets, and on required liquidity characteristics of assets. Apart from the currency and liquidity implications of these policy objectives, the central bank can however still make asset choices affecting risk and return, i.e. some leeway for investment decisions remains. Most Latin American countries typically fall under this category. Third, there are central banks which would like to avoid appreciation of their currency, and thereby purchase systematically over time foreign assets, such as many Asian central banks have done it in an unprecedented way for several years. Such reserve accumulation puts little constraint in terms of liquidity on the foreign assets (as there is only a marginal likelihood of the need to sell the reserves under time pressure), but can have, due to the amounts involved, pervasive consequences for the overall length of and risks in the central bank balance sheet. To take an example: the People's Bank of China reached at end 2007 a level of foreign reserves amounting to USD 1.5 trillion. A 10 per cent appreciation of the Yuan would thus mean losses to the central bank of USD 150 billion, which is much more than the capital of any central bank of the world. These risks in themselves are however obviously not constraining

Table 1.2 Different reasons for holding foreign exchange reserves – importance attributed by reserve managers according to a JPMorgan survey in April 2007

	Very important	Important	Somewhat important	Total
Conduct FX policies (interventions)	44%	23%	23%	91%
Crisis insurance	37%	28%	7%	72%
Serve external debt obligations	23%	12%	19%	53%
Ensure import coverage	21%	9%	19%	49%
Support country's credit standing	12%	14%	23%	49%
Build national wealth (future generations)	12%	19%	28%	59%
Other	5%	0%	2%	7%

Source: JPMorgan 'New trends in reserve management – Central bank survey', February 2008.

investment, and thus central banks of this type face very important investment choices due to the mere size of their assets.[1]

Table 1.2 overviews how central banks perceive the relevance of different motives to hold reserves as obtained by JPMorgan in a survey conducted in 2007.[2]

The existence of large unhedged foreign exchange reserves explains the rather peculiar relative weights of different risk types in the case of central banks. For large universal banks (like e.g. Citigroup, JPMorgan Chase, or Deutsche Bank), credit risk tends to clearly outweigh market risk. This may be explained by the fact that market risk can be diversified away and be hedged to a considerable extent, while credit risks eventually needs to be assumed to a considerable extent by banks (even if some diversification is possible as well, and credit risk can be transferred partially through derivatives like credit default swaps). For central banks, the opposite holds: market risks tend to outweigh very considerably credit risks (but see Chapter 3 reflecting that

[1] Indeed, the size of central bank assets is in those cases not constrained by the size of banknotes in circulation and reserve requirements. Such central banks typically have to absorb domestic excess liquidity, implying that the foreign reserves are then countered on the asset side by the sum of banknotes and liquidity absorbing domestic ('monetary policy') operations.

[2] I wish to thank JPMorgan for allowing me to use the results from their 2007 central bank survey for this chapter. The survey results were compiled from participants at the JPMorgan central bank seminar in April 2007. The responses are those of the individuals who participated in the seminar and not necessarily those of the institutions they represented. Overall, forty-four reserve managers replied to the survey. The total value of reserves managed by the respondents to this survey was USD 4,828 billion or roughly 90 per cent of global official reserves as of December 2006. The sample had a balanced mix of central banks from industrialized and emerging market economies from all parts of the world, but was biased toward central banks with large reserve holdings (the average size of reserve holdings was USD 110 billion versus roughly USD 27 billion for all central banks in the world).

Table 1.3 Risk quantification and economic capital, in billions of EUR, as at end 2005

	Deutsche Bank	ECB, 99.9% VaR, one-year horizon
Credit risk	7.1	≈0.2
Market risk (incl. private equity risk)	3.0	≈10.0
Of which: Interest rate risk	50%	5%
Equity price risk	30%	–
Exchange rate risk	12%	95%
Commodity	9%	
Operational/business risk	2.7	?
Total economic capital need	12.3	?

Sources: Deutsche Bank (2005); ECB's own calculations

credit risk may have become more important for central banks over the last years). When decomposing further market risks taken by for instance the ECB, such as done in the table below, the exceptionally high share of market risks can be traced back to foreign exchange rate and commodity risks (the latter relating to gold). Table 1.3 also reveals that in contrast to that, private banks, as in this case Deutsche Bank, hold only to a very low extent foreign exchange rate risk. Instead, interest rate and, to a lesser extent, equity price risks dominate.

One may also observe that Deutsche Bank's main risks are risks which are remunerated (for which the holder earns a risk premium), while the very dominating risk of the ECB, exchange rate risk, is a risk for which no premium is earned. It derives from one of the main policy tasks of a central bank, namely to hold foreign reserves for intervention purposes. From the naïve perspective of a private bank risk manager, central bank risk taking could therefore appear somewhat schizophrenic: holding huge non-remunerated risks, but being highly risk averse on remunerated risks. While the former is easily defended by a policy mandate and the large implicit financial buffers of central banks, the latter may be more debatable.

3.4 Financial stability functions

Under normal circumstances, central banks tend to be more risk averse in their asset choices than the normal investor (e.g. typically, central bank fixed-income portfolios have a modified duration below the one of the

market portfolio).[3] However, the role of the central bank to contribute to financial stability may imply a specific temporary shift of risk aversion, namely to take in circumstances of financial instability particular risks that no private player is willing or able to take, thereby rescuing the liquidity of some bank(s) and avoiding domino effects in the financial sector that would create substantial economic damages and welfare losses for society. Such tasks obviously have implications on the assets held by the central bank at least temporarily, and thereby in some sense reduce the leeway of the central bank as investor. While in principle, financial crisis management operations should still be done prudently with a full awareness of associated risks and with the central bank being independent from pressures by the Government (unless the central bank has a clear loss transfer agreement with the Government), there are numerous cases, in particular from outside the OECD (Organization for Economic Co-operation and Development), in which emergency assistance to banks or even non-financial industries has created large non-performing central bank assets and made effective central bank capital negative. It may also be worth mentioning that in the 1930s the Bank of England was a large share-holder of industrial companies with the aim to avoid, through support and restructuring, industry failures.[4] This is another example on how policy tasks may reduce the leeway of central banks to act as an investor. Still, it is to be admitted that these are policy constraints that today very rarely matter for central bank investments in industrialized countries. Financial crisis management of central banks is discussed in more detail from a risk management perspective in Chapter 11.

4. Optimal degree of diversification of public institutions' financial assets

Strategic (or equilibrium) approaches, in contrast to active management approaches, start from the assumption that markets are efficient and that the investor has no private information with regard to wrongly priced

[3] In principle, risk aversion of an investor should imply keeping limited the duration mismatch between assets and liabilities. In so far, a short duration of a central bank investment portfolio should reflect at the same time a view of the central bank that the liabilities associated with the assets have a short duration, or are not relevant.

[4] See Sayers (1976, vol. 1, chapter 14; 1976, vol. 2, chapter 20, section G). Sayers (1976, vol. 1, 314) writes: 'The intrusion of the Bank into problems of industrial organisation is one of the oddest episodes in its history: entirely out of character with all previous development of the Bank. . .eventually becoming one of the most characteristic activities of the Bank in the inter-war decades. It resulted from no grand design of policy, nor was the Bank dragged unwillingly into it.'

assets. 'Passive portfolio management' may correspondingly be defined as portfolio management being based on the idea that asset prices are broadly adequate, such that the main aim is to engineer an appropriate diversification and risk–return combination taking into account risk preferences and the non-alienable risk–return profile of the investor. The capital asset pricing model (CAPM) is normally the theoretical underpinning of such an approach: accordingly, investors should hold a combination of the risk-free asset and the 'market portfolio', which is simply the portfolio of outstanding financial assets weighted according to their actual weights in the market. What does this market portfolio look like? According to a JPMorgan estimate, 60% of outstanding global financial assets were in September 2006 equity, followed by fixed income with 35% and 'alternatives' of 5%. Of course, the relative share of equity vis-à-vis fixed income fluctuates over time with market prices, i.e. in particular with equity prices. According to Bandourian and Winkelmann (2003, 100), the share of equity recently reached a minimum in October 1992 with 47% and a maximum of 63% in March 2000. According to the *Lehman Global aggregate index*, as of 30 September 2006, Governments bonds constitute one-half of outstanding global fixed-income securities, followed by MBSs (18%), Corporates (17%) and Agencies (6%). It is also interesting to look at the credit quality distribution globally. According to the Lehman global aggregate index, 57% of outstanding fixed-income securities would be AAA rated, 26% would be AA, 13% A and only 4% would be BBB. Moreover, one can present the split-ups according to asset classes, sectors and credit quality by currency (e.g., USD, EUR and JPY). For the USD for instance, 34% of fixed-income securities are mortgages, 26% Treasuries, 20% Corporates, and 11% Agencies (according to the sector allocation according to the Lehman US aggregate index, September 30, 2006).

Finally, it is interesting to look at modified duration of the market portfolio in different currencies and across different sectors of the fixed-income universe. The related information, which is based on securities data in Bloomberg, is summarized in Table 1.4.[5] As one can see, the different currency and sectoral components of the market portfolio for which duration figures could be estimated relatively easily are all in the range from three to five years, which appears to reflect preferences of investors and debtors.

Obviously, public institutions do *not* come close to holding the market portfolio (combined with the risk-free asset), but this observation also holds

[5] I wish to thank Hervé Bourquin for this analysis.

Tabel 1.4 Modified duration of fixed-income market portfolios (as far as relevant)

	USD	EUR	JPY
Central government	3.0	4.2	4.0
Agencies	3.6	5.1	3.8
Corporates	4.5	3.3	3.5

Sources: Bloomberg; ECB's own calculations.

for other investors. Thus, one may first want to ask why many investors tend to diversify so little, and therefore often seem to diversify too little into credit risk. Obviously, the assumptions underlying the CAPM are not adequate, whereby the following five assumptions appear most relevant in the context of the optimal degree of diversification for public investors:

(1) **No private information**. There will always be private information in financial markets and as a consequence, in microeconomic terms, a non-trivial price-discovery process. The existence of private information is implied by the need to provide incentives for the production of information (Grossman and Stiglitz, 1980). If private information is in a market, and an investor belongs to the uninformed market participants (i.e. acts like a 'noise trader'), then he is likely to pay a price to the informed trader, e.g. in the form of a bid–ask spread as modelled by Treynor (1987), or Glosten and Milgrom (1985). This is a powerful argument to stay away from markets about which one knows little. If public institutions were comparably less efficient in decision making than private institutional investors, and had less leeway in remunerating analysts and portfolio managers, one could argue generally against competitiveness of public institutions to operate in markets with a big potential for private information.

(2) **No transaction, fixed set-up and maintenance costs**. Transaction costs take at least the following three forms: costs of purchasing or selling assets (the bid–ask spread being one part of those, the own costs of handling the deal the other), fixed one-off set-up costs for being able to understand and trade an instrument type, and fixed regular costs, e.g. costs to maintain the necessary systems and knowledge. Fixed costs arise in the front, middle and back office since the relevant human capital and IT systems need to be made available. Fixed set-up costs imply that investors will stay completely out of certain asset classes, despite the law of risk management that adding small uncorrelated risks

does not increase total risk taking at all. Fixed set-up costs also imply that the larger a portfolio, the more diversification is optimal. Portfolio optimization with fixed costs can be done in the 'brute force' way by just running a normal optimization for the different combinations of asset classes (an asset class being defined as a set of assets for which one set-up investment has to be done), shifting then the efficient frontiers by the fixed set-up costs to the left (considering the size of the portfolio), choosing the optimal portfolio, and then selecting the best amongst these optimal portfolios. While this implies that central banks with large investment portfolios are more diversified in their investment than those with smaller portfolio size, it is interesting to observe that this does not explain everything. In the Eurosystem for instance, only two NCBs have diversified their USD assets into corporate bonds. Large central bank investors may also be forced by the size of their reserves to diversify to avoid an impact of their purchases on asset prices (e.g. China with its reserves of over USD 1 billion).

(3) **No 'non-alienable' risks**. Each investor is likely to have some 'non-alienable' risk–return factors on his balance sheet. In the case of human beings, a major such risk factor is normally human capital (some have estimated human capital to constitute more than 90 per cent of wealth in the US, see Bandourian and Winkelmann (2003, 102)). In the case of central banks, the risk and returns resulting from non-alienable policy tasks are discussed further in Section 6.

(4) **No liquidity risk**. If investors have liquidity needs, because with a certain probability they need to liquidate assets, they will possibly deviate from the market portfolio in the sense of underweighting illiquid assets. This may be very relevant for e.g. central banks holding an intervention portfolio likely to be used.

(5) **No reputation risk**. Reputation risk may also be classified as non-alienable risk factor being implied by policy tasks.

When considering a diversification into a new asset category, a public institution should thus not only make an analysis of the shift in the feasible frontier that can be achieved by adding a new asset class in a portfolio optimizer. It is a rather unsurprising outcome that the frontier will shift to the left when adding a new asset class, but concluding from this that the public institution should invest into the asset class would mean basing decisions on a tautology. The above list of factors implying a divergence from the market portfolio for every investor, and public institutions in particular, should be analysed one by one for any envisaged diversification

Table 1.5 Asset classes used by central banks in their foreign reserves management

Asset class	Estimated share of asset class in total reserves
Central Government bonds	73%
US Agencies debt	18%
Corporate bonds	3%
ABS/MBS	5%
Deposits	20%
Gold	10%

Source: Wooldridge (2006).

project, and an attempt should be made to quantify each of the factors before drawing an overall conclusion.

Finally, the 'positive externalities' argument in favour of active portfolio management by central bank (see the excerpts in IMF 2005, section 3.5), could also be applied for diversification of central bank portfolios. If a central bank is invested into a financial instrument itself, it is more likely that it will have really understood it, and thus it will understand its possible role for monetary policy implementation or financial stability. In so far, the positive externality argument would drag the central bank's investment portfolio towards the market portfolio.

Tables 1.5 and 1.6 provide a survey of the degree of diversification across instrument classes that have been achieved by central banks in foreign reserves. Table 1.5 provides estimates of the share of different asset classes, according to Wooldridge (2006), who suggests that monetary authorities have since the 1970s gradually diversified into higher-yielding, higher-risk instruments, but nevertheless reserves are still invested mostly in very liquid assets, with limited credit risk.

Table 1.6 provides the results of the JPMorgan reserve managers' survey on asset classes the central banks is currently allowed to use for planning.

Obviously, all or almost all central banks invest their foreign reserves into sovereign bonds in the relevant currency. Also, a large majority invests in (AAA rated) US Agency debt and Supranational bonds, whereby the weight of the latter in the market portfolio is very small. All other major types of bonds (corporate, MBS/ABS, Pfandbriefe, bank bonds) are eligible for up to around 50 per cent of central banks. Outside fixed-income securities, deposits are of course used by most central banks, while equity has been made eligible only by around 10 per cent of central banks.

Table 1.6 Asset classes currently allowed or planned to be allowed according to a JPMorgan survey conducted amongst reserve managers in April 2007

	Approved	Planned
Gold	91%	0%
Deposits	100%	0%
US Treasuries	98%	0%
Euro govies	98%	0%
Japan and other OECD govies	77%	5%
US Agencies	88%	7%
TIPs	37%	9%
Supra/Sovereigns	98%	0%
Covered bonds	51%	12%
ABS/MBS	42%	16%
High-grade credit	35%	9%
High-yield credit	2%	2%
Emerging markets credits	12%	7%
Equities	9%	5%
Non-gold commodities	5%	2%
Hedge funds	2%	2%
Private equity	2%	2%
Real estate	5%	2%
Other	5%	0%

Source: JPMorgan 'New trends in reserve management – Central bank survey', February 2008.

According to Wooldridge (2006), the share of deposits is distributed rather heterogeneously across central banks. For instance India would have held 76% of its reserves in the form of deposits in 2006, and Russia 69%. Gold reserves still constituted 60% of foreign reserves in 1980. Currency composition of international foreign reserves would have been in 2006: around 65% in USD, 25% in EUR, JPY, GBP, and all the rest around 3% each.

Finally, it is interesting to see which derivatives are used by central banks in foreign reserve management, even if derivatives are by nature not in themselves part of the market portfolio. Derivatives may be used for efficiently replicating a benchmark (i.e. within a passive portfolio management context), for hedging, or for active position taking. Table 1.7 provides results on derivative use by central banks from the 2007 JPMorgan survey.

Table 1.7 Derivatives currently allowed or planned to be allowed according to a JPMorgan survey conducted amongst thirty-eight reserve managers in April 2007

	Approved	Planned
Gold swaps	32%	3%
Gold options	24%	3%
FX forwards	76%	3%
FX swaps	63%	5%
FX futures	18%	3%
FX options	26%	16%
Interest rate futures	61%	13%
Interest rate swaps	53%	16%
Interest rate options	18%	18%
Credit derivatives	8%	11%
Equity derivatives	8%	0%
Non-gold commodity derivatives	5%	0%

Source: JPMorgan 'New trends in reserve management – Central bank survey', February 2008.

5. How actively should public institutions manage their financial assets?

5.1 The general usefulness and 'industrial organization' of active portfolio management

Whether active management pays or not, being often associated with the question of whether financial markets are efficient or not has been the topic of extensive academic debate (for surveys of the topic see e.g. Grinold and Kahn 2000, chapter 20; Cochrane 2001, 389; see also e.g. Ippolito 1989; Berk and Green 2002; Engström 2004). Elton *et al.* (2003, 680) remain as agnostic to say that 'the case for passive versus active management certainly will not be settled during the life of the present edition of the book, if ever'. Obviously, active management creates extra costs, namely: the costs of the analysis on which the private views and forecasts are based; the cost of diversifiable risk – active portfolios, by their nature, have often more diversifiable risk than an index fund; higher transaction costs being due to a higher turn-over of securities; higher governance costs (need of an additional investment committee, etc.). It has been argued that these extra costs will not be easily recovered if the efficiency of the market is sufficiently high. The ongoing

debate on the usefulness of active management may in fact appear surprising, since already Grossman and Stiglitz had shown convincingly in their seminal paper of 1980 that the question of the *general* usefulness of active management is misplaced. Instead, active management needs to be part of a competitive equilibrium itself:

> If competitive equilibrium is defined as a situation in which prices are such that all arbitrage profits are eliminated, is it possible that a competitive economy always be in equilibrium? Clearly not, for then those who arbitrage make no (private) return from their (privately) costly activity. . .We propose here a model in which there is an equilibrium degree of disequilibrium: prices reflect the information of informed individuals (arbitrageurs) but only partially, so that those who expend resources to obtain information do receive compensation. (Grossman and Stiglitz 1980, 393)

Taking some complementary assumptions, Grossman and Stiglitz concretely model the associated equilibrium, being characterized by an amount of resources invested into informational activities, and a degree of efficiency of market prices.[6] In equilibrium, a population of active portfolio managers with comparative advantages in this profession will emerge, in which the least productive active manager will just be at the margin in terms of earning the costs associated to him. Berk and Green (2002) develop an equilibrium model in which the flows of funds towards successful active managers explain why in equilibrium, the different qualities of managers do not imply that excess returns of actively managed funds are predictable. They assume that returns to active management are decreasing with the volume of funds managed, and the equality of marginal returns is then simply ensured by the higher volumes of new funds flowing to the successful managers.

In a noisy real world equilibrium with risk, there will always be a significant share of active managers who will have generated a loss, ex post. In equilibrium, anyway, active portfolio management will not be an arbitrage, i.e. the decision to invest money in a passively or in an actively managed fund will be similar to the decision to invest in two different stocks: it will be a matter of diversification, and in a world with transaction costs and thus imperfect diversification, probably also of personal knowledge and risk

[6] In contrast to this view, Sharpe (1991) argues that necessarily, '(1) before costs, the return on the average actively managed dollar will equal the return on the average passively managed dollar and (2) after costs, the return on the average actively managed dollar will be less than the return on the average passively managed dollar'. He proves his assertion by defining passive management as strict index tracking, and active management as all the rest. The two views can probably be reconciled when introducing some kind of noise traders into the model, as it is done frequently in micro-structure market models with insider information (see e.g. Kyle 1985).

aversion. It appears plausible that any large investor should, for the sake of diversification, at least partially invest in actively managed portfolios.[7] This however does not imply that all large investors should do active management *themselves*. For analysing whether public institutions should be involved in active portfolio management, it is obviously relevant to understand in general how in equilibrium the portfolio management industry should look like. Some factors will favour *specialization of the asset management industry into active and passive management*. At the extreme, one may imagine an industry structure made up only of two distinct types of funds: pure hedge funds and passive funds. This may be due to the fact that different management styles require a different technology, different people, and different administration. The two activities would not be mixed within one company exactly as a car maker does not horizontally integrate into e.g. consumer electronics (e.g. Coase 1937; Williamson 1985). It would just not be organizationally efficient to pack into one company such diverse activities as passive management and active management.

Other factors may favour non-specialization, i.e. that *each investment portfolio is complemented by some degree of active management*. Indeed, the general aim of diversification could argue to always add at least a bit of active management, as limited amounts add only marginal risk, especially since the returns of active management tend to be uncorrelated to returns of other assets. In the case of a hedge fund, in contrast, there is little of such diversification, as the risks from active management are not pooled with the general market risks. It could also be argued that active management is preferably done by pooling lots of bets (views), instead of basing all the views on a few bets. One might thus argue that by letting each portfolio manager think about views/bets, more comes out than if one just asks a few, even if those are, on a one-by-one comparison basis, the better ones. Creativity in discovering arbitrages may be a resource too decentralized over the population of all portfolio managers to narrow down the use of this resource just to a small subset of them. Expressed differently, the marginal returns of active management by individuals may be quickly decreasing, such that specialization would have its limits.

[7] Interestingly, the literature tends to conclude that index funds tend to outperform most actively managed funds, after costs (e.g. Elton *et al.* 2003). This might itself be explained as an equilibrium result in some CAPM like world (because returns of actively managed funds are so weakly correlated to returns of the market portfolio). This extension of the CAPM of course raises a series of issues, in particular: The CAPM assumes homogenous expectations – how can this be reconciled with active management? Probably, an actively managed portfolio is not too different from any other company who earns its money through informational activities, and the speciality that an actively managed portfolio deals with assets which are themselves in the market portfolio should after all not matter.

Finally, one needs to recognize that portfolio management firms do not tend to manage only one portfolio, but several ones, such that one firm may have both actively and passively managed portfolios. It can then also pack active and passive portfolios together into mixed portfolios, following e.g. a so-called 'core-satellite' approach. Having actively and passively managed portfolios in one firm may have the disadvantage mentioned above of putting together two different production processes (like manufacturing cars and consumer electronics), but at the same time it has the advantage to allow for coordination within a hierarchical organization (e.g. in an optimized core-satellite approach).

As the fund management industry is made up of hedge funds, mixed funds (i.e. tracking an index with some intermediate risk budget to deviate from it, being organized or not in a core-satellite way), and passively managed funds, it seems that neither of the two factors working in different directions completely dominates the other. It is in any case important to retain that for every investor, the decision to have some funds dedicated to active management is at least to some extent independent of whether it should do active management itself. In other words: once an investor has decided to dedicate funds to active management, he still faces the 'make or buy' decision determining the efficient boundaries of a firm. While the former decision is one which is to be modelled mainly with the tools of portfolio theory, the latter is one in which tools from the industrial organization literature would need to be applied.

5.2 Public institutions and central banks as active investors

The fact that public institutions tend to have more complex and rigid decision-making procedures, and less leeway in the selection and compensation of portfolio managers due to rules governing the employment of public servants, could be seen as argument *against* genuine active portfolio management. Being in competition with less constrained players also looking for mispriced assets, active management could thus end up in losses, at least if the fixed costs are correctly accounted for. Alternatively, it could be argued that the private sector should not be overestimated either and that there are enough financial market inefficiencies to allow also the central bank to make additional money by position taking. Eventually, a good performance attribution model should be in place to decide on whether active management contributes positively to performance (see Chapter 7 of this book). The conclusion may be different for different types of positions

taken. What is important is to come eventually to a net cost–benefit assessment, i.e. one in which the full costs of active management enter the analysis. Also, possible positive externalities (point 9 in Section 2) would have to be considered in this analysis. Also the fact that *central banks are insiders inter alia on interest rates* could be seen to argue against active management. It may however be possible to remedy this issue partially with a Chinese wall, behavioural rules, or by precluding the kind of positions which are most subject to potential use of insider information, namely yield curve and duration position in the central bank's own currency. These measures could also be combined, on the basis of some assumptions on the likely relevance of insider information for different types of positions. The specificity that *public institutions do not have the task to maximize their income, but social welfare*, would also argue against active management of central banks, as it is at least partially a re-distributive, zero-sum activity (the argument of Hirshleifer (1971)). One could argue that active portfolio management, being based on private information, is by definition not compatible with transparency and accountability standards that should apply for public institutions, and thus are not a natural activity for public institutions. Related to that, one may argue that active management unavoidably represents a *source of reputation risk*. Central banks have special reasons to develop market intelligence, since they need to implement monetary policy in an efficient way, and need to stand ready to operate as lender of last resort. Active management could be an instrument contributing to the central bank's best possible understanding of financial markets, which is useful for other core central bank tasks, such as monetary policy implementation or the contribution to financial stability. A counter-argument could be that intelligent *passive* portfolio management, using a variety of different instruments, also force the staff to understand.

These specificities are affected by an *outsourcing of active management* to private investment companies to a different extent. Depending on which weight is given to the different pro- and con-active management specificities, one thus may or may not find outsourcing attractive. It has also been argued that a *partial* outsourcing is attractive, as it provides a further reference for assessing the performance of active managers in general. On the other side one may argue that outsourcing is itself labour intensive, and would thus be worth it only if the outsourced amounts are substantial. It has been estimated that at least two-thirds of central banks use external managers for some portion or even all of their reserves.

Table 1.8 Trading styles of central bank reserves managers according to a JPMorgan survey conducted amongst forty-two reserve managers in April 2007

	Describes exactly style	Describes somewhat style
Buy and hold	12%	19%
Passive benchmark tracking	12%	24%
Active benchmark trading	64%	21%
Active total return trading	10%	12%
Segregated alpha trading	2%	7%

Source: JPMorgan 'New trends in reserve management – Central bank survey', February 2008.

One may try to summarize the discussion on the suitability of active portfolio management for central banks and other public investors as follows. First, genuine active management is based on the idea that private information, or private analysis, allows detecting wrongly priced assets. Over- or underweighting those relative to the market portfolio then allows increasing expected returns, without necessarily implying increased risk. There is no doubt that in equilibrium, active management has a sensible role in financial markets. Second, while it is plausible as well that in equilibrium, large investors will hold at least some actively managed portfolios, it is not likely that every portfolio should be managed actively. In other words, it is important to separate the issue of diversification of investors into active management from the industrial organization issue of *which* portfolio managers should take up this business. Indeed, hedge funds, passively managed funds and mixed funds coexist in reality. Third, a number of central bank specificities could appear to argue against central banks being amongst the active portfolio managers. There is, however, one potentially important argument in favour of central banks being active managers, namely the implied incentives to develop market intelligence. As it is difficult to weigh the different arguments, it is not obvious to draw general conclusions. Eventually, central bank investment practice has emerged to include some active management, mostly undertaken by the staff of the central bank itself, and sometimes being outsourced.

Table 1.8 provides a self-assessment of forty-two central bank reserves managers with regard to the degree of activism of their trading style, such as collected in the JPMorgan reserve managers survey. It appears that the style called in the survey 'active benchmark trading', i.e. benchmark tracking with position taking in the framework of a relatively limited risk budget, is predominant amongst central banks.

6. Policy-related risk factors

Section 8 of this chapter will develop the idea of an integrated risk management for central banks. An integrated risk management obviously needs to look at the entire balance sheet of a central bank, and at all major risk factors, including the *non-alienable* risk factors (i.e. the risk factors relating to policy tasks). This section discusses four key non-alienable risk factors of central banks. While Section 3 explained how the underlying policy tasks have made large-scale investors out of central banks, the present section looks at them from the perspective of integrated central bank risk management.

Genuine threats to the structural profitability of central banks, which are often linked to policy tasks, have been discussed mainly in the literature on central bank capital. A specific model of central bank capital, namely the one of Bindseil *et al.* (2004a), will be presented in the following section. Here, we briefly review the threats to profitability that have been mentioned in this literature. Stella (1997; 2002) was one of the first to analyse the fact that several central banks had incurred such large losses due to policy tasks that they had to be recapitalizd by the government. For instance in Uruguay in the late 1980s, the central bank's losses were equal to 3% of GDP; in Paraguay the central bank's losses were 4% of GDP in 1995; in Nicaragua losses were a staggering 13.8% of GDP in 1989. By the end of 2000, the Central Bank of Costa Rica had negative capital equal to 6% of GDP.[8] Martínez-Resano (2004) surveys the full range of risks that a central bank's balance sheet is subject to. He concludes that, in the long run, central banks' financial independence should be secure as long as demand for banknotes is maintained. According to Dalton and Dziobek (2005, 3):

> Under normal circumstances, a central bank should be able to operate at a profit with a core level of earnings derived from seigniorage. Losses would have, however, arisen in several central banks from a range of activities including: open market operations; sterilization of foreign currency inflows; domestic and foreign investments, credit, and guarantees; costs associated with financial sector restructuring; direct or implicit interest subsidies; and non-core activities of a fiscal or quasi-fiscal nature.

In a recent comprehensive study, Schobert[9] analyses 108 central banks' financial statements over a total of 1880 years. Out of those, 43 central

[8] See also Leone (1993), Dalton and Dziobek (2005).

[9] See Schobert, F. 2007. 'Risk management at central banks', unpublished presentation given in a central banking course at Deutsche Bundesbank.

banks recorded at least once an annual loss, and 146 years of losses were observed in total. She attributes 41 per cent of loss years to the need to sterilize excess liquidity (which is typically due to large foreign exchange flows into the central bank balance sheet), and 33 per cent to FX valuation changes (i.e. devaluation of foreign reserves). Only 3 per cent would be attributed to credit losses, and there is no separate category regarding losses due to market price changes other than foreign exchange rate changes. In other words, interest rate risks were not considered a relevant category, probably because never was an end of year loss driven by changes of interest rates. These findings confirm that policies, and in particular foreign exchange rate policies, are the real threat to central bank profitability and capital, and not interest rate and credit risks; although those are the types of risks to which central bank risk managers devote most of their time, as these are the risks that are controlled through financial risk management decisions, while the others are largely implied by policy considerations, which may be seen to be outside the reach of financial risk management. However, even if a total priority of policy considerations would be accepted, still the lesson from the findings of Schobert and others is that when optimizing the financial assets of a central bank from the financial risk management perspective, one should never ignore the policy risk factors and how they correlate with the classical financial risk factors. In the following, the four main identified policy risk factors are discussed in more depth.

6.1 Banknotes, seignorage, and liquidation risk

The privilege to issue banknotes is a fundamental component of central bank profitability, and hence the scenario that this privilege will lose its relevance is one of the real long-term risk factors for central banks. For a very long time, central bankers and academics have speculated about a future decline in the demand for banknotes. Woodford (2001, section 2) provides an overview of recent literature on the topic. One may summarize: while there are a variety of reasons why improvements in information technology (like the more systematic use of smart cards) might be expected to reduce the demand for banknotes, it does not appear that those developments are in real competition to the main uses of banknotes, which explain the high amounts of banknotes in circulation (of around EUR 1500 per capita in the euro area). Moreover, the actual use of e.g. smart cards has progressed only slowly, while e.g. credit cards have been in circulation for a long time. Goodhart (2000), for example, argues that the popularity of

currency will never wane – at least in the black-market transactions that arguably account for a large fraction of aggregate currency demand – owing to its distinctive advantages in allowing for unrecorded transactions. The evolution of banknotes in circulation over the last years in both USD and EUR, has not given any indication of a decreasing trend (more the contrary). Another indicator against the hypothesis of a forthcoming disappearance of the circulation of banknotes and coins in the case of the euro area is a look at denominations in circulations. In fact, only 83 out of EUR 620 billion of currency in circulation was denominated in banknotes with a nominal value of less than EUR 50 or in coins, i.e. the typical transaction balances. More than 50 per cent of the value of currency in circulation was in denominations above EUR 50, i.e. denominations one rarely is confronted with. In line with this observation, Fischer *et al.* (2004) conclude that the various methods they apply all indicate rather low levels of transaction balances used within the euro area, namely of around 25–35 per cent of total currency.

In case one would be able to establish a probability distribution for the evolution of banknotes in circulation over a certain, say ten-year horizon – how to integrate exactly the risk factor 'volume of banknotes in circulation' into a risk management model for central banks? Two main types of risks may be distinguished in relation to banknotes. First, a decline of banknotes would imply a decline of seignorage, even if the assets counterbalancing banknotes would be perfectly liquid. This is the topic of Section 7. Second, in the short term, the decline of banknotes creates liquidation and liquidity risk. 'Liquidation risk' is the risk that assets need to be liquidated before the end of the investment horizon. In such a case, the originally assumed investment horizon would have been wrong, and accordingly the assumed optimum asset allocation would actually not have been optimal. 'Liquidity risk' is the risk that due to the need to undertake large rapid sales, prices obtained are influenced in a non-favourable manner. For the issue considered here (reversal of trend growth in banknotes in circulation), liquidation risk could appear more important than liquidity risk. It should be noted that it is not only the uncertainty about the demand for banknotes which creates liquidation risk. Similar risk factors are: (i) for domestic financial assets, the need to build up foreign reserves; (ii) for foreign reserves assets, the need to liquidate those assets for foreign reserve interventions; (iii) the need to buy assets for some other policy reasons, such as emergency liquidity assistance to banks. The case with uncertainty of non-maturing liabilities has been modelled e.g. by Kalkbrener and Willing (2004),

who propose a general quantitative framework for liquidity risk and interest rate risk management for non-maturing liabilities, i.e. allowing to model both an optimal liquidity and maturity structure of assets on the basis of the stochastic factors (which includes interest rate risks) of liabilities. Overall, it appears that banknotes are much more important in terms of risk factors as putting seignorage at risk, than to create liquidation and liquidity risks.

6.2 Monetary policy interest rates

Another major risk factor to the long-term profitability and positive capital of the central bank is that the currency falls into the deflationary trap, as did the JPY in the 1990s, or as other currencies did e.g. in the 1930s. Monetary policy rates need to be set by a central bank according to an objective – normally to maintain price stability. To model this, it may be assumed that there is a Wicksellian relationship between inflation tomorrow and inflation today, i.e. inflation normally accelerates if interest rates on monetary policy operations are below the sum of the real rate on capital and the current inflation rate.[10] This Wicksellian relationship is also the basis for the central bankers' *Angst* from a deflationary spiral: if deflation ever reaches a momentum to be larger than the real interest rate, then negative nominal interest rates would be required to change this deflation again into price stability or inflation. As negative nominal interest rates are however in principle impossible, at least as long as banknotes in their current form exist, deflation would accelerate more and more, and prices could never stop falling again, eventually making a total monetary reform unavoidable. While long-lasting deflations in which the central bank put nominal interest rates to zero without this solving quickly the problem have indeed been observed, a true deflationary spiral which ended in ever-accelerating price decreases has not.

Modelling the deflation risk factor from a financial investment perspective requires understanding to the largest possible extent the factors determining the setting of the policy rate by central banks, i.e. what the macro model of the central bank looks like, and how the exogenous variables deemed relevant by the central bank will evolve and possibly exert shocks pushing the system into deflation. The model in Section 7 provides more detailed insights into how one may imagine a stylized relationship between macroeconomic shocks, the monetary policy strategy of the central

[10] See Woodford (2003) for a discussion of such Wicksellian inflation functions.

bank, and the financial situation of the central bank. What should be retained here is that on average, monetary policy rates will reflect the sum of the real interest rate and the central bank's inflation target. *Real interest rates* fluctuate with the business cycle, and may be exposed to a certain downward trend in an aging society (on this, see for instance Saarenheimo (2005) who predicts as a result of ageing a possible decline of worldwide real interest rates by 70 basis points, or possibly more in case of restrictive changes in the pension system). For a credible central bank, average *inflation rates* should equal the inflation target (or benchmark inflation rate). A higher inflation rate is in principle better for central bank income than a lower inflation target. However, of course, the choice of the inflation target should be dominated by monetary policy considerations. Moreover, the amount of banknotes in circulation will depend on the expected inflation rate, i.e. the central bank will face a Laffer curve in the demand for banknotes (see e.g. Calvo and Leiderman 1992; Guitierrez and Vazquez 2004). Therefore, the income maximizing inflation rate will not be infinite.

For a proper modelling of the short-term interest rate and its impact on the real wealth of the central bank (including correlation with other risk factors), it will be relevant to also distinguish shocks to the real rate from shocks to the inflation rate. This is an issue often neglected by investors.

6.3 Foreign exchange reserves and exchange rate changes

Foreign exchange rate policy is one of the traditional elements of central bank policy. This typically implies the holding of foreign exchange reserves, creating the risks of mark-to-market losses. ECB internal estimates show that, at conventional confidence levels, around 95 per cent of the total VaR of the ECB can be attributed to exchange rate risks (including gold). Also independently of foreign exchange rate movements, holding foreign exchange reserves is typically costly for central banks, in particular for countries which have (i) a need to mop up excess liquidity in their domestic money market; (ii) have higher domestic interest rates than the interest rate of the reserve currency; (iii) of which the currency is subject to revaluation gains. While this situation appears in contradiction with covered interest rate parity, it has actually been relevant for a large number of development and transition countries for a number of years. Rodrik (2006), for example, estimates the income loss for these countries to have been on average around 1 per cent of their GDP, whereby he also concludes that 'this does not represent too steep a price as an insurance premium against financial crises'. Interestingly the

survey of Dalton and Dziobek (2005, 8) reveals that *all* of the substantial central bank losses they detected during the 1990s, concerning Brazil, Chile, the Czech Republic, Hungary, Korea and Thailand, reflected some foreign reserves issue. In fact all of these reflected a mismatch between returns on foreign reserves assets and higher costs of absorbing domestic liquidity (reflecting both interest rate differentials and revaluation effects.). In Schobert's analysis,[11] 74 per cent of observed annual central bank losses were due to FX issues.

6.4 The central bank as financial crisis manager

Financial crisis management measures often imply particular financial risk taking by the central bank (see Chapter 11 for details). This should be factored into an integrated long-term risk management of the central bank: in bad tail events, the central bank may not only make losses with its investments, but may also have to do costly crisis management operations, implying possibly losses from two sides. As indicated by Schobert, only 3 per cent of annual central bank losses would have been driven clearly by credit losses, of which those driven by emergency liquidity assistance operations would be a subset. A typical problem in developing countries' central banks are non-performing loans to banks in central bank balance sheets which are not really originating from ELA, but from the granting of credit to banks without following prudent central banking principles, maybe upon request or order of the Government (see e.g. Dalton and Dziobek 2005, 6). Although relevant for many central banks, in particular in developing countries, the model in Section 7 does not include this risk.

7. The role of central bank capital – a simple model

Capital plays a key role in integrated risk management for any financial institution, as it constitutes the buffer against total losses and thereby protects against insolvency. The Basel accords document the importance attached to bank capital from the supervisory perspective. This section provides a short summary of a model of central bank capital by Bindseil, Manzanares and Weller (2004a), in the following referred to as 'BMW'. The

[11] Schobert, F. 2007. 'Risk management at central banks', unpublished presentation given in a central banking course at Deutsche Bundesbank.

main purpose of BMW had been to show how central bank capital may matter for the achievement of the central bank's policy tasks. The mechanisms by which central bank capital can impact on a central bank's ability to achieve price stability were illustrated in this paper by a simple model in which there is a kind of dichotomy between the level of capital and inflation performance. The model is an appropriate starting point to derive the actual reasons for the relevance of central bank capital in the most transparent way. The starting point of the model specification is the following central bank balance sheet.

Stylised balance sheet of a central bank

Assets	Liabilities
Monetary policy operations ('M')	Banknotes ('B')
Other financial assets ('F')	Capital ('C')

Banknotes are assumed to always appear on the liability side, while the three other items can be *a priori* on any side of the balance sheet. For the purpose of the model, a positive sign is given to monetary policy and other financial assets when they appear on the asset side and a positive sign to capital when it appears on the liability side. The following assumptions are taken on each of these items:

- *Monetary policy operations* can be interpreted as the residual of the balance sheet. This position is remunerated at i_M per cent, the operational target interest rate of the central bank. Assume that the central bank, when setting, follows a kind of simplified Taylor rule of the type $i_{M,t} = 4 + 1.5 (\pi_{t-1} - 2)$. According to this rule, the real rate of interest is 2 per cent and the inflation target is also 2 per cent.[12] An additional condition has also been introduced in the Taylor rule, namely that in case it would imply pushing expected inflation in the following year into negative values, the rule is modified so as to imply an expected inflation of 0 per cent. It will later be modelled that for profitability/capital reasons, i.e., reasons not relating directly to its core task, the central bank may also deviate from this interest rate setting rule.
- *Other financial assets* contain foreign exchange reserves including gold but possibly also domestic financial assets clearly not relating to monetary policy. Assume it is remunerated at i_F per cent. The rate i_F per cent may

[12] See e.g. Woodford 2003 for a discussion of the properties of such policy rules.

be higher or lower than i_M per cent, which depends *inter alia* on the yield curve, international imbalances in economic conditions, the share (if any) of gold in F, etc. Also, F can be assumed to produce revaluation gains/losses each year. One may assume that $i_{F,t} = i_{M,t} + \rho + \omega_t$ with normally, but not necessarily, $\rho > 0$, implying that the rate of return on F would tend to be higher than the interest rate applied to the monetary policy instruments, and ω_t is a random variable with zero mean reflecting the associated risks. F can in principle be determined by the central bank, but it may also be partially imposed on the central bank through its secondary functions or ad hoc requests of the Government. Indeed, F may include, especially in developing countries, claims resulting from bank bailouts or from direct lending to the Government, etc. Typically, such assets are remunerated at below market interest rates, such that one would obtain $\rho > 0$. The model treats financial assets in the most simplistic way, but this is obviously where the traditional central bank risk management would be very differentiated about (while ignoring the three other balance sheet items).

- *Banknotes* are assumed to depend on inflation and normally follow some increasing trend over time, growing faster when inflation is high. Assume that $B_t = B_{t-1} + B_{t-1}(2 + \pi_t)/100 + B_{t-1}\varepsilon_t$, whereby π_t is the inflation rate, '2' is the assumed real interest or growth rate and ε_t is a noise term. It is assumed that the real interest rate is exogenous. Despite the development of new retail payment technologies over many years and speculation that banknotes could vanish in the long run, banknotes have continued to increase in most countries at approximately the rate of growth of nominal GDP. Our stylized balance sheet does not contain reserves (deposits) of banks with the central bank, but it can be assumed alternatively that reserves are implicitly contained in banknotes (which may thus be interpreted as the monetary base). The irrelevance of the particular distribution of demand between banknotes in circulation and reserves with the central bank would thus add robustness to this assumption on the dynamics of the monetary base.[13]
- *Capital* depends on the previous year's capital, the previous year's profit (or loss), and the profit sharing rule between the central bank and the Government. In the basic model setting, it is assumed that the profit

[13] A switch from banknotes holdings to reserve holdings would imply that seignorage revenues would in the first case stem from a general tax to the holders of banknotes, while in the second case they would be comparable to a tax on the banking sector.

sharing rule is as follows: if profit is positive, i.e. $P_{t-1}>0$ then $C_t = C_{t-1} + aP_{t-1}$ (with $0 < a < 1$) else $C_t = C_{t-1} + P_{t-1}$, and a is set to 0.5. Profits depend on the returns on the different balance sheet positions and on operating costs. With regard to operating costs, q, it may be assumed that they grow over time at the inflation rate. Profit and thus Capital is likely to contain a further random element, which reflects that extraordinary costs may arise to the central bank when the Government manages to assign additional duties to the bank. In the less industrialized countries, these costs could typically be the support of insolvent banks, or the forced granting of credit to the Government. As mentioned above, such factors can also be modelled as affecting the remuneration rate of financial assets.

An equation that explains the evolution across time of the inflation rate completes the model. A Wicksellian relationship between inflation tomorrow and inflation today is assumed, i.e. $\pi_{t+1} = \pi_t + \beta(2 + \pi_t - i_{M,t}) + \mu_t$, i.e. inflation normally accelerates if interest rates on monetary policy operations are below the sum of the real rate on capital (2 per cent) and the current inflation rate. The noise term μ_t means that inflation is never fully controlled. The equation also implies that there is a risk of ending in a deflationary trap: when $\pi_t < -2$, then, due to the zero constraint to interest rates, prices should start falling further and further, even if interest rates are zero. If $\mu_t \sim N(0, \sigma_\mu^2)$, this can always happen theoretically, but of course the likelihood decreases rapidly when the sum of the present inflation and of the real rate is high. Adding a time index t for the year, the time series are thus determined as follows over time:[14]

$$\pi_t = \pi_{t-1} + \beta(2 + \pi_{t-1} - i_{M,t-1}) + \mu_t \tag{1.1}$$

$$q_t = (1 + \pi_t/100)q_{t-1} \tag{1.2}$$

$$F_t = F \tag{1.3}$$

$$\begin{aligned} &\textit{if } P_{t-1} \geq 0 \textit{ then } C_t = C_{t-1} + aP_{t-1} \textit{ (with } 0 < a < 1), \\ &\textit{else } C_t = C_{t-1} + P_{t-1} \end{aligned} \tag{1.4}$$

$$B_t = B_{t-1} + B_{t-1}(2 + \pi_t)/100 + \varepsilon_t \tag{1.5}$$

[14] The order of the equations, although irrelevant from a conceptual point of view, reflects how the eight variables can be updated sequentially and thus how simulations can be obtained.

$$\text{if } \max(4 + 1.5(\pi_{t-1} - 2), 0) < \frac{\pi_{t-1}}{\beta} + 2 + \pi_{t-1} \text{ then}$$
$$i_{M,t} = \max(4 + 1.5(\pi_{t-1} - 2), 0), \text{ else } i_M = \frac{\pi_{t-1}}{\beta} + 2 + \pi_{t-1} \quad (1.6)$$

$$i_{F,t} = i_{M,t} + \rho + \omega_t \quad (1.7)$$

$$M_t = B_t + C_t - F_t \quad (1.8)$$

$$P_t = i_{M,t} M_t + i_{F,t} F_t - q_t \quad (1.9)$$

This simple modelling framework captures all basic factors relevant for the profit situation of a central bank and the related need for central bank capital. It can also be used to analyse the interaction between the central bank balance sheet, interest rates and inflation. It should be noted that, from equation (1.1) and $i_{M,t} = 4 + 1.5(\pi_{t-1} - 2)$, a second-order differences equation can be derived of the form $\pi_{t+1} - (1 + \beta)\pi_{t-1} + 1.5\beta\pi_{t-2} = \beta + \mu_t$. Disregarding the stochastic component, μ, this equation has a non-divergent solution whenever $-2/3 < \beta < 2/3$. The constant solution $\pi_t = 2, \forall t$, is a priori a solution in the deterministic setting. However, it has probability 0 when considering again the shocks μ_t.

Simulations can be performed to calculate the likelihood of profitability problems arising under various circumstances. The model can be calibrated for any central bank and for any macroeconomic environment. The impact of capital on the central bank's profitability and hence financial independence is now briefly discussed. First, as long as bankruptcy of the central bank is excluded, by definition, negative capital is not a problem per se. Indeed, as long as the central bank can issue the legal tender, it is not clear what could cause bankruptcy. By substitution, using the balance sheet identity, one obtains the profit function:

$$P_t = i_{M,t}(B_t + C_t) + (i_F - i_M).F_t - q_t \quad (1.10)$$

Therefore, a higher capital means higher profits since it increases the size of the (cost-free) liability side. For given values of the other parameters, one may therefore calculate a critical value of central bank capital, which is

needed to make the central bank profitable at a specific moment in time:

$$P_t > 0 \Rightarrow C_t > -\frac{(i_F - i_M)}{i_M} F_t + \frac{1}{i_M} q_t - B_t \tag{1.11}$$

Unsurprisingly, the higher the monetary policy interest rates, the lower the critical level of capital required to avoid losses, since the central bank does not pay interest on banknotes (or excess reserves, i.e. reserve holdings in excess of the required reserves). A priori this level of capital can be positive or negative, i.e. positive capital is neither sufficient nor necessary for a central bank to be profitable. It would also be possible for a central bank with positive capital to suffer losses over a long period, which could eventually result in negative capital. Likewise, a central bank with negative capital could have permanent profits, which would eventually lead to positive capital. Moreover, when considering the longer-term profitability outlook of a central bank in this deterministic set-up, it will turn out that initial conditions for capital and other balance sheet factors are irrelevant and the only crucial aspect is given by the growth rate of banknotes as compared with the growth rate of operating costs. The intuition for this result (stated in proposition 1 below) is that, when considering only the long term, in the end the growth rate of banknotes needs to dominate the growth rate of costs, independently of other initial conditions.

When running Monte Carlo simulations of the model (see Bindseil *et al.* 2004a section 4), the starting value of the array (M_0, F_0, B_0, C_0, π_0, i_0) as well as the level of the parameters (α, β, ρ, σ^2_ε, σ^2_ω, σ^2_μ) will be crucial for determining the likelihood that a central bank will be at a certain moment in time in the domain of positive capital and profitability.

Having shown that in the model above, a perfect dichotomy exists between the central bank's balance sheet and its monetary performance, BMW continue by asking how one explains the observation, made for instance by Stella (2003), that many financially weak central banks are associated with high inflation rates. It is likely that there is another set of factors, related to the institutional environment in which the central bank exists, that is causing a relationship between the weakness in the central bank's financial position and its inability to control inflation. BMW argue that the relevance of capital for the achievement of price stability can be explained by considering what exactly happens in case the privilege to issue legal tender is withdrawn from the central bank. If the central bank lost the right to issue currency, it would still need to pay its expenses (salaries, etc.)

in a new legal tender that it does not issue. Also, banknotes and outstanding credits would need to be redeemed in the new currency at a certain fixed exchange rate. Consider the two cases of central banks with positive and with negative capital with a very simple balance sheet consisting only in Capital, Banknotes, and monetary policy operations.

Two central banks, before their right to issue legal tender is withdrawn

Positive Capital Central Bank

Monetary policy operations	Banknotes
	Capital

Negative Capital Central Bank

Capital (negative)	Banknotes
	Monetary policy operations

After the withdrawal of the right to issue legal tender, both central banks become normal financial institutions. After liquidating their banknotes and monetary policy operations, their balance sheets take the following shape:

Two central banks, after their right to issue legal tender is withdrawn

Positive Capital (former) Central Bank

Financial assets	Capital

Negative Capital (former) Central Bank

Capital (negative)	Financial debt

Obviously, the second institution is bankrupt, and the holders of its banknotes and of liquidity absorbing monetary policy operations are not likely to recover their claims. Also, the institution will immediately have to stop paying salaries and pensions, etc. In case of a positive probability of withdrawal of the right to issue legal tender, central bank capital and profitability will thus matter. In the case of negative capital, the staff and decision-making bodies of the central bank thus have incentives to get out of the negative capital situation by lowering interest rates below the neutral level, which in turn triggers inflation, and eventually an increase of the monetary base up till positive capital is restored.

One may thus conclude that the higher the likelihood of a central bank to lose its right to issue legal tender, the more important central bank capital becomes. As the likelihood of such an event will however never be zero, central bank capital will always matter. Once this conclusion is drawn, one can start deriving, through simulations, which level of central bank capital is adequate to ensure a monetary policy aiming exclusively at maintaining

price stability. Assuming that the central bank will thus normally care about profitability and positive capital, one may, *in the case of negative capital*, substitute the interest rate generated by the Taylor rule $i_{M,t}$ by an interest rate $\tilde{i}_{M,t}$determined as follows (with $\theta<0$ a constant):

$$\tilde{i}_{M,t} = \min(4 + \theta, i_{M,t})$$

The functional form given to the capital term in this equation is, of course, ad hoc. It implies that if capital is negative, the central bank no longer reacts to an increase of inflation (reflected in the suppression of the inflation term) and even reduces rates further, by an amount corresponding to θ. Assuming that central banks will thus follow inflationary policies when having negative capital, and introducing the possibility of large negative shock to profit (due e.g. to a foreign exchange revaluation or 'contingent liabilities' as formulated by Blejer and Schumacher (2000)) in the simple model above, allows deriving a positive relationship between capital and inflation performance. One may then calculate the 'value at risk' of the central bank and determine a capital that with, say, a 95 per cent probability ensures that within one year capital will not be exhausted. This is the approach basically taken by Ernhagen *et al.* (2002) without however the comprehensive modelling framework proposed by BMW.

8. Integrated risk management for public investors

8.1 Integrated financial risk management in general

Integrated risk management is the holy grail of risk management in any bank. It means essentially being comprehensive and consistent in terms of the risk–return analysis and management of all of the institution's activities. Often, the term is also used for corporates, and it is stressed that it includes not only financial risks, but all other sorts of risks, like business or operational risks. Sometimes, the term integrated risk management is also associated with 'best practice' concepts and 'firm-wide risk management' such as done by Jorion (2003, chapter 27). Accordingly, integrated / firm-wide risk management would rest on three pillars, namely best practice policies (clear mission statement, well-defined risk tolerance and philosophy, responsible policies endorsed and understood by the board of directors); best practice methodologies (analytical methods to measure, control, and manage financial risks), and best practice infrastructures (IT systems,

organizational design). Focusing here on the issue of integrating financial risk management in the narrow sense, one may structure the key inputs to this approach in the following somewhat theoretical three categories:

(1) The starting point of an integrated risk management of a bank is a **business model of the bank and of the relevant business lines**. For each business line, a sort of 'production function' has to be assumed, which maps input factors into outputs, and which allows, knowing input and output prices, to calculate a profit contribution function and eventually an optimal size of the different activities.
(2) **Risk factors** have to be introduced into this, whereby these may concern both market prices and the production processes themselves. A description of the stochastics includes ideally the joint probability distributions of all risk factors, or, more pragmatically, some descriptive parameters like variances and covariances.
(3) The **relationship between overall risk taking, leveraging, capital and refinancing costs** has to be established. Choosing certain values of these variables, taking into account the relationship between them, means at the same time accepting a certain probability of default, which is also relevant for the relationship with other, non-financial stakeholders.

On the basis of these inputs, one can then in theory derive simultaneously the following elements of an optimum: First, one may establish the *efficient frontier* of the company in the expected profit–risk plane. Second, by matching the efficient frontier with the risk–return preferences of the company's owners (or other stakeholders), one may find the *point on the efficient set which maximizes the utility function of the owner* (and/or other stakeholders). In this, taxation considerations should be taken into account as well, as taxation is normally not linear (but convex), and therefore makes expected profits shrink when volatility of gross profits increases. Third, in line with the chosen point on the efficient frontier, one obtains optimal amounts of business activities (or asset sizes etc.) in the different business lines, and, accordingly *allocates a risk budget to those business lines and activities*. Moreover, one may implement specific tools of integrated risk management, such as *RAROC* (risk-adjusted return on capital), which allows to check ex post whether capital (or the risk budget) allocation is optimal or not, and which can be used for evaluation and compensation of business units and staff. Finally, another element of the optimum is a degree of local costly risk-mitigating measures in each of the business activities. Integrated risk management is certainly a tough challenge for any company. In view of its complexities and methodological problems (e.g. sub-additivity

of risks), it will in practice always be based on a number of 'heroic' and questionable assumptions in particular regarding risk factors. Moreover, it will be opposed by business lines that would be negatively affected from applying its conclusions.

8.2 Integrated risk management issues for public investors

For public investors, additional difficulties arise from the predominance of policy goals against pure return (risk) considerations, the perceived importance of reputation risks, the difficulties to derive from basic business economics an overall risk tolerance, etc. Consider first the three main inputs to integrated risk management, such as applying specifically to public investors:[15]

(1) **Business model and 'production function' of the public institution and of the relevant business lines**. Public institutions have often only a very limited number of activities which may be deemed to be of a 'business' nature. One of them is the investment of financial assets as far as unconstrained by policy requirements.[16] The overall extent of this activity is given by the amount of funds available, i.e. there is little perceived freedom of deciding on the overall scope of investment activities. However, there is of course room for decisions at the sub-business lines level, like how much to invest into which currency, what asset types to invest in, etc. On a first look, all this may appear to be a relatively simple portfolio optimization problem. However if one takes into account set-up costs per asset type and risk management activities per asset type, the analysis has to be enriched with elements from a more standard business decision of a corporate (or a bank), which has to decide on what businesses to go into or to specialize into. When making the list of a public institution's activities to be considered in an integrated risk management, one should also not forget the policy tasks discussed in Section 6.

(2) **Risk factors**. The risk factors relevant for the public investor are mainly the classical financial risk factors like domestic and foreign interest rate, spread, credit, foreign exchange rate and commodity risk. In addition,

[15] See Sangmanee and Raenkhum 2000 for a first paper on integrated central bank risk management.

[16] Others might be (1) provision of payments or securities settlement systems; (2) provision or reserve management services for other central banks or public institutions; (3) cash handling services. In any case, the correlation structure of these other business lines with investment management by the central bank are of limited relevance such that it is fair to analyse the investment issues separately.

there are the idiosyncratic policy-related risk factors, such as those described in Section 6 for central banks.

(3) **Relationship between total risk budget, leveraging, capital and refinancing costs**. In the case of *commercial banks*, a common way to think is that for a certain business model, there is an optimal rating, e.g. AA, which is associated with a certain probability of default. The bank has thus to look at capital costs, and the risk–return profile of business opportunities, to come to an optimal amount of capital and activities leading to an overall risk–return profile being compatible with the envisaged probability of default. For public investors, the first step of such an approach, namely to set total risks as a function of the desired rating and the capital cost function, seems to be the most difficult one. For instance central banks have, virtually regardless of the risks in their balance sheet, the rating of the relevant central government. Indeed, they can normally cope with substantial negative capital for a long time, as they should in the long run normally return to positive capital (see Section 7). Before simply accepting an ad-hoc risk constraint, one could aim at the following indirect approach, which relies on two assumptions, namely that (1) an implicit franchise capital of central banks can be calculated and (2) that this franchise capital should correspond to economic capital needs for ensuring the relevant sovereign issuer rating. Franchise capital can be calculated from the discounted expected income of a central bank due to its franchise to issue banknotes over a certain period, say ten years. Of course, the choice of the parameters and horizons underlying such a calculus will appear ad hoc, and it should thus more be considered as an illustration of the issue than as an approach to be followed.

Eventually, the company's efficient frontier needs to be matched with the risk–return preferences of the company's owners. For individuals, risk–return preferences are normally derived from the concavity of their utility function, risk aversion following from Jensen's law. For commercial banks, as for institutions in general, assuming a utility function would be too ad hoc. Instead, risk–return preferences should be derived from the business model of the company, the environment in which it operates, and preferences of stakeholders. Typical sources of risk aversion with regard to public institution's profits may be: (i) *Taking the specific perspective of the residual claimant, the Government*: The Government has an interest in the stability of the transfer payments from the central bank for budget purposes. In particular, it will dislike the case that by surprise, the public institution's

payments will be zero in some year, or even more that it would have to recapitalize the public institution. (ii) *Taking the specific perspectives of the Board of the public institution*. Typically, risk aversion of companies also stems from the profit–loss asymmetries implied by a progressive taxation schedule (as average taxes paid increase with the volatilities of profits). In the case of public institutions, a progressive profit transfer function has similar implications: losses are typically kept by the public institution, and a small fixed amount of profits can be kept for some provisions or reserves, but profits in excess of some threshold are distributed fully to the Government. A public institution wishing to increase its capital in the wide sense (including reserves and provisions) over time will try to ensure that it always earns enough to accumulate capital as much as it can, but would not care about how high its profits are beyond that threshold (although it has in practice an interest to keep the Government happy for the sake of its independence). (iii) For companies in general, risk aversion may be implied by *Financial distress in case of large losses*: liquidity costs of fire asset sales, financing premia for replacing capital, general demotivation of stakeholders when the probability of default increases beyond the optimum for the business model. For public institutions, this is probably a less relevant source of risk aversion, since financial distress tends to remain remote. (iv) *Reputation costs being associated with large losses*. This holds for any company, but maybe even more for a public institution, for which the public or the Government may assume that any large losses are due to irresponsible behaviour.

As mentioned before, the relevance of *reputation risks* for public institutions will drive apart the apparent risk preferences of public institutions for tasks assigned to them directly through their statutes, and indirectly derived tasks reflecting a largely unconstrained choice. For the former, only large losses affecting the Government's finances in a substantive way should matter and drive risk aversion, while for the latter, even very small losses are painful. The general aversion of central banks against credit exposures illustrates the issue: a default event affecting a corporate exposure, even if underweighted, is perceived by central bankers to be associated with headline risk, which is often quoted as reason to avoid such exposures. It is not clear how to handle reputation risks in an integrated central bank risk management framework. One could try to quantify the reputation risk associated to the different financial risks, and to formulate one overall risk budget and allocate it in an optimal way. Alternatively, one could argue that reputation risks after all cannot be quantified well and that therefore, for

each class of risks which are homogenous in terms of reputation risk, a separate risk budged for pure financial risks should be set up. The financial risk budget should be the lower, the associated reputation risk the higher. Finally, and this may be most realistic, reputation risk may be handled in an ad hoc way: first by excluding certain types of risk taking, second by keeping some activities low scale.

Since an integrated risk management by definition aims at avoiding too narrow optimization, it also has to take seriously the issue that public institutions should aim at *social welfare*, which is not a priori equal to its own profit and loss. This caveat does not only hold in terms of expected income, but also in terms of allocation of risk in society. For example, central banks tend to hold fixed-income portfolios with a modified duration below the one of the fixed-income market portfolio, implying that central banks hold sub-proportionally little interest rate risk. But is this result plausible from a social welfare optimization point of view? In view of its substantial financial buffers and the large other financial risks central banks assume easily (in particular FX risks), why would it not be expected to behave at least like an average fixed-income investor? One motivation of risk management seems to be particularly in contradiction with social welfare maximization, namely those related to effects driven by the *accounting and profit sharing framework*, as this often relates to zero-sum games within the state sector.

On the basis of this, the following nine concrete commandments on integrated risk management may be formulated for public institutions:

(1) Avoid segregating organizationally the financial risk management of different areas within the public institution (in the case of central banks: domestic financial assets, foreign reserves, monetary policy operations, etc.) or across risk types (credit risk versus market risk). Instead, have one central risk management unit being responsible for the consistent analysis of all financial risks in all areas.
(2) Establish one comprehensive handbook of risk management policies and procedures for the public institution; as such, a summary and overview makes the need of overall consistency obvious (apart from being useful for documentation purposes).
(3) Draw the complete list of risk factors affecting the profit and loss of the public institution, in the short run, but *also in the long run*, and establish how they affect the balance sheet and profit and loss. Such a list should contribute to avoid that important risk factors are forgotten when

optimizing the risk–return features of the central bank's balance sheet. Also, it should ensure that focusing too much on the short term is avoided.

(4) Simulate or conduct a scenario analysis on the medium- and long-run profitability and capital levels of the public institution, such as to identify the risk factors that really matter, and to get a feeling for the risks of having negative capital over a sustained period.

(5) Avoid too narrow, segregated risk–return optimizations, as this by definition leads to sub-optimal results. For instance setting a risk constraint on a domestic fixed-income portfolio leading to a low modified duration may make little sense if actually the investment horizon is long (because there is no reason to expect that the banknote demand will collapse) and there is anyway considerable reinvestment risk (as monetary policy operations are short-term operations).

(6) For the purpose of being able to measure and allocate risk consistently across risk types and business lines, establish methods for obtaining consistent risk measures (e.g. VaR or expected shortfall) that can be applied to all types of financial risks. Report regularly these risk measures, covering comprehensively risk types and business lines, such as to establish an awareness of the reports' addressees regarding the proportions between the different risks, and the need to take a comparative perspective on their justification.

(7) In case of an apparently distorted allocation of the risk budget (e.g. a lot of non-remunerated exchange rate risk, little or no interest rate risk, little or no credit risk, etc.), think carefully about what policy considerations, or associated indirect risks, such as reputation risks, may justify this. Review the outcome if no such reasons can found.

(8) Aim at deriving the concrete parameters of the risk control framework (eligibility, limit setting formulas, valuation principles, haircuts, etc.) on the basis of basic principles, the risk–return preferences of the institution, and appropriate analytical methods. Do not apply without good reasons different assumptions or methodologies to different areas (e.g. investment operations versus monetary policy operations).

(9) When assessing the usefulness of risk-taking activities, do not only look at benefits, but also at costs, before concluding that they are beneficial in terms of risk-adjusted returns.

Implementing each of these points is certainly rewarding, and even the attempt will probably provide useful insights.

9. Conclusions

This chapter started by explaining why public institutions, and in particular central banks, can be understood today also as important financial investors: this is because for many central banks, the large share of foreign exchange reserves is no longer expected to be used for interventions, and because also domestic monetary policy, namely the steering of short-term interest rates, is considered today to take place 'at the margin' thus not constraining the large part of domestic financial assets. After noting this new characteristic of central bank as large scale, only limitedly constrained investors, the chapter provided a broad overview of central bank investment and risk management issues, trying to highlight in particular two aspects often overlooked, namely (1) public institutions specificities that should not be ignored when applying concepts from private financial institution's risk management to them, and (2) issues relating to an integrated risk management for public investors. While the main single risk management techniques of private banks (portfolio optimization, risk measurement and reporting through VaR, limit setting, compliance monitoring, etc.) can normally be transferred one by one in a meaningful way to public investors, this is less obvious for the broad framework of integrated risk management. Ignoring e.g. central bank specificities in this area often means to optimize in a way that is inferior when taking a comprehensive perspective. This chapter discussed the main relating issues and proposed some tools to address them, as the difficulties with integrated risk management should never be taken as an excuse, not even to aim at such an all-encompassing approach. Rejecting the concept means rejecting the goals of comprehensiveness and consistency, which cannot be right.

2 Strategic asset allocation for fixed-income investors

Matti Koivu, Fernando Monar Lora, and Ken Nyholm

1. Introduction

The goal of strategic asset allocation (SAA) is to find an optimal allocation of funds across different asset classes subject to a relatively long investment horizon. The optimal allocation of funds should always reflect the risk–return preferences of an institution and the machinery underlying the strategic asset allocation decisions should be based on a transparent and accountable process with which such allocations can be determined and reviewed at regular intervals. Often 'modern portfolio theory' is presented following Markowitz (1959) and Sharpe (1964) in the context of the Capital Asset Pricing Model (CAPM) and mean-variance portfolio analysis as the basic theory for how equity markets behave in equilibrium and how investors should position themselves on the efficient frontier, depending on their risk aversion.[1] This theory is central to the understanding of modern finance and thus important for students and market practitioners alike. However, when it comes to actual portfolio allocation decisions and the practical implementation of portfolio allocation decisions in public and private investment organizations the CAPM leaves, quite understandably, many questions unanswered. It is some of these missing answers that the present chapter aims at addressing. In doing so, the viewpoint of a strategic investor is taken; however, elements relevant for tactical asset allocation and portfolio managers are also touched upon. In particular, the focal point of the exposition is that of a central bank's reserves management. This perspective naturally narrows the investment universe considerably. As a consequence, the following discussion focuses mainly on an investment universe comprising fixed-income securities, although credit markets are treated to some extent and some of our remarks generalize easily.

[1] See among others Ingersoll 1987; Huang and Litzenberger 1988; Campbell *et al.* 1997.

The main contributions of this chapter are to: (a) present a consistent framework supporting strategic asset allocation decisions; (b) outline and give a detailed and practitioner-oriented account for a selection of quantitative models that support strategic asset allocation decisions; (c) combine the models to form an accountable framework that easily can be expanded to include equity and other assets; and (d) show how the framework allows for integration of credit risk and exchange rate risk.

The rest of the chapter is organized as follows. Section 2 gives a primer on strategic asset allocation; presents a review of the theory underlying strategic asset allocation decisions; introduces different strategic asset allocation approaches and principles that are applied by public wealth managers; and discusses how the theoretical asset allocation models need to be adapted to fit the particular needs of strategic investors. Section 3 describes important components of the ECB investment process from a normative viewpoint. In Sections 4 and 5 it is demonstrated how quantitative techniques can be used to generate expected returns for the asset classes of interest and how the final asset allocation, i.e. the instrument weights, can be estimated. Section 6 shows through an illustrative example how the ECB uses these techniques, which should neither be taken to represent concrete investment advice nor as an 'information package' endorsed by the ECB.

2. A primer on strategic asset allocation

As mentioned above, the term 'strategic asset allocation' refers to a portfolio that through its asset composition reflects the long-term risk–return preferences, investment universe and general investment constraints of the organization in question. This portfolio serves as a yard-stick for the performance of the active layers in the investment process.

In the following, Section 2.1 provides an overview of the general principles underlying strategic asset allocation methodologies, and Section 2.2 outlines the central dimensions comprised by a strategic asset allocation framework. Special attention is paid to methodologies that are rooted in modern portfolio theory. As indicated in the introduction, modern portfolio theory is a natural point of departure for any discussion and presentation of SAA techniques; however, it is not necessarily the end goal and it does not answer all relevant questions. Therefore, the section also contains hints on how to build more ambitious framework specifications, including the approach applied by the ECB, which is presented in the subsequent sections.

2.1 General principles of SAA methodologies

It is difficult to conclude that there is a unified theory for strategic asset allocation. It seems natural to draw on financial theory and econometric techniques to form return expectations and other input variables necessary for the quantitative and qualitative techniques implemented by financial organizations. Equally natural is the predominant reliance on portfolio optimization techniques, either in the form of traditional Markowitz optimization or more elaborate methodologies resting e.g. on Bayesian methods.

The academic community is in general relatively silent on normative considerations regarding SAA techniques useable by public investors. For example, a search on an online library database suggests that no more than five to ten text books are dedicated to the topic (e.g. Leibowitz *et al.* 1995; Campbell and Viceira 2002; Meucci 2005; Satchell 2007). Public organizations have filled this gap to some extent, see e.g. Bank of England (Nugée 2000), Danmarks Nationalbank (2004) and ECB (Bernadell *et al.* 2004). The Royal Bank of Scotland has dedicated considerable efforts to study the issue of reserve management (see Pringle and Carver 2003), and currently publishes annually a report called RBS Reserve Management Trends (e.g. Pringle and Carver 2007). Other worth-mentioning publications about reserve and sovereign wealth management are Scobie and Cagliesi (2000) and Johnson-Calari and Rietveld (2007). In addition, working papers and published research articles can be found (e.g. Claessens and Kreuser 2007), as well as some details of the investment framework of the different central banks in dedicated papers and in annual reports, often available through the institutional webpages. Notwithstanding this, information regarding tools and techniques applied by public organizations as well as conceptual thoughts on framework definitions are at best disperse, and hence no single unified strategic asset allocation platform seems to exist.

One attempt to structure the SAA process is offered by the International Monetary Fund (IMF). In IMF (2005) a summary of country practices is provided together and several case studies, and IMF (2004)[2] gives 'Reserves Management Guidelines', comprising: (1) Management Objectives, Scope and Coordination; (2) Transparency and Accountability; (3) Institutional Framework; (4) Risk Management Framework; and (5) The Role of Efficient Markets. While these guidelines are formulated in very general terms and do not make concrete recommendations, they are helpful in the sense

[2] Replicated in IMF 2005, Annex 1.

that they form a map that allows organizations to manoeuvre in the SAA landscape and leaves the charting of the finer details up to the decision makers of the organization in question.

Another attempt to define the core principles of SAA in central banks was presented in a survey on 'Core principles of strategic asset allocation in the ESCB' conducted by the ECB among national central banks in 2006. From this survey conclusions were drawn, which seem to be broadly in line with the IMF guidelines. Some of these are:

- The strategic benchmark must express medium- to long-term risk return preferences of the organization (with liquidity and security considerations playing a major role in the central banks), and mimic a passive investment strategy, while being efficient enough to serve as a guide for active investment decisions, as well as constituting a portfolio that is easily replicable.
- The benchmark process (i.e. the tools, techniques and ideological background of construction and rebalancing) should be transparent and stay broadly unchanged from one year to the next, although this form of 'framework stability' should not adversely affect the adoption of new and better methodologies.

The ECB survey also detected a notable diversity regarding some central issues of the SAA framework, as the definition and role of the benchmark stability (meaning the stability in the key-risk measures of the benchmark over time e.g. the stability of the modified duration of the benchmark portfolio), the specification of the objective function, the central risk measures and constraints, the use of quantitative and qualitative techniques and the importance of explicitly forward-looking methodologies.

This diversity, according to IMF (2005), is also present in e.g. the formulation of the objectives of holding foreign reserves and the level of integration of liabilities and different risks in the SAA process.

The differences in the approaches followed by the central banking community are probably motivated by the policy and economic environment, the formulation of objectives for the portfolios, and the particular evolution in the risk and portfolio management areas of each institution.

2.2 Evolution of SAA methodologies

Figure 2.1 aims to outline some of the most central dimensions and techniques that together can constitute an SAA framework. The figure should be read from left to right illustrating a process going from simple to complex.

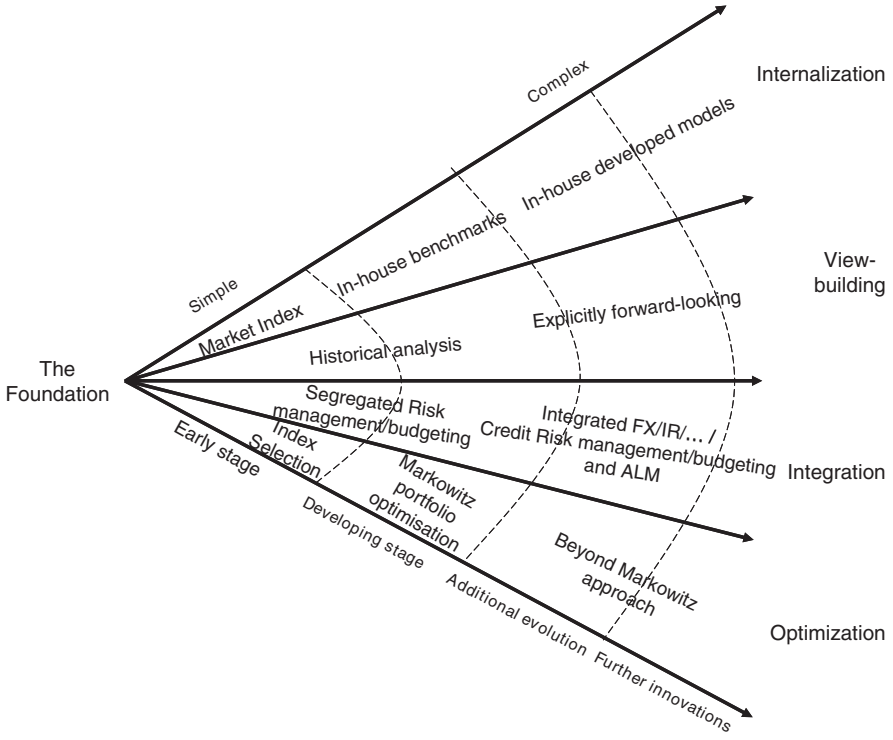

Figure 2.1. Evolution of Strategic Asset Allocation.

To the very left in Figure 2.1 the 'Foundation' is mentioned. The foundation comprises (see also IMF 2004):

- the investment objectives;
- the risk–return preferences;
- the investment horizon;
- the modelling concepts.

It is important to make the formulation of the foundation as transparent as possible and to have a clear view as to how it eventually will be implemented. Given the objectives, the organization may choose to implement additional constraints to ensure the liquidity of the portfolio(s), the diversification of the portfolio(s), and/or other more politically motivated targets such as minimum/maximum exposures to certain asset classes. A high level of liquidity is naturally of great importance if the portfolio serves as the basis for potential foreign reserve interventions, but is probably less relevant if it serves as an investment tranche, or if the funds are managed by a sovereign wealth fund. Also, it is important to clearly delegate the responsibilities

within the organization, for example, who is responsible for the strategic asset allocation decision, and who is responsible for the tactical decisions. Furthermore, one needs to decide on the investment horizon, which according to IMF (2004) is medium to long term. However, in practice it is in many cases necessary to quantify the investment horizon in terms of a given number of years, e.g. one, five or ten years, especially if an explicit forward-looking benchmarking methodology is implemented.

The different stages or degrees of complexity of the SAA process as presented in Figure 2.1 are labelled: 'Internalization', 'View-building', 'Integration' and 'Optimization'. These dimensions are interrelated and it is not possible to derive an exact mapping between them. For example, one organization can prefer to internalize the SAA process, while not paying so much attention to the level of integration of different risks. Another organization can prefer to derive its own benchmark proposals based on state-of-the-art methodologies incorporating explicitly forward-looking views and integrating different risks, but still decide to implement the benchmark through an outsourced benchmark mandate.

The complexity chosen by a given institution for its SAA framework is most likely influenced by institutional-specific features and country-specific traditions, the market developments, the evolution of regulatory requirements, advances in academic research, what peer organizations have implemented and also the natural striving for excellence. Conversely, the choice of a less complex framework can be motivated by the lack of resources, the price of complexity in terms of development and communication requirements, and a desire to obtain framework stability.

2.2.1 Internalization of the SAA process

The dimension called 'Internalization' refers to whether the SAA process is outsourced or whether in-house resources are dedicated to the establishment of quantitative and qualitative tools facilitating SAA proposals to be generated internally within the organization. This is a fundamental question without an easy answer. If an external benchmark is chosen it is often difficult for it to fully reflect the specificities of the organization as defined under the 'Foundation'. In addition, resources still need to be allocated in-house to the monitoring of investment managers and performance evaluation. Notwithstanding this, outsourcing the benchmark composition activities is surely less resource intensive compared to a situation where the benchmarking process is internalized. However, a number of benefits follow from building strategic benchmark competences in-house. For instance, it

can be argued that internally constructed benchmarks may have a second-order positive effect on the organization as a whole by increasing the knowledge and technical skills of staff and decision makers. Other, perhaps more obvious advantages of internally built benchmarks, are, among other things, that they can be tailored to fully match the requirements of the organization in terms of investment universe, investment horizon, fulfilment of risk constraints and idiosyncratic investment objectives, and they facilitate full replicability by the active layers.

Addressing a fixed-income investment universe, a strategic benchmark portfolio can either be constructed as a combination of existing indices published by investment banks and other index providers, or it can be generated in-house and tailored specifically to the preferences of senior management. A general trade-off exists between these two choices: on the one hand, a combination of external indices can easily be made to match the institution's target for modified duration, Value at Risk (VaR) or whichever summary risk statistic the organization relies on, without employing significant in-house resources. However, such a non-tailored benchmark may not fully reflect the investment universe of the organization – an issue of particular importance to central banks. Some external providers may offer customized benchmarks as a solution for this problem. However, external benchmarks often contain a high number of different bonds which makes replication by the active layers on a bond-by-bond basis difficult or even impossible. This is the case even if the individual bond weights are known on a daily basis. In effect, if an external benchmark is used, the active layers can only go neutral against the benchmark by implementing a modified duration and convexity approximation to the benchmark. Such hedges are rarely perfect, and thus a certain unknown amount of residual risk is borne, even when a manager wants to be neutral. Imperfect hedging is also an issue when dealing with synthetic benchmarks. On the other hand, by using in-house resources, a tailor-made index can be made to match fully the preferences of senior management while using instruments comprised by the particular investment universe of the organization in a number that allows for bond-by-bond replication. Whether the active management of the portfolios is outsourced or kept in-house is an important question to be considered when deciding on making or buying, since with external managers the in-house maintenance of this benchmark will not be as operative as the use of external benchmarks, while the appropriate selection of the benchmark and the risk budgets will probably be even more important.

2.2.2 Building views into the SAA framework

As Figure 2.1 further indicates, once the decision has been made to construct internal benchmarks, the level of technical sophistication has to be decided upon. In practice there are naturally no hard limits between modelling philosophies, however the figure suggests separating the technical level required by decision support frameworks on the basis of whether pure historical analyses are performed or whether an explicit forward-looking approach is implemented. The dimension labelled 'View-building' has then an obvious impact on the rest of the dimensions, since the incorporation of explicit views requires a certain amount of methodological sophistication. While a historical analysis could rely on expected returns generated as a simple average of observed historical yields for the relevant investment universe, an explicit forward-looking methodology for generating expected returns would probably require a forecasting model of some sort.

This is an issue that is often passed over too easily in financial textbooks. A standard textbook treatment of portfolio theory will most likely focus on the mathematical underpinnings and be relatively silent on how input parameters should be generated. In any practical implementation, while the mathematics should be understood and correctly applied, the input parameter values are hugely important. For example, when generating *expected* returns as input for the asset allocation process it is necessary to tackle the issue of the investment horizon also. That is, how long a future time period do the returns apply to? A usual solution is to calculate the expected returns from a given portion of historical data collected in the database available to the analyst. Such an approach naturally hinges on the assumption that the future period under investigation is identical in distribution to past observed data. This may be reasonable if the investment horizon tends to infinity and the historical data covers a long enough time period. Otherwise, it may be necessary to build a return projection model, for example relying on market consensus expectations to factors that correlate with the returns of interest.

It is worth emphasizing that any sort of view formation has a translation into an explicitly forward-looking methodology, since the assumptions regarding the expected returns imply a closed subset of evolutions for the yield curves, exchange rates and asset prices. Thus, while an explicitly forward-looking methodology aims to offer more plausible distributions of returns given the current and projected economic environment (conditional

projections), a historical analysis would imply projections of the variables, which may be consistent with past observations, but will rarely fit into the current or expected future economic environment. Other approaches are also sometimes presented as being no-view or no forward-looking, as the naïve no-change projections for the yield curve (based on the assumption of the yields following a random-walk process), but they still rely on strong assumptions (no-change) regarding the future evolution of the variables, and thus, cannot be seen as no-view forecasts, but rather as strong-view forecasts.

Some of the strategies for the formation of views rely on the evolution of prices implied by the current prices (and rates), other strategies link the expected evolution of the asset prices to the consensus expectations on some variables (typically macroeconomic variables), while others rely on the notion of market equilibrium. An evolved approach for integrating views in the asset allocation framework consists in 'mixing' some subjective or model-based priors with historical or market equilibrium based anchors (see. e.g. Black and Litterman 1992).

Further innovations regarding the formation of views may serve the purpose of adapting an explicit forward-looking approach to provide additional information to the decision-making bodies of the organization in question. This can be done, for instance, including risk-neutral or arbitrage-free considerations in the framework to consistently model spreads and credit risk or the term-structure of interest rates.

2.2.3 Integration of different risks

This dimension relates to whether liabilities are modelled explicitly in an asset-liability management (ALM) framework, and to what extent different risk sources (e.g. currency, interest rate and credit risk) and portfolios are modelled separately or jointly.

Regarding the level of integration of different risks and portfolios, typically only market risk, and sometimes only interest rate risk, is modelled in the SAA framework, whereas other types of risks, such as liquidity and credit risk, are mainly taken into account by imposing constraints on the portfolio optimization problem. This is quite often also the case for exchange rate risk, since significant uncertainties prevail when jointly modelling of currency and interest rate risk, and because the currency distribution is typically seen as a policy decision. As a result, many central banks look at different SAA problems for different portfolios typically split on the basis of

different objectives (e.g. Hong Kong SAR Backing portfolio vs. Investment portfolio) or different currencies (e.g. ECB: EUR investment portfolio and USD and JPY foreign reserves portfolios).

Regarding the use of ALM approaches, according to IMF (2005), many central banks (including those of Canada, New Zealand and the United Kingdom) apply some sort of ALM for their foreign reserves. It is worth mentioning that, depending on what is meant by the term 'liabilities', completely different approaches can be classified under the heading ALM.[3]

Explicitly forward-looking simulation-based approaches as the one presented in the following sections of this chapter may integrate quite easily different risks and liabilities.

2.2.4 Portfolio optimization: Markowitz and beyond[4]

The last dimension seen in Figure 2.1 is 'Optimization'. To some extent this dimension relies on the choices made for the other dimensions, and on the way the overall SAA problem is formulated; in particular it depends on the general specification of the asset allocation framework, such as the central risk measures used and the specified objective function. The first stage of development regarding this dimension would comprise basic 'index selection', possibly based on qualitative or historical analysis. Modern portfolio theory represents a further level of complexity in comparison to index selection.

Regardless of whether an in-house or an external benchmark is used, any organization needs a decision support framework that is consistent, accountable and that complies reasonably well with up-to-date financial modelling methods and techniques. A natural starting point for such a framework is the mean-variance portfolio theory, as developed by Markowitz, Sharpe, Lintner and Mossin, in combination with a methodology for how to generate expected returns and risk estimates for the relevant investment universe, and an appropriately formulated utility function. Section 4 details different methods for generating expected returns and accompanying risks for a fixed-income investment universe that could fit in a mean-variance formulation of the SAA problem.

Often, in textbooks treating portfolio theory, a picture similar to Figure 2.2 is used to illustrate the optimal trade-off between risk and return and thus

[3] An example of an ALM framework can be found in Claessens and Kreuser (2007).
[4] This section draws on material from Huang and Litzenberger (1988).

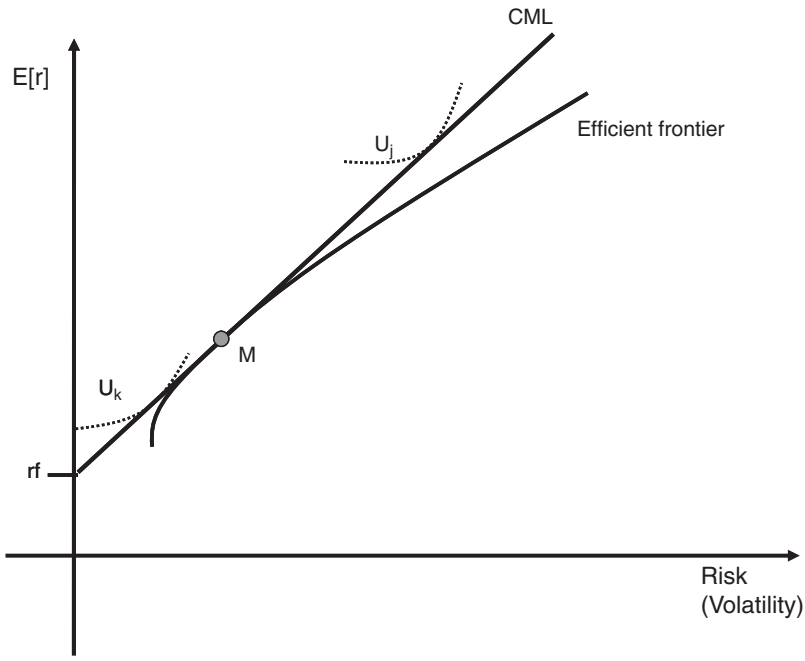

Figure 2.2 The efficient frontier.

how an investor should position themself. The concave curve in the graph is the 'efficient frontier', which traces all mean-variance efficient portfolios. In this context, efficiency refers to the fact that these portfolios are the ones that offer the highest level of expected return for a given level of risk, and the lowest level of risk for a given level of return. If an asset can be found which is risk-less, i.e. uncorrelated with the rest of the assets in the investment universe the linear line in Figure 2.2 can be generated. This line is also referred to as the Capital Market Line (CML). In the case that a risk-less asset exists and is part of the investment universe, a rational investor will choose a portfolio on the CML. Such a portfolio can be generated as a linear combination of the portfolio M (the market portfolio) and the risk-less asset or risk-free rate (*rf*) at any point along the line connecting *rf* and M, and beyond M in the case the portfolio is levered and it is possible to borrow at *rf*, so as to meet the preferences of the investor.

Strategic (as well as tactical) asset allocation would be easy if the real world was adequately reflected by Figure 2.2. Strategic asset allocation would amount to choosing a point on M that matches the institution's risk–return preferences and buy M and the risk-less bond in corresponding amounts.

As for tactical asset allocation: there would be none, if the economy is in equilibrium.

A stylized example can be used to illustrate the calculations needed to derive the efficient frontier as shown in Figure 2.2. It should be mentioned that while these calculations are not used directly in the technical part of the chapter presented in the following sections, they do constitute a cornerstone of the asset allocation theory. In the following sections we show how this foundation can be adapted in a simulation-based exercise allowing for non-normal return distributions and further relaxation of underlying assumptions.

There are two basic inputs to the investment process as illustrated in the figures above. The first is the expected portfolio return, defined as:

$$\mathrm{E}[r(p)] = w' * r$$

where p refers to a given portfolio defined by the portfolio weights collected in the vector w. The variable r is a vector collecting the expected returns for the individual assets comprised by the investment universe, and $'$ refers to the matrix operation 'transpose'. The variance of the returns for portfolio p can then be calculated by:

$$\mathrm{Var}[r(p)] = w' * C * w$$

where C is the covariance matrix of the assets comprised by the investment universe. The principle of minimizing risk for a given level of expected return, used in the above figures to derive the efficient frontier, can be expressed as a mathematical minimization problem:

$$\begin{aligned} \min \quad & \mathrm{Var}[r(p)] = 1/2 * w' * C * w \\ \mathrm{st} \quad & w' * \mathbf{1} = 1 \\ & w' * r = r(p) \end{aligned}$$

The constant '1/2' in the objective function is added for convenience. Adding a constant to the objective function will not change results since the optimal parameter values are found where the first derivative of the objective function equals zero, with respect to each of the variables. It is added because the square of the portfolio weights enter the objective function, and once we differentiate with respect to the weights the one-half will neutralize the exponents that follow from the squaring of this variable.[5]

[5] If the weights were not vectors, the first derivative of the objective function would yield $d(w^{2*}c)/dw = 2^*w^*c$.

The first constraint ensures that the portfolio weights sum to unity; the variable **1** represents a vector of ones having the same dimension as *w*. This constraint is also referred to as the full investment constraint. The second constraint specifies the level of expected return for which the variance should be minimized. By varying $r(p)$ we can see how the whole efficient frontier can be traced out.

We now proceed in a standard way by constructing the Lagrange function by substituting the constraints into the objective function:

$$\min \quad L\{w,f,g\} = \tfrac{1}{2} * w' * C * w + f * (r(p) - w' * r) + g * (1 - w' * \mathbf{1})$$

The solution to the Lagrange function is found by taking the first derivative with respect to each of the parameters, set the derivative equal to zero and solve for the parameter of interest. There are three parameters $\{w,f,g\}$. Let *d* denote the partial derivative, then:

$$dL/dw = C * w - f * r - g * \mathbf{1} = \mathbf{0}$$
$$dL/df = r(p) - w' * r = 0$$
$$dL/dg = 1 - w' * \mathbf{1} = 0$$

The system above constitutes $n+2$ equations with $n+2$ unknowns, if there are *n* assets in the eligible investment universe. Although the *n* asset returns may be correlated (this is in particular the case for a fixed-income investment universe), none of them can be perfectly correlated.[6] Because of this *C* has full rank and is thus invertible. This leads to a solution for the first of the equations above:

$$C * w - f * r - g * \mathbf{1} = \mathbf{0}$$
$$\Rightarrow \quad w = f * (C^{-1} * r) + g * (C^{-1} * \mathbf{1})$$

To make this equation operational we need to know the values of *f* and *g*. These can be derived from the last two derivatives of the Lagrange function i.e. *dL/df* and *dL/dg*. To this end it is helpful to define the following entities:

$$X = r' * C^{-1} * r$$
$$Y = r' * C^{-1} * \mathbf{1} = \mathbf{1}' * C^{-1} * r$$
$$Z = \mathbf{1}' * C^{-1} * \mathbf{1}$$
$$D = X * Z - Y^2$$

[6] If two assets were perfectly correlated they would be indistinguishable in financial terms and would hence not trade as separate entities.

We now proceed by pre-multiplying the above solution for w by r' and $\mathbf{1}'$.

$$r' * w = f * (r' * C^{-1} * r) + g * (r' * C^{-1} * \mathbf{1})$$
$$\Rightarrow r(p) = f * X + g * Y$$

and

$$\mathbf{1}' * w = f * (\mathbf{1}' * C^{-1} * r) + g * (\mathbf{1}' * C^{-1} * \mathbf{1})$$
$$\Rightarrow \quad 1 = f * Y + g * Z$$

Collecting the derived expressions above gives a system of two equations with two unknowns:

$$\begin{bmatrix} r(p) \\ 1 \end{bmatrix} = \begin{bmatrix} X & Y \\ Y & Z \end{bmatrix} * \begin{bmatrix} f \\ g \end{bmatrix}$$

which can be solved by Cramer's rule. So:

$$f = \frac{\det \begin{bmatrix} r(p) & Y \\ 1 & Z \end{bmatrix}}{\det \begin{bmatrix} X & Y \\ Y & Z \end{bmatrix}} = \frac{r(p) * Z - 1 * Y}{X * Z - Y^2} = \frac{Z * r(p) - Y}{D}$$

$$g = \frac{\det \begin{bmatrix} X & r(p) \\ Y & 1 \end{bmatrix}}{\det \begin{bmatrix} X & Y \\ Y & Z \end{bmatrix}} = \frac{X * 1 - Y * r(p)}{X * Z - Y^2} = \frac{X - Y * r(p)}{D}$$

These solutions for f and g can be substituted into the expression for the weights from above:

$$w \quad = f * (C^{-1} * r) + g * (C^{-1} * \mathbf{1})$$
$$\Rightarrow w = Z - (r(p)/D) * (C^{-1} * r) + X - (Y * r(p)/D) * (C^{-1} * \mathbf{1})$$
$$\Rightarrow w = u + p * r(p)$$

where

$$u = (1/D) * (X * C^{-1} * \mathbf{1} - Y * C^{-1} * r)$$
$$p = (1/D) * (Z * C^{-1} * r - Y * C^{-1} * \mathbf{1})$$

This shows that the set of weights that span all efficient frontier portfolios can be calculated by varying $r(p)$, if one believes the assumptions as outlined above.

However, in practice the situation is different from what is assumed in Figure 2.2. Below is a list some of the differences between theory and practice:

A. The financial markets are probably never in equilibrium, in the sense assumed in the figure above. If equilibrium emerges then there will be no trading because all agents agree on the pricing of the instruments. A casual look on Bloomberg suggests that trading commences continuously. This means that while equilibrium never materializes the collected trading of the market is all the time striving to reach equilibrium but is faced by a perpetual flow of new information that has to be interpreted and translated into monetary terms.
B. The theory underlying Figure 2.2 assumes that agents in the economy only are concerned with two properties of the traded assets, namely their expected return and risk measured by the standard deviation. In reality, financial agents are concerned about other features of returns as well: for example, the VaR or the Conditional VaR[7] (CVaR) on the basis of the empirical return distribution of individual instruments and portfolios, draw-down risk and liquidity risk, to name a few. Actually, the increasing popularity of tail-risk measures (VaR, CVaR) as opposed to volatility is not only the logical consequence of the risk aversion of some institutions, but it is also reinforced by its use in banking regulation and supervision.
C. In reality there does not exist one single risk-free rate. The 'risk-freeness' depends on the investment horizon. For example, a short investment horizon can warrant the use of a short-term government bond as the risk-free rate, while a longer investment horizon warrants the use of a long-term inflation-linked bond as the risk-free rate.
D. Utility functions are not as easily quantified and homogenous across investors as the figure suggests. Some investors cannot even write down explicitly their utility function yet plot it in a two-dimensional diagram.
E. Investors may not have homogenous expectations to the asset returns.
F. All investors do not have the same one-period investment horizon as is suggested by the figure: the length of investment horizons differ quite often, and so does the frequency of rebalancing and benchmark review, i.e. the number of periods to consider in a multi-period optimization.
G. The investment universe is not defined similarly for all investors, e.g. fixed-income managers do not invest in equities. Hence, for fixed-income

[7] Conditional Value-at-Risk is also known as Expected Shortfall.

investors Figure 2.2 represents a bond universe. Even within a bond investment universe investors are diverse. For example, central banks can often invest only in highly rated government bonds and are generally restricted when it comes to investments into lower credit grade bonds.

H. Not all investors are allowed to engage in short selling.

I. Return distributions of individual instruments and portfolios are not necessarily normally distributed as assumed above. For example, given that nominal yields are constrained at zero, i.e. nominal yield cannot be negative, the bond return distributions calculated in low yield environments must exhibit a certain degree of skewness. This phenomenon also generally pertains to return distributions of instruments having short time-to-maturity, on account of the pull-to-par effect, even when yields are at normal levels.

Nevertheless, modern portfolio theory serves as a very good frame of thought and can help shape and stream the investment process. Below, more details are presented on how the mindset of the efficient frontier can be applied to a fixed-income investment universe. This will also serve as background for the material presented in the technical sections of the chapter.

Figure 2.3 shows an adapted version of Figure 2.2 by addressing some of the caveats mentioned above. In particular, the risk-free rate has been deleted following (C) and an additional line representing a VaR constraint has been added. A utility function with such a constraint can be integrated into the investment process by specifying that portfolios are valid for investment as long as their VaR return is above zero.[8] Naturally, other minimum return levels can be imagined, for the purpose of this example, a threshold value of zero is used. In Figure 2.3 it is further assumed that individual m aims to maximize expected return and thus chooses portfolio Z. The dashed line in the first graph represents the VaR constraint, calculated according to:

$$\mathrm{VaR}(k) = E[r(k)] - N^{-1}(a) * \sigma(k)$$

[8] In future references to VaR (and CVaR) in this chapter, a positive figure will represent expected gains at the specified confidence level, while a negative figure will represent expected losses. This interpretation of VaR is better defined by the expression VaR return or return on the tail, since VaR as the well-known risk measure is always presented as a positive number measuring losses.

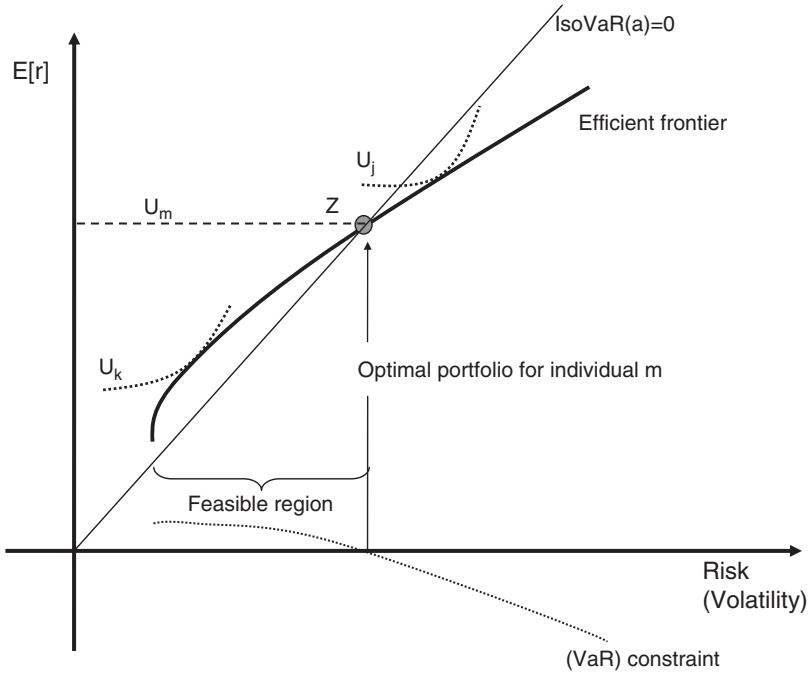

Figure 2.3 Adapted efficient frontier and VaR constraint.

where k counts portfolios on the efficient frontier, VaR is the Value-at-Risk return, $\mathrm{E}[r]$ is the expected return, σ is the standard deviation and $N^{-1}(a)$ is the inverse cumulative standard normal distribution at confidence level a. An additional line (IsoVaR) has been added representing all the different combinations of return and volatility yielding the same VaR for a given confidence level a. Combinations yielding a higher VaR will fall to the left of the IsoVaR, and those yielding a lower VaR will fall to the right. The IsoVaR line, for a VaR with value zero, is defined by the straight line:

$$\mathrm{E}[r] = N^{-1}(a) * \sigma$$

To maximize the expected utility, and assuming at least some part of the efficient frontier lies in the feasible area falling to the left or over the IsoVaR line, individual m will choose the feasible portfolio yielding a higher expected return. When facing a 'normal' efficient frontier, as the one presented in the graph, this portfolio will be determined by the higher intersection of the IsoVaR and the efficient frontier line. The utility of individuals j, k and m (U_j, U_k and U_m) are also shown in Figure 2.3.

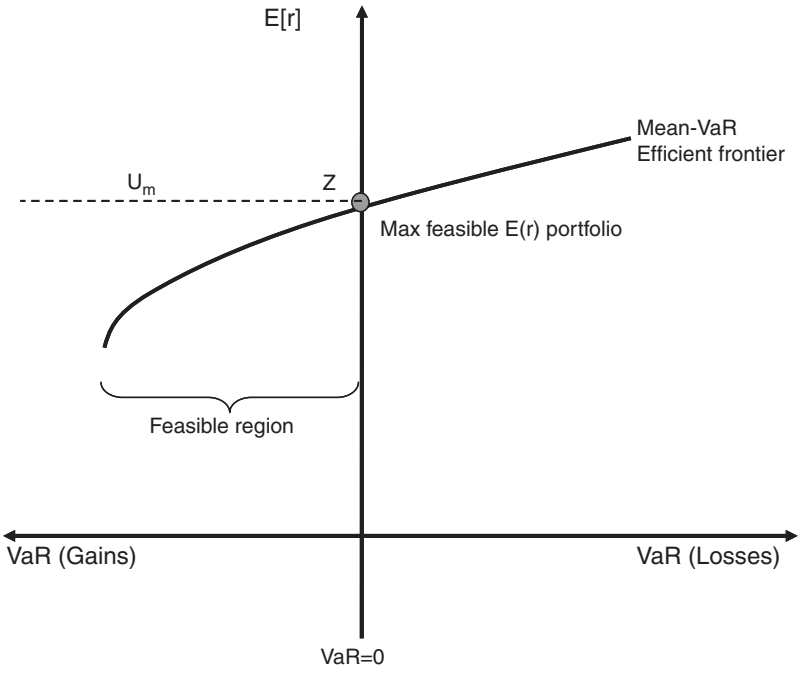

Figure 2.4 Efficient frontier in E[r]–VaR space.

It can be argued[9] that in a CAPM world where individuals have *well-behaved* utility functions (i.e. complete, reflexive, transitive, monotone, convex and continuous indifference curves, as those drawn for individuals *j* and *k*), the inclusion of a VaR constraint may result in the selection of sub-optimal portfolios, and thus, that the VaR constraint may have a shadow-cost. Nevertheless, such utility functions are not as easily quantified and are probably not homogenous across investors. Thus, the choice of a VaR constraint for an expected return maximization strategy does not have to be seen as a costly restriction preventing the achievement of an optimal allocation in relation to a hypothetical *standard* investor exhibiting a *well-behaved* utility function. Rather, it should be seen as a clear specification of a discontinuous utility function for a *non-standard* investor, who is concerned about tail risk instead of volatility.

The risk–return space for this investor could be better represented by Figure 2.4, in which the VaR return does not necessarily need to be derived parametrically under the normality assumption, which has been pointed out

[9] See e.g. Sentana 2003.

as one of the weaknesses of the Markowitz theory, but can represent the empirical distribution observed in historical data or obtained via simulation, as in the framework to be presented in the following sections. It can be seen that the relevant risk–return space for this sort of investor is not the traditional mean-variance space, but rather a mean-VaR/shortfall space.

Further criticism has been raised against optimization with a VaR-based constraint or objective function regarding the fact that portfolios optimized using the historically observed empirical distribution (historical VaR) overfit the data, but do not perform so well out-of-sample. A simulation-based approach relying on the consistent risk measure Conditional VaR (see Pflug 2000) is presented below to illustrate how potentially one can to overcome these problems.

A problem that may appear when using a VaR/shortfall approach is the unfeasibility of the whole efficient frontier for a given (C)VaR[10] constraint, due to special market conditions, the inclusion of harder constraints or the integration of different risks in the optimization exercise. To solve this problem, the (C)VaR constraint could be specified using a different confidence level, or an alternative approach based on the maximization of the (C)VaR return of the portfolio for the selected confidence level could be used.

A utility function corresponding to a general VaR/shortfall approach based on the maximization of return subject to a (C)VaR constraint when the efficient frontier is feasible, or the maximization of the (C)VaR in other case, could be defined as a discontinuous function of the form:

$$U = u(\mathrm{r}, (\mathrm{C})\mathrm{VaR}(\alpha))$$

$$u(r, (C)VaR) = \begin{cases} \mathrm{r}, \text{if } (\mathrm{C})\mathrm{VaR}(\alpha) \geq 0 \\ (\mathrm{C})\mathrm{VaR}, \text{if } (\mathrm{C})\mathrm{VaR}(\alpha) < 0 \end{cases}$$

This sort of utility function corresponds to an investor who uses a certain threshold (risk budget) to consider risk exposures as acceptable or unacceptable, and consequently behave as a risk-neutral or extremely risk-averse investor. Such an approach will be developed in Section 5 of this chapter, due to its widespread use in the central banking community and

[10] In the shortfall approach presented in the following sections, CVaR will be used as a more appropriate risk-measure, although its interpretation in terms of regular VaR will also be shown. Consequently, we have opted for presenting a general formulation in which the expression (C)VaR can refer to either the VaR or the CVaR.

since it corresponds to the ECB specification of the portfolio optimization problem, as shown in Section 6.

Other problems may sometimes be encountered when using the Markowitz portfolio optimization model in practice (see e.g. Scherer 2002). Since Markowitz optimization relies only on expected returns and the covariance of returns, it is possible to find optimal allocations that are counterintuitive by reflecting disproportionally high weightings to only a few of the assets comprised by the eligible investment universe. This problem is sometimes referred to as 'corner solutions' and will materialize more often the higher the correlation is among the assets, and the more similar with respect to risk the assets are. Corner solutions are difficult to handle in the context of SAA, and are also a problem when using other methodologies. Often it is desirable that the optimal allocation has funds smoothly distributed among the asset classes comprised by the eligible investment universe. This is helpful to avoid unnecessary transaction costs when the portfolio is reoptimized on an e.g. annual basis and also in the context of the regular maintenance of the benchmark between the regular review dates. In addition, a central financial principle is that of diversification, which corner-solution portfolios by definition violate. Traditionally this problem has been addressed through the imposition of minimum holdings for the different modelled assets, ideally based on some transparent and theoretically appealing rule as the use of relative market capitalization weights for imposing these constraints. However, it is also possible to remedy the corner-solution issue by using e.g. resampling techniques (see e.g. Michaud 1989; Michaud 1998) or Bayesian techniques and shrinkage methods (see e.g. Black and Litterman 1992).

3. Components of the ECB investment process

Needless to say, it is a complicated task to outline a general framework underlying SAA decisions, and it is even more difficult to fill the framework with quantitative techniques that can produce output relevant to decision makers. Regardless of how such a framework is structured, it will always hinge on a number of central assumptions. In the following we build on the general principles for asset allocation in central banks as suggested by Bernadell *et al.* (2004) and IMF (2004). Although some of these principles have already been presented in Section 2, this exposition is slightly more specific and also more detailed, presenting in a normative fashion a flexible

Investment Process	Strategic benchmark	*risk budget* ↔	Tactical benchmark	*risk budget* ↔	Portfolio managers
Horizon	Longer term		Medium term		Short term
Objective	Translate risk return preferences into an actual asset allocation		Generate outperformance		Generate outperformance
Information content	Public information		Market seasonalities		Daily market sentiment
Responsibility	Decision making bodies		Investment committee		Traders
Methodology	Model based, forward looking		Model based, forward looking		Trader views

Figure 2.5 Components of an investment process.

approach to the investment process, as the one applied by the ECB. In addition, some guiding principles that may serve as a compass for organizations that are in the process of designing or revising their SAA setup are derived.

Bernadell *et al.* (2004) mention the importance of the asset allocation process being framed by an appropriate governance structure, in particular, a three-layer structure is suggested comprising (a) an oversight committee: the strategic level; (b) an investment committee: the tactical level and (c) the portfolio managers. Bernadell *et al.* emphasize that such a multilayer governance structure efficiently support an active investment style where the oversight committee sets the long-term investment strategy in accordance with the risk–return preferences of the organization; the investment committee is responsible for the medium-term investment decisions, i.e. striving to exploit medium-term market movements not foreseen by the strategic benchmark; and finally, it is the portfolio managers mandate to outperform the tactical benchmark by using superior short-term analysis skills and to exploit information that is not taken into account at the tactical level.

Figure 2.5 aims at further expounding the governance structure outlined above. Two dimensions of the investment process are specified: firstly, the three-layer management structure comprising the strategic benchmark, where attention is given to appropriate loading on risk factors (beta selection), the tactical benchmark and the portfolio managers, which both aims at generating portfolio outperformance by searching for and implementing alpha strategies. It is indicated in Figure 2.5 that risk budgets are allocated to the individual active levels. A guiding principle for the allocation of such risk budgets is the organization's overall risk appetite towards

active asset management and the trust it places in the active layers' ability to generate outperformance. Based on these considerations an overall risk budget is allocated to the tactical benchmark and the portfolio managers in sum. The subdivision of the overall limit between the two active layers must be based on the relative expected value added by each layer, i.e. on each layer's ability to generate performance on a risk-adjusted basis.[11] Secondly, some important dimensions of the investment process are outlined; these are: investment horizon, investment objective(s), information content, responsibility and methodology. Each of these dimensions is briefly described below for each of the three layers that form the governance structure.

The investment horizon is tied to the benchmark revision frequency. In the above figure, it is stated that the investment horizon for the SAA should be relatively long reflecting the strategic orientation of this layer. While it may depend on the view of the organization in question one can at least say that the investment horizon should be longer for the strategic layer than for the tactical layers. Furthermore, investment banks will probably have a shorter strategic horizon than a central bank. Taking the case of a central bank, it may be an aim to establish a strategic portfolio that 'sees through the cycle', which implies that the investment horizon probably should be longer than six months, since this is the shortest historical period classified by the NBER as a recession (in the US). Practical considerations may favour an investment horizon that is one, two, five or ten years. Depending on the eligible asset universe a revision frequency can be chosen as regular (or irregular) fix points at which it is analysed whether the previously chosen asset allocation still meets the overall risk–return preferences as defined by the decision-making bodies of the institution.

The greater the information flow to the relevant market segment on which the strategic allocation is defined, and the tighter the deviation bands between the determined risk–return preferences and the actual strategic allocation are, the more often the benchmark should be reviewed. If, on the one hand, the investment universe comprises plain vanilla fixed-income products, as it may be the case for central bank's intervention portfolios, an annual revision frequency may be appropriate. If, on the other hand, the portfolio serves as a store of national wealth, and the investment universe for this or other reasons is broader and comprise assets where new

[11] Naturally, depending on the institution in question, the allocated risk budget may also depend to a smaller or larger extent on political considerations.

information is generated more frequently than it is the case for government bonds, it is probably necessarily to revisit the portfolio allocation more often than annually.

The objective for the three individual layers of the investment process can be seen in the context of an asset pricing model, for example the CAPM or multi-factor representations such as the Arbitrage Pricing Theory (APT). It seems reasonable and intuitive that the strategic level, which accounts for the long-term risk return preferences of the organization, translates these risk–return preferences into an actual portfolio allocation through the use of an appropriate modelling framework. In the context of the CAPM/APT this means that the strategic benchmark is responsible for the overall loading on the market risk factors, in other words the selection of beta. This implies that the strategic benchmark functions as 'the market risk-neutral portfolio' for the active layers and as such provides an allocation that active management can resort to in the absence of active views; reciprocally, the strategic benchmark serves a role as a yardstick for active management by providing a 'market allocation' against which the performance of active layers are measured. It follows indirectly from this objective that the information content that feeds into the decision-making process for the strategic layer is publicly available information and that the aim of the strategic allocation is not to beat the market in any sense, but rather to track market evolutions in an optimal way, paying due attention to the policy objectives as specified by the decision makers of the organization who, as shown in Figure 2.5, are responsibility for the strategic level.

It is the responsibility of the senior management to decide the utility function and the policy constraints that together form the SAA, however, the footwork of the benchmark proposals should be prepared by an entity that is organizationally separated from the active layers i.e. the tactical benchmark and the portfolio managers. Since risk management traditionally is responsible for risk compliance monitoring and performance evaluation of the active layers of the investment process, it is natural that the framework development, analysis and implementation of the strategic benchmark process rests with risk management. This would close the 'food chain' of asset allocation by placing risk management at the top, via its role to formulate the strategic benchmark proposals, and at the bottom, via its risk monitoring and compliance roles.

To fulfil its role in the organization it is seen as crucial that the entity responsible for the methodologies which translate the long-term risk preferences into a de facto replicable portfolio allocation is organizationally

separated from the active layers and is ensured a direct and uninterrupted reporting line to senior management. Otherwise, it is very difficult to establish an accountable and transparent framework and to gain trust and recognition among external economic counterparties as well as the general public.

Another issue that ties in with the importance of accountability in the SAA process is the use of a model-based decision support framework. Rather than making long-term investment decisions based on intuition alone, it is emphasized in Figure 2.5 that the framework in place for the strategic benchmarking should be model based and forward looking. In this context 'model based' need not to be taken too literally: it simply indicates the need to formalize (and document) the details surrounding the benchmark process. This should facilitate easy communication of the benchmark process inside the organization and to external parties. In addition, it builds analysis capabilities on all involved levels of the organization and helps the understanding of the causal relationships within the sub-section of the financial market upon which the eligible investment universe is defined. Needless to say, the actual complexity of the economic, financial and econometric models that are applied to assist the strategic benchmark decisions should be chosen to fit the organization in question.

To be 'forward looking' or, even better, 'explicitly forward looking' refers to the importance of relying on expectations to the future when deciding on long-term asset allocations.

The remaining sections of Figure 2.5 that concern the tactical benchmark and the portfolio managers can be presented in a way similar to the exposition above for the strategic level. However, it is beyond the scope of the present chapter to go into detail with these layers of the investment process given the title of the chapter and its focus on SAA.

As mentioned above, the overall responsibility for SAA rests with the senior management, however, the day-to-day development work on the decision support framework and the preparation of the regular optimal asset allocation reviews should be allocated to a separate unit (e.g. the risk management division). Within the segregation of labour, senior management will decide on the acceptable level of risk to be assumed by the benchmark and otherwise stipulate the relevant policy requirements, while the unit in charge of the day-to-day benchmark process work will devise a framework that meets the specified policy requirements. Figure 2.6 illustrates such an approach and also some of the relevant policy dimensions to

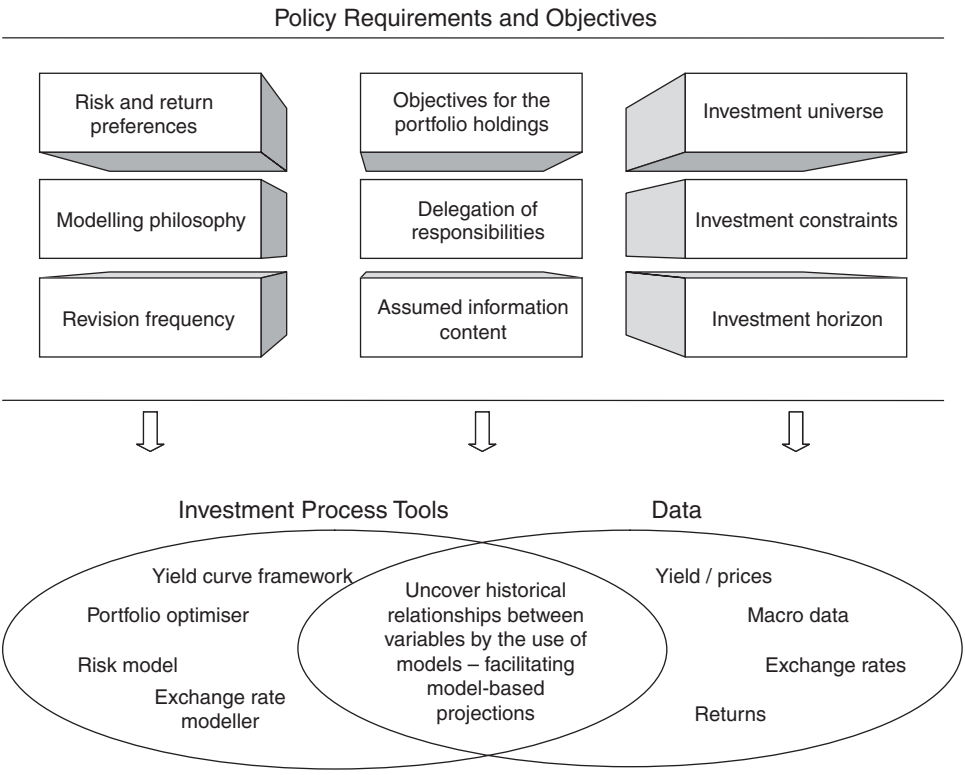

Figure 2.6 The overall policy structure of the investment process.

decide on: in the first part of the figure these high-level policy requirements are illustrated as boxes. These are:

(a) risk–return preferences or put differently, the utility function to be applied;
(b) which modelling philosophy to base the SAA decisions on;
(c) which investment horizon and revision frequency to use;
(d) what the objectives for holding reserves are – if it is a pure intervention portfolio then security and liquidity may be overriding principles, while reserves held as a store of national wealth may induce less strict liquidity and security requirements;
(e) it has to be decided how the responsibility for the organization's asset allocation decisions should be allocated i.e. who is responsible for the strategic and tactical layers in the investment chain;
(f) which information content is assumed to feed into the investment decisions at the various levels of the investment process, e.g. whether it

is appropriate that the strategic level only is based on publicly available information and what 'private' information is assumed to enter at the tactical levels;

(g) related to the 'objective for holding reserves' the eligible investment universe needs to be defined as well as the applicable investment constraints. If for example, the objective is purely intervention based it is likely that the investment universe will be more restrictive than in other cases; likewise the investment constraints will be affected by this.

The second part of Figure 2.6 illustrates the translation of the policy requirements into an accountable, efficient and workable set of models that together serve the role as a decision support framework. In order to produce the output requested by senior management a set of input data should also be chosen. These two components: models on the one hand and data on the other hand, as well as their interplay, is illustrated in the lower part of the figure. To the left in the figure the necessary building blocks are illustrated and exemplified by:

(a) an underlying yield curve model to facilitate the modelling of interest rates for different maturities and credit rating segments comprised by the eligible/potential investment universe, which can be extended to integrate the modelling of credit risk;

(b) a risk model that can be used to generate and simulate risk premia for equities and other instruments that fall outside the yield curve model mentioned in (a);

(c) a module to integrate exchange rates into the evolution of interest rates and equity returns in order to facilitate consistent calculation of returns in a common currency;

(d) a portfolio optimization module that serves the purpose of tying the ends together and which generates the optimal portfolio allocation based on the input generated by the modules referred in (a)–(c) above and the policy constraints such as the utility function and investment constraints.

The estimation of the model segments referred to above depends on data as also illustrated in Figure 2.6, and the intersection between the models and the data suggests that the primary job of the models is to uncover the historical relationship between the variables that are deemed important for the asset allocation process. This means that depending on the investment horizon, the model and the data will serve as a lens through which expected returns and other necessary inputs to the decision support framework are seen.

4. Forward-looking modelling of the stochastic factors

This section present the main building blocks of the ECB's SAA framework. In other words, it details which tools the ECB currently relies on when deciding its strategic benchmark asset allocation. At the outset it is therefore important to outline some of the central assumptions applied by the ECB because this, to a large extent, shapes the models and model framework than can be applied.

Based on the exposition in Section 3, the central policy requirements are: (i) the investment horizon should be medium to long term; (ii) the purpose of holding reserves is to ensure that, if needed, interventions can be conducted in the currency markets, hence, the investment universe comprises only very liquid instrument vehicles such as government bonds, government supported agencies and bonds issued by supranational organizations having a high credit rating; (iii) in the same vein, the risk–return preferences are specified subject to security and liquidity as maximizing expected return while ensuring that there are no losses at a given confidence level over the chosen investment horizon; (iv) it is not the purpose of the strategic benchmark allocation to generate out-performance relative to the market, but rather to serve as an internal optimal market portfolio for the active layers in the investment process and to act as an anchor for neutral positions in the event that the active layers have no views. As a consequence, it should be ensured that only publicly available information enters the SAA process.

Against this policy background it seems natural that a fundamental paradigm of the ECB investment process for the SAA is that of 'conditional forecasting' based on publicly available information. The crux of the approach is to employ a set of transparent and well-documented models that can help generate return distributions for the eligible investment universe on the basis of externally generated predictions of the key macroeconomic variables; these return distributions are then fed into the portfolio optimization module, treated in Section 5, which translates the input data into an optimal allocation complying with the specified risk–return preferences. In this context macroeconomic variables and their expected future time-series behaviour are important because a central premise is that yield curves, and thus fixed-income returns, mainly are functions of the state of the economy, especially at the long-term forecasting horizon that is relevant for the ECB. The 'market neutral view' is implemented by the use of

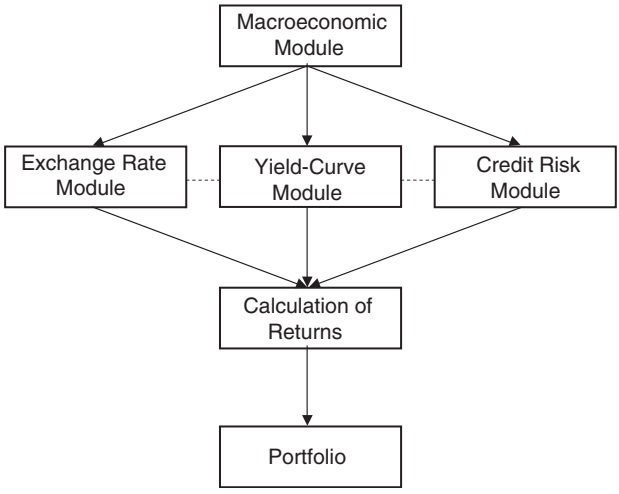

Figure 2.7 Modular structure of SAA tools.

external projections for the average time-series paths for the macroeconomic variables: GDP and CPI growth rates. The use of a simulation methodology allows that random deviations from the externally provided average projection path can be generated in accordance with historical observations. The link between the time-series evolution of the macroeconomic variables and yield curve dynamics is facilitated by a regime-switching model.

The stochastic factors are modelled using the modular structure presented in Figure 2.7, which will generate the necessary input for the portfolio optimizer together with extra summary information used in the decision-making process.

The rest of this section describes the above-mentioned modules in more detail. Section 4.1 presents a general simulation-based framework for modelling the behaviour of GDP and CPI growth on the basis of an exogenously obtained average trajectory path for these variables. Section 4.2 outlines a regime-switching yield curve model and how it can be used to generate predictions conditional on macroeconomic variables, Section 4.3 describes how bonds affected by credit risk (migration and default risk) potentially can be integrated into the framework, Section 4.4 discusses the integration of exchange rate risk and Section 4.5 ties the knot and shows how the produced information can be used to calculate expected return distributions. The portfolio optimizer is presented in Section 5.

4.1 A macro model for multiple currency areas

A first building block of the ECB framework is a model for the joint evolution of the macroeconomic variables across multiple currency areas.[12] As mentioned above, the advocated modelling approach, at least for a main economic scenario, is based on exogenously provided average trajectories for GDP and CPI growth rates. It is necessary to devise a model to properly reflect the uncertainty surrounding the future evolvement of the macro variables. When using externally provided forecasts, the straightforward way to do this would be to rely on historical forecast errors realized by the external forecast providers. Unfortunately, in most cases such data is not available. An alternative approach is therefore needed. The approach currently applied by the ECB is outlined below.

To facilitate that deviations from the exogenously provided mean trajectories for the variables are consistent with past deviations, it can be hypothesized that the GDP growth and inflation follow a vector auto-regressive (VAR) process of the form

$$x_t = c + \sum_{l=1}^{L} A_l x_{t-l} + \varepsilon_t,\ \varepsilon_t \sim N(0, \Sigma)$$

where

$$x_t = \left[g_t^1, \ldots, g_t^K, i_t^1, \ldots, i_t^K\right]'$$

GDP growth is denoted by (g_t^k) and inflation is denoted by (i_t^k) within currency area $k = 1, \ldots, K$, at time t.

x_t is then the vector of macroeconomic variables (GDP and CPI growth rates) observed at time t, and c is a vector of constants. A_l is a matrix containing the auto-regressive parameters at lag l. The residuals, ε_t, are assumed to be normally distributed with a zero mean and covariance matrix Σ.

By assuming that the externally provided forecasts are based on all the available information, i.e. the information contained in the VAR system and additional information not captured by the model, it is possible to simulate deviations around the externally provided mean path by adding to the provided mean path (f_t) the cumulative deviations (u_t):

$$\tilde{x}_t = f_t + u_t$$

[12] This model can naturally also be applied to single currency areas.

The data-generating process used to simulate the cumulative deviations follow the VAR structure:

$$u_t = \sum_{l=1}^{L} A_l u_{t-l} + \varepsilon_t$$

in which the shocks ε_t are sampled from a multivariate normal distribution with zero mean and covariance matrix Σ, and A_l is the estimated matrix containing the autoregressive parameters for the VAR on macro variables.

The modelling approach suggested above generates a very comprehensive set of macroeconomic scenarios with an average trajectory tracking that of the externally provided mean trajectories for the relevant variables. Typically this average can be considered, at least for long-term horizons, as reflecting market consensus views, and is thus also called a 'normal' macro environment. It is sometimes necessary, and it is always a good idea, to also stress test the chosen portfolio allocation. Given the long investment horizon it is natural to generate such stress scenarios by varying the macroeconomic variables; e.g. to stress test the portfolio allocation in a scenario characterized by inflationary and/or recessionary pressures. In order to do this an alternative average trajectory can be used instead of the one reflecting market consensus views as provided by the forecasters.

The described modelling approach for the macro variables represents an important input to the yield-curve model developed in the following section. It facilitates a link to be created between yield-curve dynamics and the dynamic evolution of the macroeconomic variables, and as such, it supports the comprehensive simulation-based modelling approach advocated in this chapter.

Box 2.1. The VAR macro model

It is hypothesized that the GDP growth and inflation follow a vector auto-regressive process of the form

$$x_t = c + \sum_{l=1}^{L} A_l x_{t-l} + \varepsilon_t,\ \varepsilon_t \sim N(0, \Sigma) \tag{2.1}$$

Subtracting the estimated unconditional means for the macroeconomic variables equation (2.1) reads

$$x_t - m = \sum_{l=1}^{L} A_l (x_{t-l} - m) + \varepsilon_t \tag{2.2}$$

Box 2.1. (cont.)

where m represents the vector of historical unconditional means of the GDP and CPI growth rates, and is based on the estimated intercepts (c) and the estimated auto-regressive coefficients (A_l):

$$m = (I - \sum_{l=1}^{L} A_l)^{-1} c,$$

where I is the Identity Matrix with the same dimension as A_l

Define u_t as the cumulative deviations of the macroeconomic variables (x_t) from their unconditional means:

$$u_t = x_t - m$$

where (2.2) can be rewritten as:

$$u_t = \sum_{l=1}^{L} A_l u_{t-l} + \varepsilon_t \tag{2.3}$$

The provided forecast or mean path for the simulation (f_t) is considered to be the expected value or mean conditional on the current set of information, which includes the current and past values (L lags) of the macro variables ($x_{-L} \ldots x_0$) and other exogenous information (e_y) relevant for each of the forecasted periods (e.g. using annual forecasts and a five-year investment horizon, there would be five e_y observations),

$$\mathrm{E}(x_t | x_{-L} \ldots . x_0, e_y) = f_t$$

and from the definition of u_t we could also express f_t as

$$\mathrm{E}(x_t | m, u_{-L} \ldots . u_0, e_y) = f_t$$

So, all the information content of the cumulative errors in every $t \leq 0$ (before the simulation), is already contained in the current forecast or expected value for x in time t, (f_t).

Then, for the simulation, the value of u_t has to be reinitialized to zero for every $t \leq 0$, and thus, only the simulated errors and the estimated autoregressive matrix ($\tilde{A}_l$) will be used to generate the cumulative deviations around the mean path, following the structure presented in equation (2.3).

$$u_t = \begin{cases} \sum_{l=1}^{L} \tilde{A}_l u_{t-l} + \varepsilon_t, \varepsilon_t \sim N(0, \Sigma) & \forall t > 0 \\ 0 & \forall t \leq 0 \end{cases}$$

The last step in the data-generating process would yield the simulated values for x as a result of adding the simulated cumulative deviations (u_t) to the provided mean path (f_t):

$$\tilde{x}_t = f_t + u_t$$

4.2 The yield-curve model

The approach used for modelling the evolution of yields is based on the model developed in Bernadell *et al.* (2005). This modelling framework relies on a Nelson–Siegel parametric description of the shape and location of the nominal yield curve (Nelson and Sigel 1987) in combination with a three-state regime switching model (Hamilton 1994), extended with time-varying transition probabilities that depend on realizations of exogenous macro-economic variables, as mentioned in Section 4.1. Based on the evolution of such macro variables, projections can be constructed for the development of the yield curves within each currency area $k = 1, \ldots, K$. In the presentation below the currency superscript is omitted for notational convenience; the reader can easily confirm that the modelling approach generalizes to a multi-currency context.

The model set-up is based on a Kalman-filter representation with the Nelson–Siegel functional form as the observation equation and the time-series evolution of the Nelson–Siegel factors as the state equation. Regime-switches are incorporated following Kim and Nelson (1999). This setup is similar to Diebold and Li (2006) but is expanded with regime switches and yield-curve evolutions simultaneously for several yield-curve credit segments.

The formulation proposed by Nelson and Siegel (1987) expresses the vector of yields at each point in time as a function of underlying factors and factor sensitivities. These factors have an economic interpretation as yield-curve level, slope and curvature factors.

$$Y_t = [Y_{t,1}, Y_{t,2}, \ldots, Y_{t,Q}]'$$

denote the stacked vector of yield-curve observations for different market segments (e.g. corresponding to different credit ratings) $q = \{1, \ldots, Q\}$ at time t, where each yield curve q consists of $n(q)$ observations with maturities $\tau_q = \{\tau_{1q}, \ldots, \tau_{n(q)q}\}$.[13]

[13] In an effort to circumvent problems that relate to negative yields when using the model for forecasting purposes, the possibility exists to model logarithmic yields rather than the observable yields.

The vector of yields can be expressed using the Nelson–Siegel factors as

$$Y_t = H\beta_t + e_t,\ e_t \sim N(0, \Phi) \tag{2.4}$$

where $\beta_t = [\beta_{t,1}^{level}, \beta_{t,1}^{slope}, \beta_{t,1}^{curve}, ..., \beta_{t,Q}^{level}, \beta_{t,Q}^{slope}, \beta_{t,Q}^{curve}]'$ collects the Nelson–Siegel factors i.e. the level, slope and curvature, for all the considered market segments, H is a block diagonal matrix:

$$H = \begin{bmatrix} h_1 & 0 & ... & 0 \\ 0 & h_2 & 0 & \vdots \\ \vdots & \vdots & \vdots & 0 \\ 0 & ... & 0 & h_Q \end{bmatrix},$$

where the diagonal block elements are defined by the factor sensitivities

$$h_q = \begin{bmatrix} 1 & \frac{1-\exp(-\lambda_q\tau_{1q})}{\lambda_q\tau_{1q}} & \frac{1-\exp(-\lambda_q\tau_{1q})}{\lambda_q\tau_{1q}} - \exp(-\lambda_q\tau_{1q}) \\ 1 & \frac{1-\exp(-\lambda_q\tau_{2q})}{\lambda_q\tau_{2q}} & \frac{1-\exp(-\lambda_q\tau_{2q})}{\lambda_q\tau_{2q}} - \exp(-\lambda_q\tau_{2q}) \\ \vdots & \vdots & \vdots \\ 1 & \frac{1-\exp(-\lambda_q\tau_{n(q)q})}{\lambda_q\tau_{n(q)q}} & \frac{1-\exp(-\lambda\tau_{n(q)q})}{\lambda\tau_{n(q)q}} - \exp(-\lambda_q\tau_{n(q)q}) \end{bmatrix}, \tag{2.5}$$

and e_t is a vector of error-terms with a covariance structure given by Φ.

The three yield-curve factors can be interpreted as the level, i.e. the yield at infinite maturity, the negative of the yield-curve slope, i.e. the difference between the short and the long end of the yield curves, and the curvature. The parameter λ_q determines the segment specific time-decay in the maturity spectrum of factor sensitivities two and three as can be seen from the definition of h_q above. The evolution of the Nelson-Siegel factors (β_t) are assumed to follow AR(1) processes with regime-switching means. The model specification in equation (2.6) assumes three regimes (S, N, I) which imply distinct means for each Nelson–Siegel factor. The interpretation of these three states are based on the shape of the yield curve and defined as: **S**teep, **N**ormal and **I**nverse. The regime switching probabilities at time t are denoted by $\pi_t = \left[\pi_t^S\ \pi_t^N\ \pi_t^I\right]^l$ and a diagonal matrix F collects the auto-regressive parameters.

$$\beta_t = C\pi_t + F\beta_{t-1} + v_t, v_t \sim N(0, \Omega) \tag{2.6}$$

where

$$C = \begin{bmatrix} c_{N,1}^{level} & c_{S,1}^{level} & c_{I,1}^{level} \\ c_{N,1}^{slope} & c_{S,1}^{slope} & c_{I,1}^{slope} \\ c_{N,1}^{curve} & c_{S,1}^{curve} & c_{I,1}^{curve} \\ \vdots & \vdots & \vdots \\ c_{N,Q}^{level} & c_{S,Q}^{level} & c_{I,Q}^{level} \\ c_{N,Q}^{slope} & c_{S,Q}^{slope} & c_{I,Q}^{slope} \\ c_{N,Q}^{curve} & c_{S,Q}^{curve} & c_{I,Q}^{curve} \end{bmatrix}$$

The regime-switching probabilities evolve according to equation (2.7), where π_{t-1} is the regime-switching probability for the previous period. p^{Z_t} is the transition probability matrix which contains the probabilities of switching from one state to another, given the current state.

$$\pi_t = p^{Z_t}\pi_{t-1} \tag{2.7}$$

Equation (2.8) links the transition probabilities to the projected GDP growth rate g_t and the inflation rate i_t as well as threshold values for these variables ($g^\star$ and $i^\star$) which are used to identify distinct macroeconomic environments. In effect, it is hypothesized that there exist three transition probability matrices: p^2 refers to the transition matrix applicable in a recession environment (GDP growth and inflation rate below their threshold values), p^3 refers to an inflationary environment (GDP growth and inflation rate above their threshold values), and p^1 to refers to a 'residual' environment, which can be categorized either as a normal (GDP growth above and inflation rate below their threshold values) or a stagflation-type of environment (GDP growth below and inflation rate above threshold values). More precisely, define:

$$Z_t = \begin{cases} 1 & otherwise \\ 2 & if \quad g_t < g^* \text{ and } i_t < i^* \\ 3 & if \quad g_t > g^* \text{ and } i_t > i^* \end{cases} \tag{2.8}$$

Box 2.2. Transformation of yields and relative slope

In actual applications of the presented modelling philosophy it may be beneficial to apply a data transformation on the historical yield curves. It is observed by Bernadell *et al.* (2005) that it may be beneficial to express the slope of the yield curve relative to the level of the yield curve. This is particularly important because the slope factor is the main determinant of the estimated regimes in the Bernadell *et al.* paper. The transformation applied is the following:

$$\tilde{Y}_t(\tau_j) = Y_t(\tau_N) - \frac{Y_t(\tau_N) - Y_t(\tau_j)}{Y_t(\tau_N)} \forall j \in \{1, 2, \ldots, N\}$$

Analogously, the following transformation could be applied directly on the Nelson-Siegel slope and curvature factors, instead of on the observed yields:

$$\tilde{\beta}_t^{slope} = \frac{\beta_t^{slope}}{\beta_t^{level}}, \text{and}$$
$$\tilde{\beta}_t^{curve} = \frac{\beta_t^{curve}}{\beta_t^{level}}$$

Economically it makes sense to impose the above transformation because the values that the slope can assume are restricted by the yield-curve level, e.g. in a situation where a very low-yield environment prevails the slope can take values that are constrained from above by the value of the level, since the short nominal yield cannot be negative. If a classification scheme is established on the basis of the slope of the yield curve, as it is the case in Bernadell *et al.*, and the estimation period is long and thus potentially covers high- and low-yield environments, it seems necessary to apply the above transformation to control for the effect the level has on the slope factor.

4.3 A model for credit migrations

The yield-curve framework outlined above also allows for the modelling of portfolio credit risk comprising default and migration risk. By evolving forward several yield-curve segments at the same time it is possible, once the credit state of a bond or bond index is known, to price this instrument on the appropriate yield-curve segment. In a Monte Carlo setting, this allows for the calculation of price changes following bond up- and downgrades as well as losses following defaults. For example, if a bond portfolio is constituting X number of AAA bonds, Y number of AA bonds, Z number of A bonds and so forth, then it is possible to simulate the credit state of these bonds over the investment horizon, and once a downgrade is observed, e.g. a downgrade of

a AAA bond to the AA category at time t, then this particular bond will be priced on the AA yield-curve segment from time t and onwards (until it potentially is down- or upgraded), and on the AAA yield-curve segment from time 0 to $t-1$. Due to the yield spread between the AAA and AA yield-curve segments, the bond holder in question will then experiences a negative return from time $t-1$ to t due to the credit migration.[14] Once the Monte Carlo experiment is finalized the simulated losses (and gains) due to migrations and defaults are collected allowing for the calculation of the return distributions containing both credit and market gains and losses.

This section describes in more detail how the credit states of bonds can be simulated. The simulation engine requires the following inputs:

- a portfolio of *Nissuers* number of bond issuers:
  - credit ratings at the initial time for each issuer
  - exposures i.e. the position taken in each issuer
  - the maturity of the holdings in each issuer
  - the coupon rate for each issuer;
- a migration matrix M that holds the migration and default probabilities for each credit rating;
- an asset correlation describing how the credit state of issuers move together over time;
- investment horizon and its discretization of *Nyears* and *Nperiods*.

It is noted that the portfolio is expressed in terms of 'issuer' rather than 'bond' holdings. This is because the default and migration events are linked uniformly to the issuer rather than to the actual bond issues. It is naturally possible to build a model for bonds by appropriately adapting the correlation matrix, which expresses the co-movements between the issuers/bonds. However, this would increase computational time unnecessarily and not bring about more precise results. Instead, generic indices can be constructed on the basis of the bonds issued by the same issuer; these issuer-indices then reflect the characteristics of the underlying bonds, e.g. as a result of a market value weighting scheme, and show the exposure in a portfolio to the included issuers.

[14] It is worth noting that it is not necessarily guaranteed that the simulated yield curves will exhibit positive spread for decreasing credit ratings. If the variance of the innovations to the simulated paths are much greater than the estimated spreads between the credit curves, then it may be that curves cross during the simulation horizon, e.g. that the A curve at one or more time points for one or more maturities are higher than BBB or lower credit rating curves. Such dynamics seem to contradict economic intuition and can be avoided by proper model choices e.g. by modelling the spreads of AA and lower credit ratings as a function of the time-series evolution of the level for the AAA/Gov segment and constants of increasing size.

Based upon the input variables defined above the actual credit simulation follows the steps below:

(1) **Simulation of correlated random variables**. A matrix z of dimension (*Nperiods* x *Nissuers*) is drawn from a normal multivariate distribution with zero mean and a covariance (correlation) matrix Q with dimension (*Nissuers* x *Nissuers*) and showing unity on the diagonal and the asset correlation on off-diagonals. In order to get a random value that is comparable to the credit-rating thresholds implied by the used credit migration matrix, the inverse normal (N^{-1}) of the random variables are taken. This comparison determines whether a given issuer defaults, migrates or has an unchanged credit rating at the observation points covering the investment horizon.

(2) **Convert random numbers into credit ratings at each observation point**. By combining the information from step (1) with the migration matrix M it is possible to derive the credit state of the issuers comprised by the investment universe. M represents the probability over a given horizon (usually annual) that an issuer with a given credit rating upgrades, downgrades, stays unchanged or defaults. After the entries of the migration matrix have been adjusted for the time period under investigation the normal inverse function is applied to *M_adj* to make the entries comparable to z from step (1).

Conditional on the current credit state of the issuer, credit migrations are then determined by comparing the appropriate entry in z to the normal inverse of the corresponding row in *M_adj*. Denote by *Cr_state* the matrix of simulated credit states for the issuers comprised by the portfolio, and let t denote the time period, and let j denote the issuers, then the entries in *Cr_state* are found by:

$$Cr_state(t,j)|k = \left\{ \begin{array}{l} \min\{h\mathbf{1}(z(t,j) > N^{-1}(1 - M_adj(k,h)))|h \in \{1,...,k-1\}\}, \\ \max\{h\mathbf{1}(z(t,j) < N^{-1}(M_adj(k,h)))|h \in \{1,...,k-1\}\}, \\ h\mathbf{1}\left(\begin{array}{l} z(t,j) > N^{-1}(M_adj(k,h+1)) \wedge z(t,j) \\ < N^{-1}(1 - M_adj(k,h-1)) \end{array} \right)|h \in \{k\} \end{array} \right\}$$

where k is the credit state at $t-1$ and $h = 1, \ldots, \bar{H}$ is a numerical equivalent for the rating classes and $\mathbf{1}(.)$ is an indicator function.

(3) **Evolve yield curves forward**. To facilitate the pricing of the generic issuer-indices relevant yield-curve segments have to be evolved forward for the chosen planning horizon. This process is based on the yield-curve model outlined in Section 4.2; hence, a yield curve is projected

forward for each credit rating conditional on the realization of the macroeconomic variables.

(4) **Price the bonds according to their credit rating at each observation point**. Based on the characteristics of the generic issuer indices i.e. coupon rate, coupon frequency, maturity/duration the issuer indices are priced according to conventional bond pricing formulas as given in Section 4.5.

In order to account for the states of the economy included in the credit model, various approaches can be followed:

- Different means other than zero, and even different covariance matrices could be used in the generation of random numbers (step 1) to account for periods of higher correlation and/or volatility, and observed trends in the evolution of creditworthiness under different states of the economy. A calibration would be necessary in both cases to guarantee on average a zero mean and a covariance matrix with unitary standard deviation and average observed correlations.
- Different transition matrices could be used for recessionary, inflationary or normal periods (step 2).

4.4 Modelling exchange rates

A comprehensive SAA framework aiming to support long-term fixed-income investment decisions should ideally comprise a model for the contemporaneous evolution of the central risk drivers (yield curves or yield-curve factors) and exchange rates. Exchange rates are needed in the case where different portfolios are denominated in different currencies, and when one wants to exploit possible international diversification effects by optimizing all portfolios jointly.

If one contemplates doing joint portfolio optimization for international fixed-income portfolios it is naturally necessary to convert returns into a common currency denominator. Once the decision is taken to do joint optimization across several currency areas, one proceeds by calculating expected returns in a common currency for the full investment universe as well as the needed risk measure(s), i.e. the covariance in the case of Markowitz mean-variance optimization, and return distributions in the case a simulation-based approach is used. When the relevant input data have been recovered it is easy to finalize the job by finding the optimal asset allocation fulfilling the risk–return preferences and policy constraints of the organization in question. So, in theory, international fixed-income portfolio

optimization is straightforward. Unfortunately, in practice we need to deal with real-life data, and this complicates the process for at least two reasons. First, international fixed-income returns are generally highly correlated. The volatility of the foreign reserve markets is roughly seven times higher than the volatility of fixed-income market. As a result, given the high correlation of returns and the much higher volatility of the foreign bonds, it may be difficult to obtain any positive allocation to the non-domestic investment vehicles from a pure portfolio optimization exercise. Second, it has proven to be extremely difficult to produce reliable forecasts for exchange rate movements (see among others Meese and Rogoff 1983). The difficulty in predicting exchange rates may be mitigated, to some extent, if the relevant currency pairs are integrated in a joint modelling framework comprising also yield-curve factors and macroeconomic variables. Koivu *et al.* (2007) do this and report encouraging results. In fact, they hypothesize that the improved forecastability of exchange rates originates from the incorporation of yield-curve factors (they use the Nelson–Siegel yield-curve factors: level, slope and curvature) rather than just one single point on the yield curve as it is often done in empirical tests of e.g. the interest rate parities.

Even in the event that a given model parameterization, when estimated on historical data either in the form of a VAR or a co-integrated model, does not provide superior exchange rate forecasts, compared to e.g. a random walk model, it can still make a lot of sense to rely on a joint exchange rate, yield curve and macro model, when doing joint portfolio optimization. In particular, such a framework would allow for the generation of stress tests and could provide decision makers with important information conditional upon given expectations they may have to one or more of the variables included in the model.

It is clear that it is an empirical challenge to integrate exchange rates into a strategic asset allocation framework applicable to fixed-income portfolios. Above we have outlined a few ideas as to how one may generate return distributions for domestic and foreign bonds as they are relevant for SAA. However, we acknowledge that this is only an appetizer: fundamental questions still remain to be answered.

4.5 Instrument pricing and calculation of returns

This section outlines well-known fixed-income mathematics relevant for calculating bond prices and returns. The process described above is quite general in that it allows for the projection of yield curves, exchange rates and

the credit state of given issuer/bond indices. This facilitates the calculation of bond returns comprising market as well as credit risk in the local currency, and by incorporating the exchange rate changes these returns can also be expressed in a base currency. It is naturally also possible to calculate expected return distributions originating from either of the unique risk sources, if that should be of interest, as it may be in a traditional market risk analysis in local currencies. The calculation formulas presented below are general and can be used in either of these situations.

Projected yield curves are translated into prices $P_{t,j}$ and returns expressed in local currency k, $R^k_{t,j}$ for the individual instrument classes $j = \{1, \ldots, J\}$, where j subsumes instruments that have pure market risk exposure as well as instruments that have both market and credit risk exposures. The maturity is denominated by $\tau_{t,j}$. The price of an instrument at time t is a function of the instrument's maturity, its coupon C and the prevailing market yield Y as it is observed at the maturity relevant for asset class j. The price can be written as

$$P_{t,j}(C, Y) = \frac{C}{Y}\left(1 - \frac{1}{(1+Y)^N}\right) + \frac{100}{(1+Y)^N}$$

where $C = C_{t-1,j}$ denotes the coupon, $N = \tau_{t,j}$ denotes the maturity and $Y = Y_{t,j}$ denotes the yield. It is important to note that $Y_{t,j}$ refers to the relevant credit yield-curve segment at time t for the relevant maturity segment. Finally, total gross returns in the local (foreign) currency k (R^k) for the instrument classes can be calculated as

$$R^k_{t,j} = \frac{P_{t,j}(\tau_{t,j}, C_{t-1,j}, Y_{t,j})}{P_{t-1,j}(\tau_{t-1,j}, C_{t-1,j}, Y_{t-1,j})} + C_{t-1,j}\Delta_t$$

where $C_{t-1,j}\Delta_t$ is the deterministic part of the return resulting from coupon payments. In the calculations it is assumed that at time t the portfolio is always rebalanced by replacing the existing bonds with instruments issued at par at time t, thus the coupon payments correspond to the prevailing yields at $t-1$.

The presented gross returns are expressed in local currency, whereas in a multi-currency framework, in which exchange rates are modelled, the relevant returns are expressed in a base currency. To transform these gross returns in local currency into gross returns in base currency one has to multiply gross returns with gross exchange rate returns (Ψ). Denoting by ξ_k the exchange rate quoted on a direct basis (Foreign/Home), then the

exchange rate gross return (Ψ) for currency k using domestic currency as the base currency from time $t-1$ to t will be

$$\Psi_t^k = \xi_t^k / \xi_{t-1}^k$$

and the gross return in base currency (R^b) will be

$$R_{t,j}^b = R_{t,j}^k \times \Psi_t^k$$

Typically the conversion of returns expressed in local currency to returns expressed in foreign currency will not be done at the instrument level, but at the portfolio or portfolio tranche level.

5. Optimization models for SAA under a shortfall approach

This section describes how to reach an optimal asset allocation using the inputs described above, i.e. most importantly the simulated return distributions. It is the premises that the investor is interested in a relatively long investment horizon, for which the return distributions are simulated, and that the objective function expresses aversion against losses. The formulations presented below should be relevant for many central banks who aim at avoiding annual investment losses in their reserves management operations (see IMF 2005).

The particular innovation of this section is to formulate the SAA problem as a multi-stage optimization problem without imposing any particular distributional form on the return distributions as opposed to a general one-period Markowitz optimization, and relying on a shortfall approach in which the objective function will be defined as either to minimize the risks or to maximize return subject to a given risk budget. Section 2.2.4 of this chapter presented the following discontinuous utility function for a shortfall/VaR approach:

$$U = u(\mathrm{r}, (\mathrm{C})\mathrm{VaR}(\alpha))$$

$$u(r, (C)VaR) = \begin{cases} \mathrm{r}, \text{if}\,(\mathrm{C})\mathrm{VaR}(\alpha) \geq 0 \\ (\mathrm{C})\mathrm{VaR}, \text{if}\,(\mathrm{C})\mathrm{VaR}(\alpha) < 0 \end{cases}$$

Maximization of this utility function can be directly specified in the optimization problem, or be split into two different objective functions which would apply under specific conditions. For pedagogical reasons, the latter option is presented below.

Section 5.1 describes an approach for SAA in a multi-currency environment where the objective is to minimize the risks inherent in holding the foreign reserve assets from a base currency perspective hence including exchange rate risk in the analysis. The main idea being that a central bank is unwilling to take active part in the FX markets, apart from if/when interventions are undertaken, but wants to diversify its investments to minimize annual losses. This sort of approach is preferred when integrating exchange rate risk in a shortfall type of analysis, since other formulations of the objective function and optimization constraints, as the one presented in Section 5.2 for a single currency, are not feasible given the dramatic increase in the shortfall risk figures arising from the inclusion of exchange rate risk.

Following the same lines, Section 5.2 presents a model for a single currency area in which the objective is to maximize returns in local currency subject to a no-loss constraint. This second approach is the one currently applied by the ECB when managing foreign reserves and own funds portfolios, since exchange rate risk is managed independently, mainly on the basis of policy considerations, and thus, it is difficult to assess whether the pay-off of the presented joint or multi-currency portfolio optimization compensate for the increase in model risk arising from the inclusion of exchange rate risk. Notwithstanding this, the so-called 'multi-currency model' is first presented, since it serves the purpose of illustrating a general portfolio optimization framework, where the single currency model can be seen as a special case. Also, the general specification presents an objective function that complements the single-currency objective function when certain market conditions and/or hard constraints on minimum holdings prevent the projected efficient frontier to lie in the area of the mean-shortfall space which is defined as feasible by the no-loss (VaR/shortfall) constraint.[15]

The SAA decision when set in a multi-period and multi-currency setting aims at finding an optimal currency allocation as well as an allocation between instrument classes within each currency area while taking into account global and local optimization constraints. For example, a central purpose of many reserve holdings is to facilitate central bank currency interventions and this may naturally impose some restrictions on the relative distribution of investments across currencies as well as on the liquidity and risk profile of the allocations within each local currency area.

[15] An example of a case in which complementary objective function is needed is presented in Section 6.5.

Such requirements have to be accounted for in the formulation of the central bank's risk preferences used to derive the optimal asset and currency allocation. For instance, the ability to liquidate quickly and with minimum price impact a large proportion of the assets held is an important requirement for an intervention portfolio. In terms of optimization constraint, this requirement imposes a minimum portfolio allocation to highly liquid instruments.

A consistent risk measure that accounts for the concerns of central banks in relation to foreign reserve holdings is Conditional Value at Risk. This measure allows decision makers to express directly their aversion against losses of a certain magnitude at a given confidence level. It is shown below that CVaR also has an interpretation in terms of regular Value at Risk. Hence, relying on CVaR easily generalizes to cases where decision makers express their risk aversion in the form of a VaR number at a given confidence level.

By definition the objective of avoiding annual losses at some confidence level α can be related to the concept of Value at Risk (VaR). Let $R \in (0, \infty)$ denote a random growth rate of an asset, i.e. $R = 1 + r$ with r being the return, such that $R < 1$ represents a loss and $R > 1$ represents a gain, let F_R be its cumulative distribution function, i.e. $F_R(u) = \Pr(R \leq u)$ and let $F_R^{-1}(v)$ be its inverse i.e. $F_R^{-1}(v) = \inf\{u : F_R > v\}$. For a specified confidence level α, define the VaR as the $1 - \alpha$ quantile, i.e. $VaR_{1-\alpha}(R) = F_R^{-1}(1 - \alpha)$. For example, the objective of avoiding losses at a 99 per cent confidence level would be formulated as $VaR_{0.01}(R) \geq 1$.

VaR is widely used as a risk measure in the finance industry, but it has several drawbacks which limit its use as a risk measure in portfolio optimization applications (see e.g. Rockafellar and Uryasev 2000 and Krokhmal *et al.* 2002). For example, since VaR is defined as the loss at a given confidence level it does not give information about the losses beyond the chosen confidence level. Losses beyond VaR may be important especially if the loss distribution is non-normal, e.g. fat-tailed or skewed, as it is often the case for returns on financial assets. Theoretical shortcomings of VaR are (i) it is not sub-additive, i.e. diversification among financial assets may actually increase VaR, i.e. increase the risk measured, rather than decrease it as conventional portfolio theory would suggest; (ii) it is in general non-convex, which causes great practical difficulties in optimization applications due to possibly multiple local optima. These mentioned shortcomings of VaR are not shared by the closely related risk measure CVaR.

The definition of CVaR, for a specified confidence level α and a probability measure P is

$$CVaR_\alpha(R) = E^P(R : R \leqslant VaR_{1-\alpha}(R))$$

That is, $CVaR_\alpha$ equals the expected tail return below $VaR_{1-\alpha}(R)$, i.e. the expected return in the worst $(1-\alpha)$*100% of cases. See e.g. Rockafellar and Uryasev (2000) and Pflug (2000) for a detailed discussion of the theoretical and computational advantages of *CVaR* compared to *VaR* as a risk measure. *CVaR* can equivalently be defined as the solution of an optimization problem (Rockafellar and Uryasev 2000):

$$CVaR_\alpha(R) = \sup_b\{b - \frac{1}{1-\alpha}E^P[\max(b-R,0)]\}$$

where the optimal b is the $VaR_{1-\alpha}(R)$. A desirable feature of this formulation is that when the continuous probability measure P is approximated by a discrete sample of realizations from P, which is usually needed for numerical solution of the problem, it can be expressed as a system of linear constraints, see Rockafellar and Uryasev (2000). This makes *CVaR* an attractive risk measure for portfolio optimization applications and will be used to formulate the risk preferences in the following sections. In addition, this way of formulating the optimization problem ensures that the tail constituting the losses has properties that are in accordance with decision makers' wishes. In particular, it is ensured, that exactly α per cent of the probability mass is in the tail, while at the same time the *VaR* level is maximized (i.e. the least negative value of b is chosen).

An intuitive way to think about the above linkage between *VaR* and *CVaR* is that *VaR* is the quantile where a certain probability mass is ensured to be observed in the tail (in this case the left tail of the return distribution R). Hence, it is possible to optimize over the *VaR* level, denoted by b, until exactly α per cent of the distribution is in the tail, and once the correct level of b is found, then the *CVaR* can be found as the average losses that fall in the tail. As for the last term on the right-hand side in the above equation, i.e. the definition of the tail beyond (b):

$$\frac{1}{1-\alpha}E^P[\max(b-R,0)]$$

it is noted that the tail is defined as $\max(b-R, 0)$, which means that, going from the right tail to the left tail (i.e. from gains to losses), all returns are

given a zero until the level of b is reached and then afterwards observations are allocated the value b-R. Hence, the expectation is calculated as the original return minus the VaR return (b), and to get the result right, the VaR return (b) is then added again.

5.1 Multi-currency model

Consider a multi-stage portfolio optimization problem where decisions are interlinked with realizations of random variables. The investment decisions at each time stage $t = 0,1, \ldots, T$, where $t = 0$ denotes the present time, are based on the revealed information up to that point in time. This kind of interdependence of decisions and information is typical for sequential decision making under uncertainty.

Let $k \in \{1, \ldots, K\}$ denote the set of modelled currency areas with K denoting the base currency and the set of available asset classes in currency area k is indexed by $j = 1, \ldots, J(k)$. The following summarizes the used notation:

Random variables:

$\Psi_t^k = \frac{\xi_t^k}{\xi_{t-1}^k}$ = exchange rate return measured in base currency against foreign currency $k = 1, \ldots, K-1$ over period $[t-1, t]$, where $t = 1, \ldots, T$,

$R_{t,j}^k$ = total gross return/growth rate of asset class $j = 1, \ldots, J(k)$ invested in currency $k = 1, \ldots, K-1$ over period $[t-1, t]$, expressed in local currency k.

Decision variables:

$x_{t,j}^k$ proportion of wealth invested in asset class $j = 1, \ldots, J(k)$ in currency $k = 1, \ldots, K-1$ at time $t = 1, \ldots, T$,

η^k = proportion of wealth invested in currency $k = 1, \ldots, K-1$,

$b_t^k =$ cut-off point for CVaR at time $t = 1, \ldots, T$, in currency $k = 1, \ldots, K$.

Deterministic parameters:

$c_j^k =$ portfolio update limit for asset class $j = 1, \ldots, J(k)$ in currency $k = 1, \ldots, K-1$,

$l_j^k, u_j^k =$ lower and upper bounds for the proportion invested in asset class $j = 1, \ldots J(k)$ in currency $k = 1, \ldots, K-1$,

$\tilde{l}^k, \tilde{u}^k$ lower and upper bounds for the proportion invested in currency $k = 1, \ldots, K-1$,

α^k = confidence level for $CVaR$ in currency $k = 1, \ldots, K$.

As argued above, the decision maker's utility function is specified in terms of CVaR with the objective to minimize expected annual risk expressed as $CVaR_{\alpha^k}$ over the investment horizon measured in base currency at a confidence level α,

$$\max_{x,b,\eta} \sum_{t=1}^{T} (b_t^K - \frac{1}{1-\alpha^K} E^P[\max(b_t^K - \sum_{k=1}^{K-1} \Psi_t^k \sum_{j=1}^{J(k)} x_{t-1,j}^k R_{t,j}^{i,k}, 0)], \tag{2.9}$$

where P is the probability measure of the random variables and E^P denotes the expectation operator with respect to the probability measure P. Here the currency composition and co-dependence between bond returns and exchange rate returns play an important role in shaping the combined return distribution from a base currency perspective. Additional constraints that guard against annual losses in local currencies can be imposed by defining the following set of restrictions

$$\sup_b \{b_t^k - \frac{1}{1-\alpha^k} E^P[\max(b_t^k - \sum_{j=1}^{J} x_{t-1,j}^k R_{t,j}^k, 0)]\} \geq \eta^k, \tag{2.10}$$

$t = 1, \ldots, T,\ k = 1, \ldots, K-1.$

As mentioned above, a desirable feature of *CVaR*, which makes it a very attractive risk measure for portfolio optimization, is that when the continuous probability measure P is approximated by a discrete sample of realizations from P, equations (2.9) and (2.10) simplify to a linear objective function with a system of linear constraints, see Section 5.1.1. Policy constraints are easily included in the formulation. Such possible constraints are outlines below:

- Transaction cost limits that keeps the portfolio turnover in period $[t-1, t]$ under a specified tolerance level c_j^k may be relevant, i.e.:

$$\left| x_{t,j}^k - x_{t-1,j}^k \right| \leq c_j^k, t = 0, \ldots, T-1, j = 1, \ldots, J(k), k = 1, \ldots, K-1.$$

- Bounds on the currency weights can be formulated as:

$$\tilde{l}^k \leq \eta^k \leq \tilde{u}^k, k = 1, \ldots, K-1,$$

- Limits on the portfolio shares within each currency k can be expressed as,

$$l_j^k \eta^k \leq x_{t,j}^k \leq u_j^k \eta^k, t = 0, \ldots, T-1, j = 1, \ldots, J(k), k = 1, \ldots, K-1.$$

- Asset class specific bounds to account e.g. for liquidity issues can be written as

$$\hat{l}_j^k \eta^k \leq \sum_{j \in B(k)} x_{t,j}^k \leq \hat{u}_j^k \eta^k, t = 0, \ldots, T-1, k = 1, \ldots, K-1.$$

where $B(k) \subset \{1, \ldots, J(k)\}$ is some subset of the available assets in currency area $k=1, \ldots, K$.

- As a matter of definition, it is required that the portfolio shares within each currency sum up to the currency share, i.e.

$$\sum_{j=1}^{J(k)} x_{t,j}^{k} = \eta^{k}, t = 0, \ldots, T-1, k = 1, \ldots, K-1$$

- And, finally it is required that the static currency weights sum up to one:

$$\sum_{k=1}^{K} \eta^{k} \leq 1$$

In the presented model formulation it is assumed that the decisions taken are non-anticipative, which means that the decision variables of a given time stage cannot depend on the random variables whose values are observed only in later stages. The evolution of the random variables is described with continuous (multivariate) probability distribution generated by the model presented in Section 4. For such a (general) probability distribution the analytical computation of the expectation in (2.9) is practically impossible and therefore, the presented optimization problem has to be solved numerically by discretizing the continuous probability distribution P. The discrete approximation of P is generated by sampling a set of scenarios from the specified density and the sample approximation of the original problem is then solved using numerical optimization techniques. A convex deterministic equivalent formulation of the above problem is presented in the following section.

5.1.1 Discretization

In order to solve the optimization model presented above, the probability distribution P of the random variables has to be discretized and the resulting problem solved numerically. This can be done by generating N sample paths of realizations for the random variables spanning the time stages $t = 1, \ldots, T$, as also mentioned above. Each simulated path reflects a sequence of possible outcomes for the random variables over the investment horizon and the collection of sample paths gives a discrete approximation of the probability measure P. For a discretized probability measure the objective (2.9) can be formulated as a combination of a linear objective

function (2.11) and linear inequalities (2.12). The loss constraints in local currencies (2.10) can be replaced with a system of linear restrictions (2.13)–(2.14) (Rockafellar and Uryasev 2000) in the formulation below. This results in a linear stochastic optimization problem

$$\max_{x,z,b,\eta} \sum_{t=1}^{T} \left(b_t^K - (1-\alpha^K)^{-1} \sum\nolimits_{i\in\nu_t} p^i z_t^{i,K}\right) \tag{2.11}$$

subject to

$$z_t^{i,K} \geq b_t^K - \sum_{k=1}^{K-1} \Psi_t^k \sum_{j=1}^{J(k)} x_{t-1,j}^k R_{t,j}^{i,k}, \;\; t = 1,\ldots,T, i = 1,\ldots,N, \tag{2.12}$$

$$b_t^k - (1-\alpha^k)^{-1} \sum\nolimits_{i\in\nu_t} p^i z_t^{i,k} \geq \eta^k, \quad t = 1,\ldots,T,\; k = 1,\ldots,K-1 \tag{2.13}$$

$$\begin{aligned} z_t^{i,k} &\geq b_t^k - \sum_{j=1}^{J(k)} x_{t-1,j}^k R_{t,j}^{i,k}, \quad t = 1,\ldots,T, i = 1,\ldots,N, \\ & k = 1,\ldots,K-1 \end{aligned} \tag{2.14}$$

$$\sum_{j=1}^{J(k)} x_{t,j}^k = \eta^k, t = 0,\ldots,T-1, k = 1,\ldots,K-1, \tag{2.15}$$

$$\sum_{k=1}^{K} \eta^k \leq 1, \tag{2.16}$$

$$\begin{aligned} -c_j^k \eta^k \leq x_{t,j}^k - x_{t-1,j}^k \leq c_j^k \eta^k, \quad & t = 1,\ldots,T-1, k = 1,\ldots,K-1, \\ & j = 1,\ldots,J(k), \end{aligned} \tag{2.17}$$

$$\tilde{l}^k \leq \eta^k \leq \tilde{u}^k, k = 1,\ldots,K-1, \tag{2.18}$$

$$\begin{aligned} l_j^k \eta^k \leq x_{t,j}^k \leq u_j^k \eta^k, \;\; & t = 0,\ldots,T-1, k = 1,\ldots,K-1, \\ & j = 1,\ldots,J(k), \end{aligned} \tag{2.19}$$

where $z_t^{i,k}$ are dummy variables, p^i is the probability of scenario *i*. The decision variables in the model only depend on time (i.e. they are scenario independent) and therefore the solution will be non-anticipative even

though a tree structure is not used for describing the evolution of the random variables. If constraints (2.13) are active at the optimum, the corresponding optimal value b_t^k will equal the $VaR_{t,1-\alpha}^k$ and the left-hand side of (2.13) will be equal to $CVaR_{t,\alpha}^k$ for $k=1,\ldots,K-1$ and $t=1,\ldots,T$ (for details see e.g. Rockafellar and Uryasev 2000). Constraint (2.15) restricts the sum of portfolio weights within each currency to equal the share of that currency in the portfolio and (2.16) ensures that the currency weights sum up to one. Constraint (2.17) defines the annual portfolio updating limits for each asset class and (2.18) –(2.19) give the lower and upper bounds for the weights of individual currencies and portfolio shares within each currency, respectively.

A fixed-mix solution can be found if the turn-over constraints (the c_j) are set to zero from one period to the next. In this case the asset weights stay constant over the investment horizon.

5.2 Single market model

This section describes a model for SAA in a single currency setting and as such it presents a reduced version of the above general multi-currency model formulation.

The objective here is to find a dynamic decision strategy that maximizes the value of the portfolio at the end of the planning horizon, while complying with the decision makers risk–return preferences as well as other policy constraints when applicable.

For most parts the model follows the formulation presented in the previous section. The objective of the model, to maximize the expected wealth at the end of the investment horizon can be written as

$$\max_{x,W,b} E^P W_T \tag{2.20}$$

where P is the probability measure of the random variables, E^P denotes the expectation operator and the development of the portfolio value (wealth) over period $[t-1,\ t]$ is given by

$$W_t = W_{t-1} \sum_{j=1}^{J} x_{t-1,j} R_{t,j}, \quad t = 1,\ldots,T \tag{2.21}$$

Maximizing (2.20) subject to (2.21) and constraints (2.13) –(2.19) applied to a single currency case results in the following discretized stochastic optimization model:

$$\max_{x,z,b,W} \sum_{i \in N} p^i W_T^i \tag{2.22}$$

subject to

$$b_t - (1-\alpha)^{-1} \sum\nolimits_{i \in \nu_t} p^i z_t^i \geq 1, \quad t = 1, \ldots, T \tag{2.23}$$

$$z_t^i \geq b_t - \sum_{j=1}^{J} x_{t-1,j} R_{t,j}^i, \quad t = 1, \ldots, T, \ldots, N \tag{2.24}$$

$$\sum_{j=1}^{J} x_{t,j} = 1, \quad t = 0, \ldots, T-1 \tag{2.25}$$

$$W_t^i = W_{t-1}^i \sum_{j=1}^{j} x_{t-t,j} R_{i,j}^i, \quad t = 1, \ldots, T, \; i = 1, \ldots, N \tag{2.26}$$

$$-c_j \leq x_{t,j} - x_{t-1,j} \leq c_j, \quad t = 0, \ldots, T-1, j = 1, \ldots, J \tag{2.27}$$

$$l_j \leq x_{t,j} \leq u_j, \; t = 0, \ldots, T-1, j = 1, \ldots, J \tag{2.28}$$

where the *CVaR* constraints against periodic losses are defined by a system of linear restrictions (2.23)–(2.24), z_t^i are scenario-dependent dummy variables and p^i is the probability of scenario *i*. If constraint (2.23) is active at an optimal solution, the corresponding optimal value b_t will equal the $VaR_{t,1-\alpha}$ and the left-hand side of (2.23) will be equal to $CVaR_{t,\alpha}$ for stage *t*. Constraint (2.25) ensures that the sum of the portfolio weights equals one and the portfolio wealth at time *t* in scenario *i* is expressed by (2.26). Constraint (2.27) specifies the annual portfolio updating limits for each asset class and the lower and upper bounds for the portfolio shares are given by (2.28).

A fixed-mix solution can be found if the turn-over constraints (the c_j's) are set to zero from one period to the next. In this case the asset weights stay constant over the investment horizon.

6. The ECB case: an application

This section presents examples of how the techniques outlined above are used within the ECB to provide information that can aid senior management in making SAA decisions. The examples are only illustrative and should neither be taken to represent concrete investment advice nor as an 'information package' endorsed by the ECB. Rather the examples show hypothetical empirical application of the methodology advocated in the above sections.

The next section describes the investment universe; Section 6.2 presents the objective function; Section 6.3 elaborates on how the models presented in other sections are used in practise and describes some details about the specification and parameters of those models as have been used in the examples, Section 6.4 shows an application to a realistic scenario that is labelled as normal, due to the expected evolution of macroeconomic variables and the starting yield curve, and Section 6.5 shows an application to a non-normal scenario, presenting an inflationary economic situation and a starting yield curve close to the historically lowest levels observed in the US. Those scenarios have been chosen, instead of a single normal scenario, to better illustrate the effect of the starting yield curves and the projected evolution of the macroeconomic variables have on the SAA decisions support information generated by the SAA framework.

6.1 The investment universe

The definition of the eligible investment universe may to some extent reflect the objective of holding reserves. In the case where reserves are held for intervention purposes the overriding principles for the holdings will be security and liquidity and hence only very conservative investment vehicles should constitute the available investment universe. The following list of instruments could form such an eligible investment universe:

- government bonds
- agencies with government support
- BIS instruments
- bonds issued by supranational organizations
- deposits.

For modelling purposes the eligible investment universe may be subdivided into maturity buckets. A representative maturity for each bucket will be

Table 2.1 Example of the eligible investment universe for a USD portfolio

Asset class	Maturity segment	Mnemonic
Government	0–1 years	US Gov 0–1Y
	1–3 years	US Gov 1–3Y
	3–5 years	US Gov 3–5Y
	5–7 years	US Gov 5–7Y
	7–10 years	US Gov 7–10Y
Spread products	0–1 years	US Sprd 0–1Y
	1–3 years	US Sprd 1–3Y
	3–5 years	US Sprd 3–5Y
	5–7 years	US Sprd 5–7Y
	7–10 years	US Sprd 7–10Y
Cash/Depo	1 month	US Depo 1M

used in the Nelson–Siegel model to project the evolution of yields and calculate the return distributions. Table 2.1 gives an example of how such an investment universe could look like for the US dollar portfolio, and this is actually what is used in the current examples below. It is important to note that Table 2.1 shows only the first step of the process of determining the strategic benchmark. After the optimization exercise is carried out and optimal exposures are determined to each of the generic 'indices' shown below, each of these exposures are translated into actually tradable bonds facilitating 100 per cent replication of the benchmark allocation.

6.2 The objective function and constraints

The objective function and the constraints serve as a means to translate the long-term preferences of the senior management into an actual allocation. In this context the following formulation may be used. It is the objective to maximize returns in local currency (USD for a US dollar portfolio) subject to the liquidity and security of the holdings. Hence, portfolio return is maximized under the constraints that (a) there are no losses at a 99.5 per cent confidence level over a one-year horizon, where the no-loss constraint is expressed in terms of CVaR; (b) there is a minimum holding of government bonds in the portfolio, e.g. no less than 50 per cent of the portfolio must be held in government bonds; (c) there is a maximum holding allowed of certain instruments, e.g. deposits of e.g. 10 per cent; (d) relative market capitalization is used to impose minimum holdings per time bucket in each

instrument category, determining 60 per cent of the portfolio composition. The only relevant risk for the portfolio optimizer is interest rate risk, since the credit risk and exchange rate modules presented in Sections 4.3 and 4.4 are not used to generate the returns that will feed the portfolio optimizer in the current example.

6.3 Using the models

In a decision-making process, as the one applied by the ECB, there is a lag between the date that marks the beginning of the analysis period, i.e. the date after which no new information/data can be added to perform the relevant model/parameter estimations, and the date when the optimized portfolio allocation is implemented in practise. This lag is accounted for by not only modelling the evolution of the different variables through the investment horizon, but also during the lag period, effectively introducing a certain amount of uncertainty even prior to the beginning of the relevant investment horizon. In the examples presented in the next sections, the relevant investment horizon starts at the end of year X and finishes after one year, i.e. in year X + 1. It is assumed that the analysis is conducted using data from the end of September of year X.

A VAR model as the one presented in Section 4.1 is used to project year-over-year monthly observed growth rates for US GDP and CPI, given the available forecasts for these variables. First, average trajectories for the GDP and CPI growth rates are projected over the investment horizon. Then, using the covariance matrix of the residuals of the macroeconomic VAR, 10,000 contemporary shocks on both variables are generated for each month along the forecast horizon, by sampling from a multivariate normal distribution. These shocks are cumulated using the estimated auto-regressive process and added to the average trajectories as explained in Section 4.1.

A dynamic Nelson–Siegel model as presented in Section 4.2 is used to project the evolution of the different yield curves conditional to the macroeconomic realizations. These realizations are classified as Normal (including stagflation), Recession or Inflation using the classification scheme shown in Table 2.2.

The transition matrices (p) for the different yield-curve regimes conditional to the classification of the different macroeconomic environments are presented in Table 2.3.

The slope and curvature have been transformed as described in Box 2.1 of Section 4.2 to be expressed in relative terms to the level factor.

Table 2.2 Classification scheme

Macro variable	GDP YoY growth (%)	CPI YoY growth (%)
Inflation	>1.75	>3.75
Recession	<1.75	<3.75
Normal / Stagflation	Other combinations	

Table 2.3 Transition matrices

Macroeconomy	Normal (p^1)			Recession (p^2)			Inflation (p^3)		
Regime	Normal (t)	Steep (t)	Inverse (t)	Normal (t)	Steep (t)	Inverse (t)	Normal (t)	Steep (t)	Inverse (t)
Normal (t + 1)	1.00	0.05	0.05	0.95	0.00	1.00	0.95	1.00	0.00
Steep (t + 1)	0.00	0.95	0.00	0.05	1.00	0.00	0.00	0.00	0.00
Inverse (t + 1)	0.00	0.00	0.95	0.00	0.00	0.00	0.05	0.00	1.00

Table 2.4 Intercepts of the Nelson–Siegel state equation

	Regime		
Factor	Normal	Steep	Inverse
Level	0.06	0.045	0.065
Slope	− 0.024	− 0.048	0
Curvature	0.002	−0.05	0.002

The intercepts (C) and the autoregressive coefficients (F) corresponding to the US Government Yield Curve in the state equation of the Nelson–Siegel model used in this example can be seen in Table 2.4 and Table 2.5.

The parameter *lambda* of the Nelson–Siegel model has been assumed to have a value of 0.0687. The presented model parameters, after reversing the transformation of the slope and curvature factors, imply the generic *normal*, *steep* and *inverse* yield curves for the US government instruments that can be observed in Figure 2.8.

The probability of switching to (or staying in) a normal yield-curve regime under a normal economic environment converges in the limit to 100 per cent, and so, after a sufficiently long period of normal economic growth and

Table 2.5 Autoregressive coefficients of the Nelson–Siegel state equation

Factor	Level	Slope	Curvature
Level	0.99	0.00	0.00
Slope	0.00	0.92	0.00
Curvature	0.00	0.00	0.90

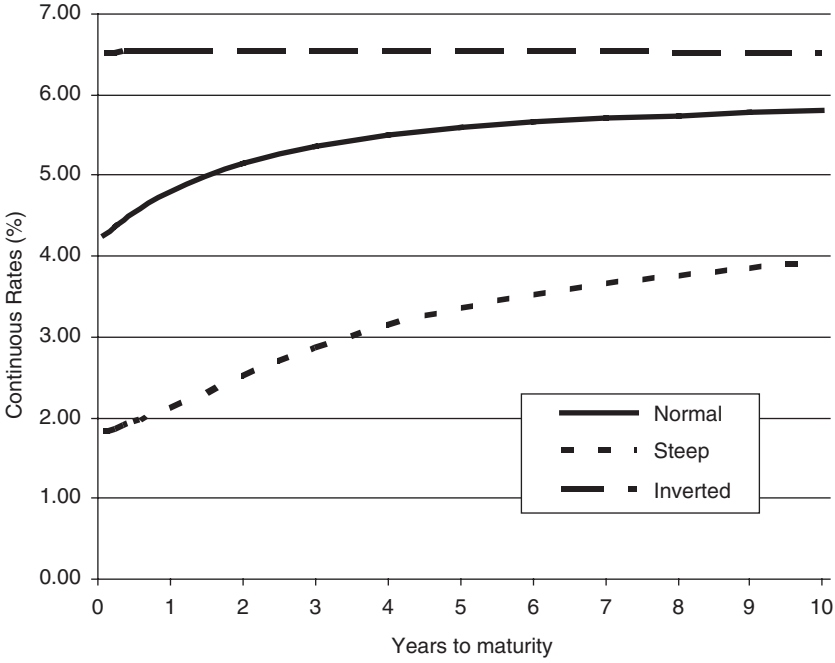

Figure 2.8 Generic yield curves.

inflation, the yield curve would be expected to converge towards the generic normal curve. A persistent recessionary period would make the probability of switching to a steep yield-curve regime converge to 100 per cent, and consequently, the yield curve will be expected to move towards the generic steep curve. After a long period of inflation, the probability of switching to an inverted yield-curve regime will converge to 100 per cent, and the yield curve will be expected to move accordingly towards the generic inverted curve.

As it has been shown, the different evolution of the macroeconomic scenarios imply a diverse evolution of the state probabilities used to weight the intercepts corresponding to each yield-curve regime in the Nelson–Siegel state equation.

Using the covariance matrix of the residuals of the estimated VAR process for the Nelson–Siegel factors, 10,000 contemporary shocks on the factors are generated for each month along the forecast horizon, sampling from a multivariate normal distribution. Besides these shocks, additional noise has been added to the simulation by modelling the error-terms of the Nelson–Siegel observation equation. If a given simulation run produces negative yields at any maturity the scenario is discarded and replaced by a new one.

Uncertainty is introduced at the level of the evolution of the macroeconomic variables, at the level of the evolution of yield-curve factors and at the level where yield-curve factors are translated into the actual yield curves. Introducing such uncertainty allows the analyst to generate realistic yield-curve scenarios that facilitate stochastic portfolio optimization.

Based on the yield-curve projections it is possible to calculate expected returns for the generic instrument classes in Table 2.2, in the fashion described in Section 4.5. Returns are expressed in local currency, i.e. in USD since this example presents a USD portfolio.

Different baseline scenarios should be investigated in order to provide decision makers with a full picture of possible future realizations of the world. Naturally, some of these scenarios may be defined by the decisions makers.

On the basis of the summary information as well as a detailed account of how scenarios are generated and in-depth analysis of what each scenario implies in terms of adherence to the risk–return preferences of the organization in question, the decision makers can then decide on the optimal asset allocation for the coming period.

6.4 Application to a normal scenario

6.4.1 Macroeconomic scenarios and starting yield curves

As a first example we present a scenario in which the yield curve stays very close to historically normal levels although relatively flatter, and is assigned a 100 per cent probability of belonging to the Normal regime. The expected evolution of the macroeconomic variables (GDP and CPI growth) can also be seen to represent a 'normal' economic environment, since there are no inflationary or recessionary pressures.

Using such expected evolutions as a baseline scenario, 10,000 simulations are used to generate the needed variable distributions. These distributions are shown in Figures 2.9a and 2.9b, where shades of grey represent the simulated probability density and the darker areas represent a higher

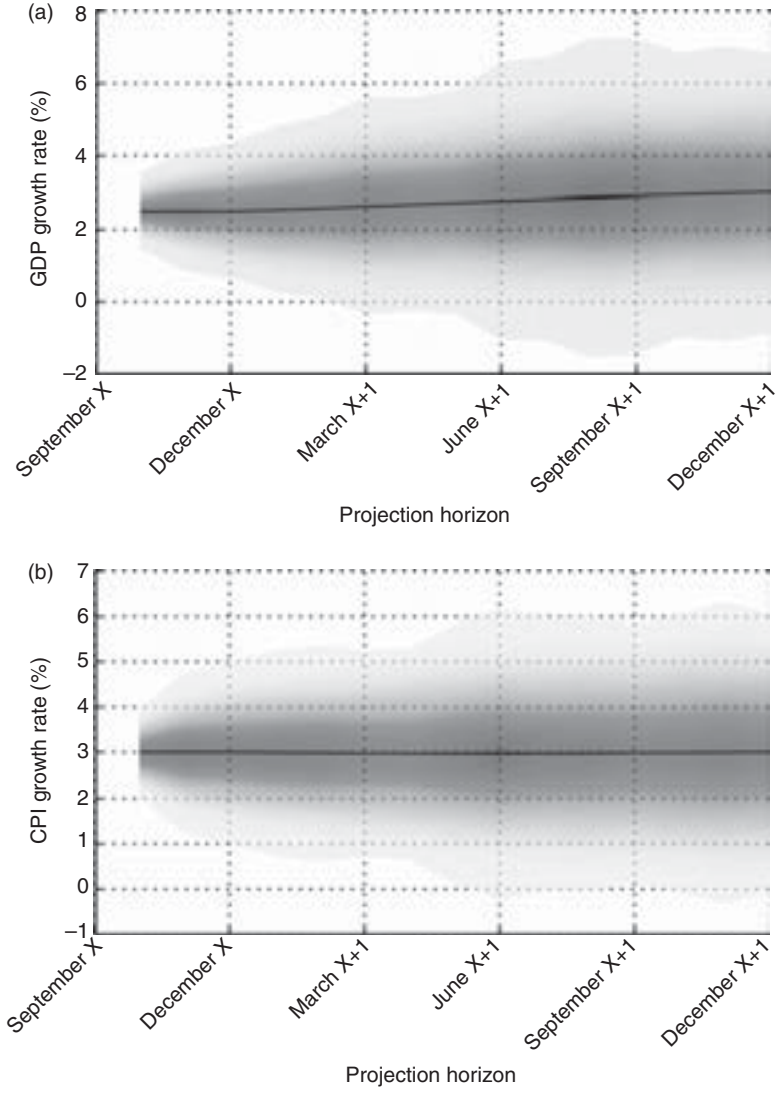

Figure 2.9 Normal macroeconomic evolution: (a) GDP YoY % Growth; (b) CPI YoY % Growth.

probability density for GDP growth (a) and CPI growth (b). The black line reflects the baseline or average evolution.

6.4.2 Yield-curve projections and expected returns

Using the simulated densities for the macro variables as input to the yield-curve modelling framework described above, allows us to derive simulated

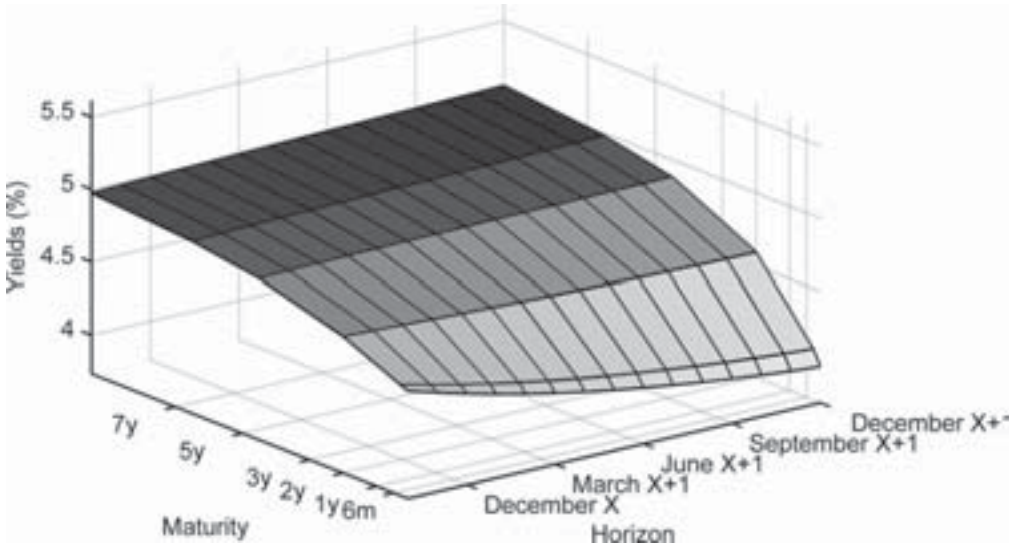

Figure 2.10 Projected average evolution of the US Government yield curve in a normal example.

distributions for the projected evolution of yield curves, starting from its shape and location at the time the projection is made to the end of the projection horizon. This is illustrated in Figure 2.10 for the US Government yield curve. Since the starting yield curve and most of the projected macroeconomic scenarios can be classified as normal, no drastic changes in the yield-curve shape/location are expected, and consequently only a slight and smooth steepening of the curve is projected on average.

It is worth noting that Figure 2.10 shows only the average yield path, i.e. the average across all 10,000 simulated yield-curve evolution paths. To gain additional insight into the simulated yield-curve distributions along the projection horizon, Figure 2.11 presents two example plots for the US Gov 0–1Y (Figure 2.11a) and for the US Gov 7–10Y (Figure 2.11b) indices.

The projected evolution of the yield curve permits us to compute the returns for those indices over the relevant investment horizon (from December X to December X + 1 in this example) and thus to generate return distributions. Table 2.6 illustrates the summary return statistics for the different indices in this example.

It is shown how the spread products outperform their maturity-matching Government products, but at the price of a higher volatility arising from the spread risk (pure credit risk has not been taken into account). It is also shown how the 1–3 segment of the curve, although it is not the more risky

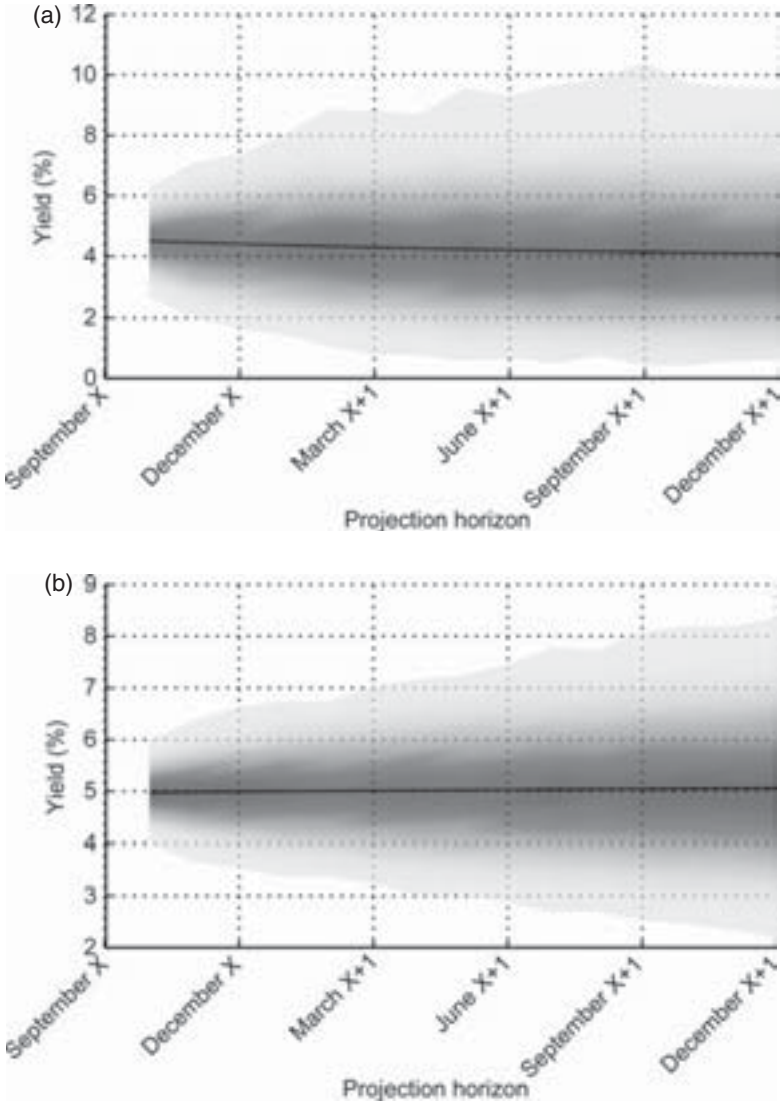

Figure 2.11 Projected distribution of yields in a normal example: (a) US Gov 0–1Y; (b) US Gov 7–10Y.

segment, is expected to be the best-performing segment, due to the projected slight steepening of the curve. The possibility of finding indices or instruments performing better than other more risky ones, can be observed even when facing a normal economic environment, particularly when the investment horizon is shorter than the average length of a business cycle.

Table 2.6 Returns in a normal example: average and standard deviation

Asset class	Maturity segment	Average (%)	Standard deviation (%)
Government	0–1 years	4.47	0.87
	1–3 years	5.13	1.93
	3–5 years	5.12	3.53
	5–7 years	4.96	4.51
	7–10 years	4.89	5.76
Spread Products	0–1 years	4.88	0.87
	1–3 years	5.38	1.96
	3–5 years	5.37	3.64
	5–7 years	5.21	4.79
	7–10 years	5.13	6.38
Cash/Depo	1 month	4.50	1.01

However, in the long run the so-called 'first law of finance' (*the higher the risk, the higher the expected return*) will on average hold, since the capital losses (gains) coming from increasing (decreasing) yields in the short run will be compensated by the coupon effect in the medium and long run, and because yield-curve movements follow the business cycle; and so, a steepening today may be followed in the future by a flattening of the curve.

Another summary statistic worth mentioning is the dispersion of the return distributions corresponding to the Cash/Depo asset class, which is higher than the US Sprd 0–1Y, although the maturity and duration of Cash is lower and both indices have been projected as being priced off the same (spread) curve. There are two explanations for this fact: first, since the investment horizon is one year, the annual return for an index with a maturity of one month may be more volatile than that corresponding to an index with a maturity of six months, which is the average maturity for the US Sprd 0–1Y index; and second, the last source of uncertainty induced in the yield-curve model serves the purpose of introducing some specific risk other than the risk arising from the evolution of the Nelson–Siegel factors. This specific risk, which is modelled after the perturbation term in the observation equation of the Nelson-Siegel model, has been parameterized as higher for the US Depo 1M index than for the US Sprd 0–1Y.

The presented returns, together with a covariance matrix, could serve as the input for a Markowitz optimization. However, since the preferred risk measure of the ECB is not volatility, but rather a tail-risk measure such as VaR and CVaR, we are also interested in other features (moments) of the

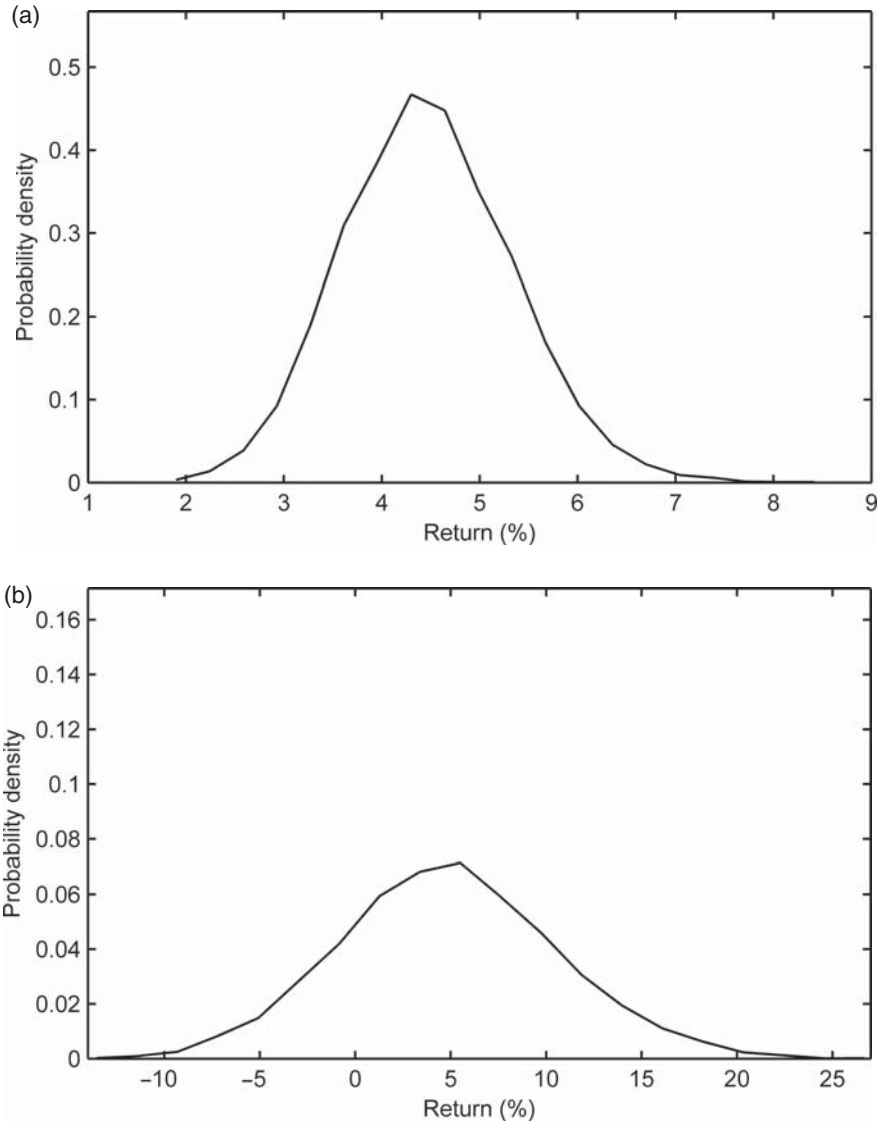

Figure 2.12 Distribution of returns in a normal example: (a) US Gov 0–1Y; (b) US Gov 7–10Y.

return distribution, such as skewness and kurtosis. To illustrate why these features may be relevant for a fixed-income investor, the asymmetric and leptokurtic distribution of the simulated returns for the US Gov 0–1Y index is presented in Figure 2.12a and the platykurtic and disperse distribution of returns for the US Gov 7–10Y in Figure 2.12b.

Table 2.7 Optimal portfolio composition in a normal example

Asset class	Maturity segment	Portfolio weights (%)
Government	0–1 Years	8
	1–3 Years	32
	3–5 Years	4
	5–7 Years	2
	7–10 Years	4
Spread Products	0–1 Years	21
	1–3 Years	15
	3–5 Years	3
	5–7 Years	3
	7–10 Years	2
Cash/Depo	1 Month	6

Table 2.8 Summary information for the optimal portfolio in a normal example

Expected Return	5.02%
Modified Duration	1.9
Volatility	2.40%
VaR (99.5%)	0.53%

6.4.3 Optimal portfolio allocations

Based on the simulated returns it is possible to apply the stochastic portfolio optimization technique. Applying this portfolio optimizer produces an optimal exposure to each generic index as illustrated in Table 2.7.

The asset allocation in this case is basically determined by the expected returns, and the imposed constraints on minimum holdings in Government instruments and minimum holdings relative to the market capitalization. The risk budget is not completely consumed, since the portfolio that maximizes expected return is feasible without the need of taking the maximum amount of risk permitted. Some summary information for this portfolio is presented in Table 2.8.

A quick comparison between this summary information and the statistics for the different indices show how this portfolio composition may seem sub-optimal, since it presents a lower expected return for a higher volatility[16]

[16] A direct comparison of those tables is not recommended anyway, since the standard deviation of the returns of the different indices is a measure of the dispersion of the annual simulated returns at the end of the forecasting period, while the volatility of the portfolio has been computed as the average volatility of the different simulated time-series of portfolio returns, taking monthly returns and annualizing them. This second measure is closest to the standard

than, e.g. the US Gov 1–3 index. This sub-optimality is in this case the price to pay for getting a smooth allocation among different indices, i.e. the cost of the holdings relative to market capitalization constraint. An institution may be willing to pay this sort of price to increase the stability of the strategic benchmarks in terms of asset allocation and modified duration. If these considerations are seen as part of the utility function of the institution, although they will typically take the form of constraints in the optimization problem instead of being an explicit part of the objective function, the constrained portfolio should then be considered as optimal.

6.5 Application to a non-normal scenario

6.5.1 Macroeconomic scenarios and starting yield curves

This second example presents an extreme scenario, where the economy faces inflationary pressure with the yield curve departing from historically low levels. The probability of the starting yield curve belonging to a steep curve regime is assessed as being equal to 100 per cent. The expected evolution of the macroeconomic variables (GDP and CPI growth) shows an inflationary scenario, in which the GDP forecasts do not show signs of risk to the growth (i.e. the scenario does not reflect a stagflation environment).

Using this expected evolution as a baseline, 10,000 simulations are run as in the previous example. These distributions are seen in Figures 2.13a and 2.13b, in which the lightest areas represent lower density (number of simulations), the darkest areas represent a higher density and the black line traces the baseline or average evolution.

6.5.2 Yield-curve projections and expected returns

The inflationary macro scenario results in an increase in the probability of observing flat yield curves, and consequently, a transition is projected from a relatively steep curve located at a very low level to a flatter curve located at a higher level. What is shown in Figure 2.14 represents the mean path of the yield curve. By introducing uncertainty into the picture, less smooth transitions are generated.

A better insight into the distribution of the evolution yields along the projection horizon is presented in Figure 2.15, containing the example plots

notion of volatility, since it is based in the evolution of returns in each scenario, rather than in the dispersion of different realizations under different scenarios.

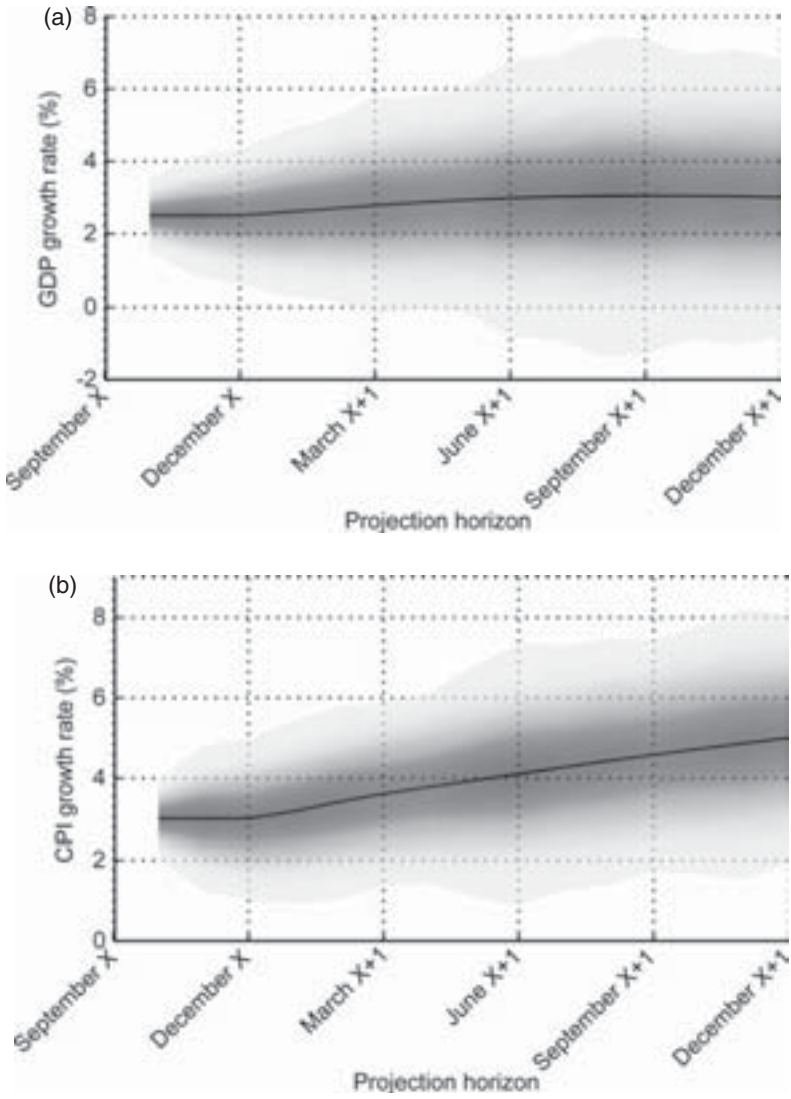

Figure 2.13 Inflationary macroeconomic evolution: (a) GDP YoY% Growth; (b) CPI YoY% Growth.

for the US Gov 0–1Y (Figure 2.15a) and for the US Gov 7–10Y (Figure 2.15b) indices.

The projected evolution of the yields corresponding to the different generic indices modelled permit us to compute the returns for those indices over the relevant investment horizon (from December X to December X + 1

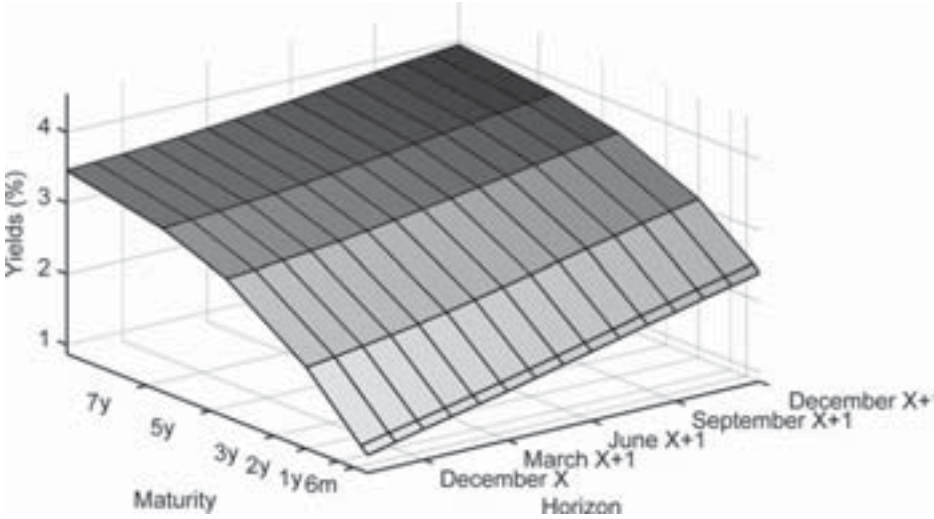

Figure 2.14 Projected average evolution of the US Government yield curve in a non-normal example.

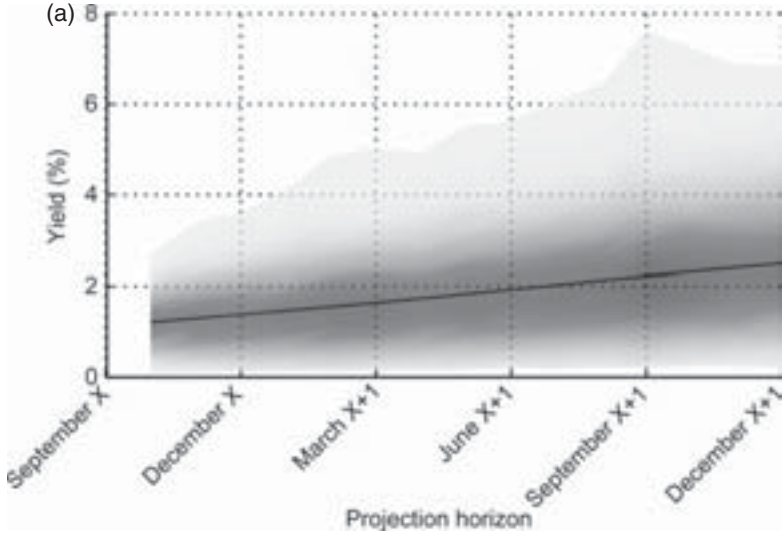

Figure 2.15 Projected distribution of yields in a non-normal example: (a) US Gov 0–1Y, (b) US Gov 7–10Y.

in this example). Table 2.9 illustrates the summary return statistics for the different indices in this example.

It is precisely in this sort of non-normal environment where the presented summary return statistics lose most of their representative power and, therefore, a better representation of the return distributions is needed. To illustrate this, the extremely right-skewed and leptokurtic distribution

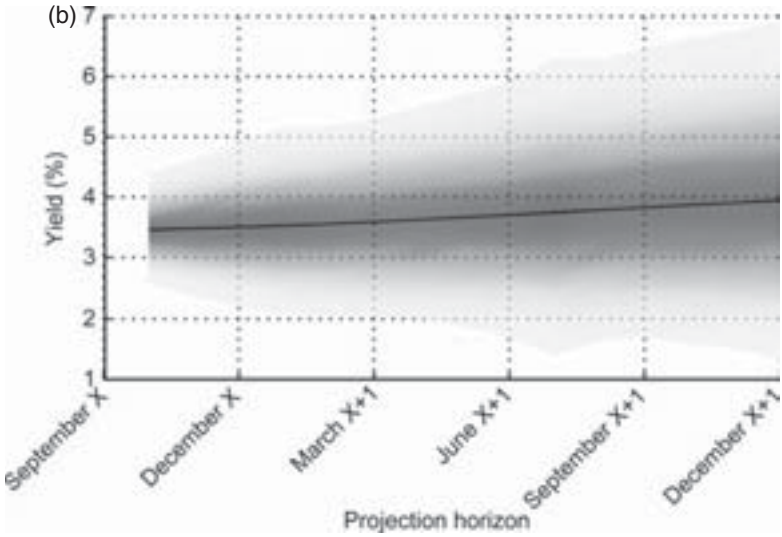

Figure 2.15 (cont.)

Table 2.9 Returns in a non-normal example: average and standard deviation

Asset class	Maturity segment	Average returns (%)	Standard deviation (%)
Government	0–1 Years	1.75	0.53
	1–3 Years	1.30	1.49
	3–5 Years	1.02	3.00
	5–7 Years	0.76	4.00
	7–10 Years	0.71	5.35
Spread Products	0–1 Years	1.88	0.56
	1–3 Years	1.54	1.53
	3–5 Years	1.50	3.17
	5–7 Years	1.50	4.37
	7–10 Years	1.80	6.12
Cash/Depo	1 Month	2.03	0.63

of the simulated returns for the US Gov 0–1Y index is presented in Figure 2.16a and the distribution of returns corresponding to the US Gov 7–10Y index in Figure 2.16b.

6.5.3 Optimal portfolio allocations

Based on these mean returns and the simulated return distributions, the portfolio optimizer finds the portfolio that maximizes expected return

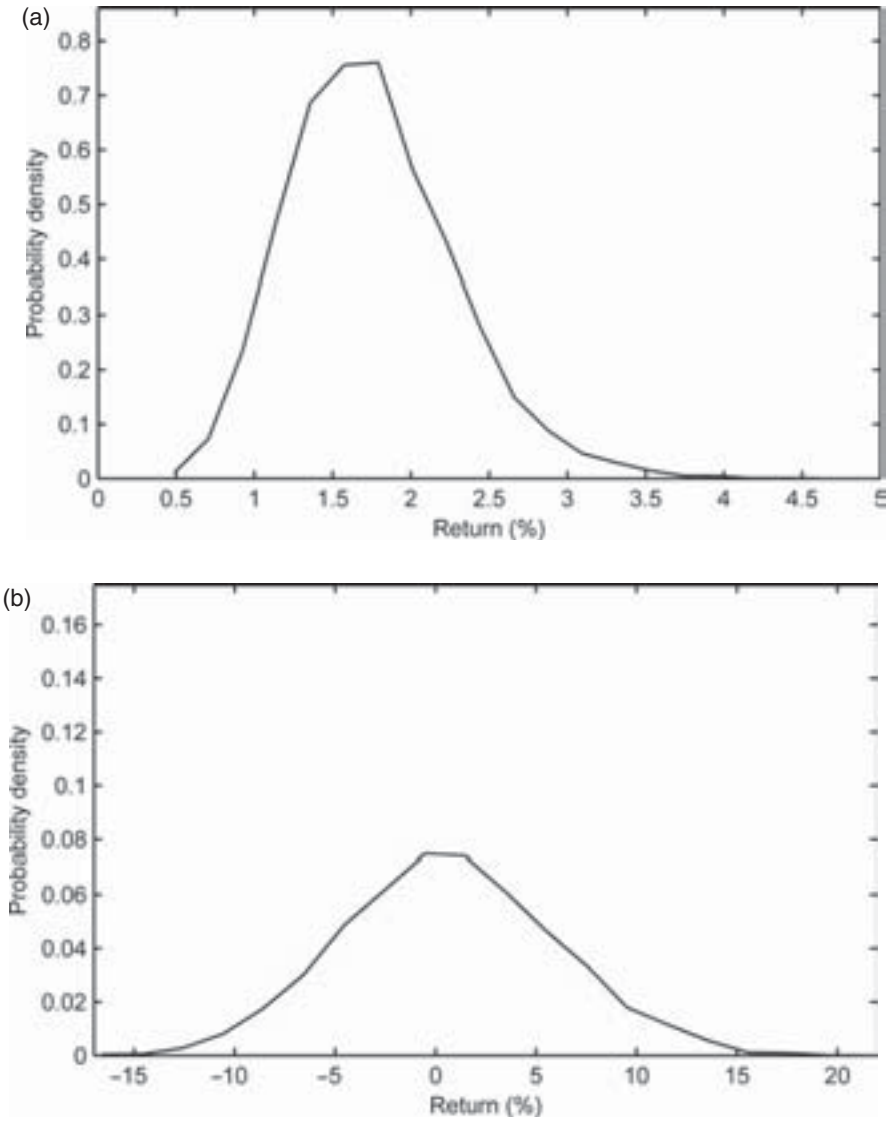

Figure 2.16 Distribution of returns in a non-normal example: (a) US Gov 0–1Y; (b) US Gov 7–10Y.

subject to a no-loss constraint specified in terms of CVaR with a confidence level of 99.5 per cent as described in Section 5.2. Unfortunately, due to the imposed constraint on minimum holdings relative to market capitalization and to the extreme scenario presented, such a portfolio does not exist. Consequently, the strategy has been changed to find the portfolio that minimizes risk (CVaR), applying to a single currency case the methodology

Table 2.10 Optimal portfolio composition in a non-normal example

Asset class	Maturity segment	Portfolio weights (%)
Government	0–1 Years	44
	1–3 Years	9
	3–5 Years	4
	5–7 Years	2
	7–10 Years	4
Spread Products	0–1 Years	12
	1–3 Years	7
	3–5 Years	3
	5–7 Years	3
	7–10 Years	2
Cash/Depo	1 Month	10

Table 2.11 Summary information for the optimal portfolio in a non-normal example

Expected Return	1.64%
Modified Duration	1.5
Volatility	1.53%
VaR (99.5%)	− 0.90%

described in Section 5.1.1. This produces an optimal exposure to each generic index as illustrated in Table 2.10.

The asset allocation is in this case basically determined by the minimum risk exposure implied by the imposed constraints on minimum and maximum holdings. Some summary information for this portfolio is presented in Table 2.11.

This portfolio violates the no-loss constraint at the specified confidence level, due to the cost of including a restriction on minimum holdings relative to the market capitalization, which prevents the Modified Duration of the portfolio to fall below 1.5. The possible trade-off between different objectives, as the specification of a no-loss constraint and the benchmark stability provided by the minimum holdings constraint, has to be treated carefully in the optimization set-up. In this case benchmark stability and liquidity and credit exposure limit compliance have been imposed as primary goals, while the no-loss constraint has been instrumented to play a secondary role, and finally, the maximization of return has been integrated as the third objective in order of importance.

3 Credit risk modelling for public institutions' investment portfolios

Han van der Hoorn

1. Introduction

Credit risk may be defined as the potential that an obligor (borrower or counterparty) is unwilling or unable to meet his financial obligations in a timely manner. Credit risk is a dynamic and broad concept as it encompasses default risk (i.e. an obligor being unwilling or unable to repay his debt) as well as changes in the quality of the credit (e.g. a rating change). Credit risk in central banks comes from two sources. The first is related to policy operations and is discussed in Chapters 7 to 10. The second source of credit risk comes from investment activities, and is the topic of this chapter.

Credit risk is the dominant source of financial risk in a typical commercial bank, whose traditional role is of an intermediary between lenders and borrowers. In contrast, and as illustrated by Table 1.3 in Chapter 1, the typical central bank has only a very limited exposure to credit risk, in particular when compared with currency and gold price risk. The picture for public institutions and also institutional investors is more mixed. For some of them, lending is a core activity, and their credit risk profile may resemble that of a commercial bank. Examples include the European Investment Bank (EIB) or the European Bank for Reconstruction and Development (EBRD). But for others, including (state) pension funds and sovereign wealth funds, credit is not necessarily a natural asset class and their exposures too tend to be rather modest. Many of the topics in this chapter therefore also apply to this wider audience.

The investment universe of central banks is expanding gradually, also into instruments with higher credit risk, but credit risk modelling is still in its early development phase. This chapter reflects some of these developments and is organized as follows. Section 2 starts with a discussion of some

of the arguments that explain why credit risk is increasing in central bank portfolios. Section 3 presents the ECB's approach towards credit risk modelling. In this section, the main parameters of the model will be discussed and compared with a peer group of Eurosystem National Central Banks (NCBs). An empirical analysis is done for two different portfolios, with the aim of comparing simulation results and estimating sensitivities to parameter changes. The results are presented in Section 4. Section 5 concludes.

2. Credit risk in central bank and other public investors' portfolios

Traditionally, central banks have been very conservative investors. Although, on a mark-to-market basis, the typical central bank balance sheet is very risky, there has been little if any appetite for credit risk. This is because the dominant sources of risk, currency and gold price risk, are direct consequences of a central bank's mandate to maintain price stability and are therefore regarded as (at least partly) 'unavoidable' or 'inescapable'. Adding credit risk, in contrast, may improve the risk–return profile of the balance sheet, but comes at the expense of security and liquidity. Since return maximization is not the primary objective of a central bank, it should come as no surprise that the amount of credit risk in any central bank portfolio is below the 'optimal' level from a pure investment perspective.

Nevertheless, diversification into 'non-traditional' assets which bear credit risk is increasing gradually, for central banks as well as public wealth funds. Wooldridge (2006) estimates that the proportion of foreign official institutions' long-term USD debt securities held in corporate, other non-government and non-agency debt had risen to 4.2 per cent in 2005, significantly less than for instance the share of corporate (non-securitized) debt in e.g. the Lehman Global Aggregate Index (approximately 17 per cent at the end of 2007), but still twice as much as five years earlier. Pringle and Carver (2005), in one of their annual surveys of reserve management trends, observe that 'The single most important risk facing central banks in 2005 is seen as *market risk* (reflecting expectations of volatility in securities markets and exchange rates). However, large central banks view *credit risk* as likely to be equally if not more important for them as diversification of asset classes increases their exposure to a wider range of borrowers/investments.' The ECB has also gradually expanded its investment universe, primarily for its

(domestic) own funds portfolio, thereby cautiously adding some credit risk to its investment portfolios (ECB 2006a).

There are a number of explanations for this trend. As already discussed in Chapter 1, central bank reserves have grown rapidly in recent years, in particular in Asia. To the extent that some of these reserves may not be directly needed to fulfill public duties (e.g. be used to fund interventions), the public increasingly demands a decent investment return on assets. At the same time, until recently, expected returns have diminished, as a result of lower interest rates and risk premia. Credit instruments may offer attractive investment opportunities with higher expected returns than traditional assets such as government debt, at only modest additional risk. This is the argument brought forward by, amongst others, de Beaufort *et al.* (2002) and Grava (2004). At the same time, the rapid growth of the market for credit derivatives has lowered 'barriers to entry' to the credit market for non-traditional financiers. This trend has in particular enabled investors to 'buy' exposure in sectors to which they otherwise would not have had access (such as small- and medium-sized enterprises – SMEs). This last argument is particularly relevant for other public and private investors, as central banks mostly shy away from derivatives.

Moreover, several studies argue not only that the expected return on investment grade credit is higher than the expected return on similar government bonds, but that the risk within a single currency market is also lower, as a result of negative correlations between spreads and the level of government yields (see, for instance, Loeys and Coughlan 1999), although it is not clear if this view is maintained in light of the recent financial markets turmoil. Credit risk can also be a hedge for currency risk, and vice versa, as demonstrated by Gould and Jiltsov (2004). Given the large amounts of currency risk in a typical central bank balance sheet, this result is potentially very relevant for central banks. The intuition is that certain currencies act as a safe haven and are in strong demand after a credit event in other currency markets. A particularly good hedge was found in the Swiss franc versus USD corporate bonds.

In both of these studies, risk is measured by the standard deviation of return and, hence, it is implicitly assumed that portfolio returns are normally distributed. This is not necessarily appropriate for credit risk – indeed, this is the motivation for devoting a separate chapter to credit risk – although Loeys and Coughlan (1999) argue that the return distribution of a well-diversified high-quality credit portfolio is not dissimilar from

government bond portfolios. But even if the assumption of normality is dropped and the risk in the tail of the return distribution is measured, it can be shown that, under certain conditions, even a high credit quality portfolio may show a considerable amount of credit risk, once the confidence level of common risk measures like value at risk due to credit risk (CreditVaR, but simply referred to as VaR in the remainder of this chapter, except where confusion might arise) approaches 100 per cent. It turns out that diversification into assets that are more risky in isolation may reduce risks at the portfolio level.

Clearly, investing in credit instruments is not a free lunch, even if it reduces risks under most circumstances. Liquidity and security are lower than in government securities with similar durations. Moreover, the payoff of credit instruments such as corporate bonds is highly asymmetric. The upside is limited (redemption at par), whereas the downside in the event of a default is much larger. Although the downside risk may be somewhat mitigated by diversification and a semi-active investment style that tries to avoid defaults by selling bonds once they have been downgraded beyond a certain threshold (naturally, at the expense of giving up higher expected returns), default should remain a concern to any central bank, in particular because it may harm its reputation. The determinants of credit spreads and expected excess returns – in particular the gap between spreads and expected default losses ('credit spread puzzle') – have attracted a lot of research. Some evidence seems to indicate that an investor is mainly compensated for insufficient diversification opportunities and, hence, for tail risk events, i.e. defaults, even of investment grade issuers (Box 3.1).

Moreover, investing in credit instruments has resource implications, not only for front office and risk management areas, but all the way to the decision-making bodies, that need to spend time and effort in understanding different and more complex instruments than traditionally used. Although there are positive spin-offs from this 'intellectual investment', the time spent on non-core tasks such as credit investments naturally reduces the time that can be devoted to the core activities of the central bank. This explains why a central bank credit portfolio is likely to consist at most of fairly 'plain-vanilla' credit instruments only, such as deposits, corporate bonds and, more remote, credit default swaps (CDS). The same is true for other conservative investors, in particular if investing is not core business.

Box 3.1. Credit spreads and the limitations to diversification

It is widely accepted that diversification of credit portfolios is more difficult and, at the same time, more important than for e.g. equity portfolios. It is more important because of the asymmetric return distribution of credits: a single default can easily offset positive returns from all other assets in the portfolio, even when the number of 'names' in the portfolio is large and individual obligor weights are correspondingly small. In an equity portfolio, the probability of a single stock suffering large losses may be higher, but such losses are more likely to be compensated by positive returns on other stocks, since their upside is much higher than for credits.

Diversification of credit risk comes in at least three dimensions: by sector, by region and by individual name. Each of these is probably more difficult than for equity exposure, but it seems particularly problematic for sector diversification. The euro corporate bond market is dominated by financials, which cover approximately one-third of the overall investment grade market. Excluding BBB-rated issuers – which may be relevant for a conservative investor like a central bank – the dominance of financials becomes even stronger: around 50 per cent of the AAA–A corporate debt in euros is issued by financials (*source*: iBoxx EUR corporates senior index, December 31, 2007), even though non-financial companies are increasingly being rated and gaining in importance (ECB 2004b). At present, the three largest sectors – financials, utilities and telecom – cover around 75 per cent of the investment grade market. Clearly, correlations within sectors are typically higher than among issuers in different sectors. For this reason, the 'market portfolio', the cornerstone of the Capital Asset Pricing Model (CAPM), may not be the optimal investment portfolio. This becomes apparent also when one realizes that the market portfolio is to a large extent supply driven: heavily indebted and therefore risky companies have larger weights in the market and in market indices.

The difficulties of diversification have triggered research into the question whether credit spreads reflect idiosyncratic risk. According to the CAPM, an equity investor is rewarded for exposure to (general) market risk only. For credits, a compensation for specific risk may be justified if idiosyncratic risks cannot be fully diversified away. Proponents of this theory are, among other people, Amato and Remolona (2003), who essentially study name concentration. They argue that even a credit portfolio with 300 names may be poorly diversified (based on hypothetical portfolios, with individual assets comparable to BBB). While other factors – expected losses, liquidity, taxes – help to explain some of the credit spread, they argue that the limitations to diversification are the largest contributor.

The aim of this section is not to discuss the pros and cons of credit in central bank portfolios at length. Rather, it is noted that there may be good arguments to invest some of the central bank reserves in credit instruments, and that this is increasingly happening in practice. The arguments for and against are not the same for all central banks and depend, inter alia, on the size of reserves, the risk tolerance and resources of the central bank. The

ECB's investment-related exposure to credit risk is limited and mainly comes from short-term bank deposits and investments in agencies, senior unsecured bank bonds and covered bonds. Evidently, the risk rises if government bonds can no longer be considered credit risk free, which may be a valid approach for stress testing, as e.g. recommended in principle 13 of the BCBS (2000b) *Principles for the management of credit risk*: 'Banks should . . . assess their credit risk exposures under stressful conditions.' This assumption will also be used for the simulation exercises in this chapter.

A central bank exposed to credit risk should consider, like any other investor, the purchase or in-house development of a portfolio credit risk model that captures the asymmetry and fat tails of credit return distributions. As outlined in the next section, the structure of these models is very different from market risk models, that have by now become commonplace also in central banks. Within the Eurosystem, only a few central banks have practical experience with credit risk modelling, and so far they are used mainly for reporting. In the future, they are likely to be used for a variety of other reasons as well, including limit setting and strategic asset allocation decisions, thereby making the trade-off between market and credit risks.

3. The ECB's approach towards credit risk modelling: issues and parameter choices

3.1 Motivation

Credit risk models are generally very different in nature from the market risk models that are discussed in Chapters 2 and 4 of this book. Credit risk models also suffer from serious data limitations: defaults are rare events and *correlated* defaults are even rarer. This makes it problematic to derive statistically robust and reliable estimates of credit risk. For portfolios dominated by government bonds, the data problem is even more challenging. Moreover, the impact of a credit event – default or downgrade – is potentially very large and can easily erase one year of performance or more. Given the limited upside of credit instruments, the return distribution of credit instruments is very asymmetric to the downside and has a fat tail. While the normal distribution may be a reasonable assumption for the return of many 'market' instruments with approximately linear pay-off (i.e. non option-like) structures, this is clearly inappropriate for credit risk, except perhaps under very special circumstances.

There are a number of competing approaches to model credit risk, each with their own strengths and weaknesses. The purpose of this chapter is not to give an overview of various approaches – the interested reader is referred to a growing list of textbooks[1] – but to describe and motivate the ECB's approach, with a focus on issues and parameters that are particularly relevant for central banks and other public investors.

The selection of a credit risk model at the ECB was driven by theoretical as well as practical considerations. It was foreseen that the model would primarily be used for ex post risk assessments, and that it might be integrated in strategic asset allocation decisions (see Chapter 2) and be employed for limit setting (Chapter 4), but it was considered unlikely that the model would be used for trading. This setting has a number of implications, two of which are worth mentioning here. The first is that speed of calculations does not have a very high priority. A simulation-based approach can therefore be used, which can be made very flexible and intuitive, also for non-insiders, even if the technical details may be complex. The second implication is that there is no need for a very precise pricing model; for our purpose, a crude approximation of (relative) prices based on ratings and generic credit curves is sufficient. Given that, in addition, all issuers and counterparties of the ECB are rated by at least one of the major rating agencies, it is natural to use a ratings- and simulation-based approach for credit risk modelling.

At the time the decision was made to model credit risk more formally – around 2005 – an off-the-shelf system already existed that seemed to fulfill most of the ECB's requirements and needs, CreditManager® from the RiskMetrics Group, based on the well-known CreditMetrics™ methodology (Gupton *et al.* 1997). It was, however, decided to develop an in-house system. Aside from the more general considerations regarding the choice between build and buy, discussed in Chapter 4, a particular argument in favour of an in-house development was the learning experience in terms of modelling and improved understanding of credit markets. At the time, these were fairly underdeveloped areas of expertise, which deserved more attention. Moreover, commercial systems seemed primarily targeted at 'pure' asset managers and investors, and were considered not necessarily optimal for central bank portfolios.

At the same time, it was recognized that an in-house model does not undergo rigorous testing by the market. It was therefore decided to

[1] This list includes Bluhm *et al.* (2003), Cossin and Pirotte (2007), Duffie and Singleton (2003), Lando (2004), Saunders and Allen (2002). A particularly good introduction for practitioners is Ramaswamy (2004a).

'benchmark' the model against similar models used by several Eurosystem NCBs. This benchmark study has culminated in an ECB Occasional Paper by a Task Force of the Markets Operations Committee of the European System of Central Banks (2007). Some of its main findings are also discussed in this chapter.

There are some apparent limitations to the use of external ratings, most of which have been well known for many years (see for example BCBS 2000a). Despite these limitations, ratings are still believed to add value and are considered an efficient instrument for resource-constrained investors, such as central banks, although they should not and are not taken as a substitute for own risk analysis. When treated with care, an external rating can be used as an initial assessment of credit risk. In order to reduce the risk of using an overly optimistic rating, a second-best rating rule is in place (see also Chapter 4).

In what follows, a risk horizon of one year is assumed, although longer or shorter horizons can also be considered. Note that rating agencies claim that their ratings reflect a 'through-the-cycle' opinion of credit risk, which obviously impacts the migration probabilities. Actual probabilities over a one-year point-in-time horizon fluctuate around the through-the-cycle averages and depend on economic and financial market conditions. The ECB's credit risk model distinguishes eight rating levels: AAA, AA, A, BBB, BB, B, CCC-C and D (= default, all using the S&P/Fitch classification). As explained in detail in Chapter 4, one of the eligibility criteria for issuers and counterparties is that they have at least one rating (of a certain level) by one of the major rating agencies. Consequently, the initial rating of each obligor in the portfolios is known and can be mapped onto one of the eight rating levels. Probabilities of default and up- or downgrades for all obligors are readily obtained from historical default and migration studies and summarized in so-called 'migration (or transition) matrices', provided the maturities of the positions exceed the horizon of the migration probabilities (if not, adjustments will have to be made that are discussed in the next section). These matrices are published and updated at least annually by all the rating agencies. It is important to realize that the ECB model does not provide PD estimates; instead, PDs are important input parameters.[2] Given also the magnitude of credit spreads and recovery rates, it is relatively easy to estimate the expected value and, hence, expected credit loss over a given horizon for every obligor in portfolio. The horizon is set at one year. The

[2] The most common alternative approaches are estimating these probabilities from bond prices and spreads, using a reduced form model, and from the volatility of stock prices, using a structured model in the spirit of Merton (1974).

derivation of summary statistics other than expected loss involves more steps and is discussed in the next sections. A distinction is made between those results that can be derived analytically and those for which simulation is needed.

3.2 Analytical results

The core of the ECB's credit risk model is a large-scale simulation engine, but some results – expected loss and unexpected loss – are also derived analytically. Whenever available, analytical results are preferred over simulation results, which are essentially random and therefore subject to finite-sample noise. Moreover, analytical results play an important role in the validation of simulation results, since expected and unexpected losses are also estimated from the simulation output. In order to formalize the analytical derivation of expected loss, already touched upon in the previous section, it is useful to introduce a number of concepts that will facilitate notation, in particular for the derivation of unexpected loss.

Define the *forward value FV* as the value of a position that is found by moving one year forward in time, while keeping the rating of the obligor unchanged. It is computed by discounting all cash flows at the relevant discount rate. Formally, for a position in obligor i and portfolio P:

$$FV_i = \sum_j CF_{ij} df^{icr_i}(t_{ij}) \tag{3.1}$$

$$FV_P = \sum_{i=1}^{n} FV_i \tag{3.2}$$

Here, CF_{ij} represents the jth cash flow (in EUR) by obligor i, t_{ij} is the time (in years) of the cash flow and $df^{cr}(t)$ is the one-year forward discount factor for a cash flow at time t from an obligor with a credit rating equal to cr (icr_i is the initial credit rating of obligor i). This discount rate is derived from the relevant spot (zero coupon) rates y at maturities 1 and t years. Assuming, in addition, that any cash flows received during the year are not reinvested, so that the value of any of these cash flows at time $t = 1$ is simply equal to the cash flow itself, the expression for the forward discount factors, using continuous compounding, is as follows:

$$df^{cr}(t) = \begin{cases} \exp[y^{cr}(1) - t \cdot y^{cr}(t)], & t > 1 \\ 1, & t \leq 1 \end{cases} \tag{3.3}$$

The *conditional* forward value *CFV* is the forward value of a position, conditional upon a rating change:

$$CFV_i^{fcr} = \begin{cases} \sum_j CF_{ij} df^{fcr}(t_{ij}), & fcr \neq D \\ Nom_i rr_i, & fcr = D \end{cases} \tag{3.4}$$

In this expression, *fcr* is the forward credit rating (*D* is default), Nom_i is the nominal investment (in EUR) in obligor *i* and rr_i is the recovery rate (in per cent) of obligor *i*.

Finally, the *expected* forward value *EFV* is the forward value of the position, taking into account all expected rating changes. It is therefore a weighted average of conditional forward values, with weights equal to the probabilities of migration *p*:

$$EFV_i = \sum_{fcr} p(fcr|icr_i) CFV_i^{fcr} \tag{3.5}$$

With these concepts, we can simply write the expected (forward) loss *EL* as the differences between the forward and expected portfolio value:

$$EL_P = FV_P - \sum_{i=1}^{n} EFV_i \tag{3.6}$$

Note that expected loss is defined as the difference between two portfolio values, both at time $t = 1$. Hence, the current market value of the portfolio is not used; if it were, expected loss would be 'biased' by the time return (carry and possibly roll-down) of the portfolio. Defining expected loss as in equation (3.6) ensures a 'pure' credit risk concept. It is useful to decompose expected loss into the contribution of migration (up- and downgrades) and default. Substituting (3.1), (3.2), (3.4) and (3.5) in (3.6) and rearranging, it is easy to verify that

$$EL_P = \sum_{i=1}^{n} \left\{ \underbrace{\sum_{fcr \neq D} p(fcr|icr_i) \sum_j CF_{ij} \left[df^{icr_i}(t_{ij}) - df^{fcr}(t_{ij}) \right]}_{\text{contribution of migration}} \right.$$

$$\left. + \underbrace{p(D|icr_i) \left(\sum_j CF_{ij} df^{icr_i}(t_{ij}) - Nom_i rr_i \right)}_{\text{contribution of default}} \right\} \tag{3.7}$$

The first element of the right-hand side of equation (3.7) represents the contribution of migration per obligor. It is equal to a probability-weighted average of the change in forward value of each cash flow. For high-quality portfolios, a reasonably good first-order approximation of this expression is usually found by multiplying the modified duration of the bond one year forward by the change in the forward credit spread. The second element is the contribution of default.

Unexpected loss UL, defined as the standard deviation of losses in excess of the expected loss, is derived in a similar way, although the calculations are more involved and need assumptions on the co-movement of ratings. Building on the concepts already defined, a convenient way to compute unexpected loss analytically involves the computation of standard deviations of all two-obligor subportfolios (of which there are $n\,[n-1]\,/\,2$), as well as individual obligor standard deviations. First note that, by analogy of expected loss, the variance (unexpected loss squared) of each individual position is given by

$$UL_i^2 = \sum_{fcr} p(fcr|icr_i)\left(CFV_i^{fcr}\right)^2 - EFV_i^2 \tag{3.8}$$

In this formula, it is assumed that there is uncertainty only in the ratings one year forward, and that conditional forward values of each position are known. It could be argued that there is also uncertainty in these values, in particular the recovery value, in which case the standard deviation needs to be added to the conditional forward values.

A similar calculation can be made for each two-obligor portfolio, but the probabilities of migration to each of the 8×8 possible rating combinations depend on the joint probability distributions of ratings. Rather than modelling this directly, it is common and convenient to assume that rating changes are driven by an underlying asset return x and to model joint asset returns as standard bivariate normal with a given correlation ρ, known as asset correlation. The intuition of this approach should become clear in the next section on simulation. The joint probability of migrating to ratings fcr_i and fcr_j, given initial ratings icr_i and icr_j, and correlation p_{ij} equals

$$p\left(fcr_i, fcr_j \middle| icr_i, icr_j; p_{ij}\right) = \int_{b^-_{fcr_i|icr_i}}^{b^+_{fcr_i|icr_i}} \int_{b^-_{fcr_j|icr_j}}^{b^+_{fcr_j|icr_j}} \frac{1}{2\pi\sqrt{1-\rho_{ij}^2}} \exp\left(-\tfrac{1}{2}\left(x_i^2 + x_j^2 - 2\rho_{ij}x_i x_j\right)\right) dx_j dx_i \tag{3.9}$$

where the b represent the boundaries for rating migrations from a standard normal distribution (also explained in the next section). The probabilities allow the variance computation for each two-obligor portfolio:

$$UL_{i+j}^2 = \sum_{fcr_i} \sum_{fcr_j} p\left(fcr_i, fcr_j \middle| icr_i, icr_j; p_{ij}\right) \left(CFV_i^{fcr_i} + CFV_j^{fcr_j}\right)^2 - \left(EFV_i + EFV_j\right)^2 \tag{3.10}$$

With the results from equations (3.8)–(3.10), it is easy to compute the unexpected loss of the portfolio:[3]

$$UL_P = \sqrt{\sum_{i=1}^{n-1} \sum_{j=i+1}^{n} UL_{i+j}^2 - (n-2) \sum_{i=1}^{n} UL_i^2} \tag{3.11}$$

3.3 Simulation approach

Expected and unexpected loss capture only the first two moments of the credit loss distribution. In fact, formally, expected loss is not even a risk measure, as risk is by definition restricted to unexpected events. Unexpected loss measures the standard deviation of losses and is also of relatively limited use, as a credit loss distribution is skewed and fat tailed. In order to derive more meaningful (tail) risk measures such as VaR and ES (expected shortfall), a simulation approach is needed.

To understand the simulation approach at the portfolio level, consider first a single bond with a known initial rating, e.g. A. Over a one-year period, there is a high probability that the rating remains unchanged. There is also a (smaller) probability that the bond is upgraded to AA, or even to AAA, and, conversely, that it is downgraded or even that the issuer defaults. Assume that rating changes are driven by an underlying asset value of the issuer, the same as was used for the analytical derivation of unexpected loss in the previous section. The bond is upgraded if the asset value increases

[3] This result is derived from a standard result in statistics. If $X_1, \ldots, X_n$ are all normal random variables with variances σ_i^2 and covariances σ_{ij}, then $Y = \Sigma X_i$ is also normal and has variance equal to $\sigma_Y^2 = \sum_{i=1}^{n} \sigma_i^2 + 2 \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} \sigma_{ij}$. Rearranging the formula for a two-asset portfolio $X_i + X_j$ yields an expression for each covariance pair: $\sigma_{ij} = \frac{1}{2}\left(\sigma_{i+j}^2 - \sigma_i^2 - \sigma_j^2\right)$ which, when substituted back into the formula for σ_Y^2, gives the desired result.

Table 3.1 Migration probabilities and standard normal boundaries for bond with initial rating A

Rating after one year (fcr)	D	CCC-C	B	BB	BBB	A	AA	AAA
Migration probability	0.06%	0.03%	0.17%	0.41%	5.63%	91.49%	2.17%	0.04%
Cumulative probability	0.06%	0.09%	0.26%	0.67%	6.30%	97.79%	99.96%	100.00%
Lower migration boundary ($b^{-}_{fcr\|A}$)	$-\infty$	−3.24	−3.12	−2.79	−2.47	−1.53	2.01	3.35
Upper migration boundary ($b^{+}_{fcr\|A}$)	−3.24	−3.12	−2.79	−2.47	−1.53	2.01	3.35	$+\infty$

Source: Standard & Poor's (2008a, Table 6 – adjusted for withdrawn ratings).

beyond a certain threshold and downgraded in case of a large decrease in asset value. The thresholds are set such that the ratings derived from the (pseudo-) random asset returns converge to the migration probabilities, given a certain density for the asset returns. For the latter, it is common to assume standard normally distributed asset returns, and this is also the approach of the ECB model, although other densities may be used as well.[4] Note that normal asset returns do *not* imply that migrations and therefore bond prices are normally distributed. The process is illustrated using actual migration probabilities from a recent S&P study in Table 3.1.

The table shows, for instance, that the simulated rating remains unchanged, for simulated asset returns between thresholds $b^{-}_{A|A} = -1.53$ and $b^{+}_{A|A} = 2.01$. If the simulated asset return is between $b^{-}_{AA|A} = 2.01$ and $b^{+}_{AA|A} = 3.35$, the bond is upgraded to AA, and if it exceeds $b^{-}_{AAA|A} = 3.35$, the bond is upgraded to AAA. Note that $b^{+}_{fcr|icr} = b^{-}_{fcr+1|icr}$ for any combination of initial credit rating *icr* and forward credit rating *fcr* (where '+1' refers to the next highest rating). The same levels are also used in the analytical derivation of unexpected loss in equation (3.9). The mechanism applies to downgrades. It is easy to verify that the simulated frequency of ratings should converge to the migration probabilities. Combining simulated ratings with yield spreads and recovery rates yields asymptotically the

[4] The normal distribution of asset returns is merely used for convenience, because the only determinant of co-dependence is the correlation. It is quite common to use the normal distribution, but in theory alternative probability distributions for asset returns can also be used. These do, however, increase the complexity of the model.

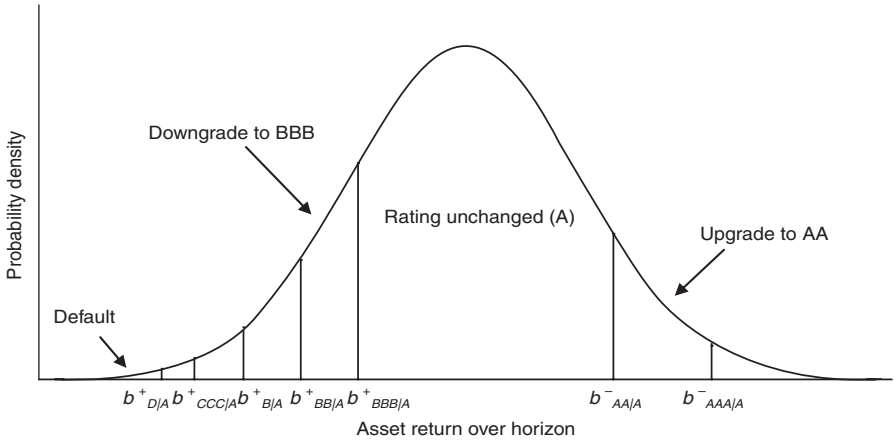

Figure 3.1 Asset value and migration (probabilities not according to scale).

same expected loss as in the analytic approach. The simulation approach is illustrated graphically in Figure 3.1.

At the portfolio level, random asset returns are sampled independently and subsequently transformed into correlated returns. This is achieved via a Cholesky decomposition of the correlation matrix into an upper and a lower triangular matrix and by subsequently pre-multiplying the vector of independent asset returns (of length n) by the lower triangular matrix (of dimension $n \times n$).[5] The result is a vector of correlated asset returns, each of which is converted into a rating using the approach outlined above. Assuming deterministic spread curves, the ratings are subsequently used to reprice each position in the portfolio. The model can also be used for multi-step simulations, whereby the year is broken down in several subperiods, or

[5] A correlation matrix $\mathbf{R}$ is decomposed into a lower triangular $\mathbf{L}$ and an upper triangular matrix $\mathbf{L}'$ in such a way that $\mathbf{R} = \mathbf{L}\mathbf{L}'$. A vector of independent random returns $\mathbf{x}$ is transformed into a vector of correlated returns $\mathbf{x}_c = \mathbf{L}\mathbf{x}$. It is easy to see that $\mathbf{x}_c$ has zero mean, because $\mathbf{x}$ has zero mean, and a correlation matrix equal to $E(\mathbf{x}^c(\mathbf{x}^c)') = E(\mathbf{L}\mathbf{x}\mathbf{x}'\mathbf{L}')' = \mathbf{L}E(\mathbf{x}\mathbf{x}')\mathbf{L} = \mathbf{L}\mathbf{I}\mathbf{L}' = \mathbf{L}\mathbf{L}' = \mathbf{R}$, as desired. Since correlation matrices are symmetric and positive-definite, the Cholesky decomposition exists. Note, however, that the decomposition is not unique. It is, for example, easily verified that if $\mathbf{L} = \begin{pmatrix} l_{11} & 0 & \cdots & 0 \\ l_{21} & l_{22} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ l_{n1} & \cdots & l_{n2} & l_{nn} \end{pmatrix}$ is a valid lower triangular matrix, then so is $\mathbf{L}^* = \begin{pmatrix} l_{11} & 0 & \cdots & 0 \\ l_{21} & l_{22} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ l_{n1} & \cdots & l_{n2} & -l_{nn} \end{pmatrix}$. Any of these may be used to transform uncorrelated returns into correlated returns.

where the horizon consists of several one-year periods. In those cases, the vector of returns becomes a matrix (of dimension $n \times$ # periods), but otherwise the approach is essentially the same as for a one-step simulation. As shown in Chapter 2, it is also possible to use stochastic spreads, thus integrating market (spread) and credit risk, but this chapter considers deterministic spreads only.

In order to generate reliable estimates of (tail) risk measures, a large number of iterations are needed, but the number can be reduced by applying importance sampling techniques. Importance sampling is based on the idea that one is really only concerned with the tail of the distribution, and should therefore sample more observations from the tail than from the rest of the distribution. With importance sampling, the original distribution from which observations are drawn is transformed into a distribution which increases the likelihood that 'important' observations are drawn. These observations are then weighted by the likelihood ratio to ensure that estimates are unbiased. The transformation is done by shifting the mean of the distribution. Technical details of importance sampling are discussed in Chapter 10.

The simulation approach is summarized in the following steps:

Step 0 Create a matrix (of dimension # names $\times$ # ratings), consisting of the conditional forward values of the investment in each obligor under each possible rating realization, as given by equation (3.4).

Step 1 Generate n independent (pseudo-) random returns from a standard normal distribution, but sampling with a higher probability from the tail of the distribution. Store the results in a vector $\mathbf{x}$.

Step 2 Transform the vector of independent returns into a vector of correlated returns $\mathbf{x}^c$ via $\mathbf{x}^c = \mathbf{L}\mathbf{x}$, where $\mathbf{L}\mathbf{L}' = \mathbf{R} = \begin{pmatrix} 1 & \rho_{21} & \cdots & \rho_{1n} \\ \rho_{21} & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \rho_{n-1,n} \\ \rho_{n1} & \cdots & \rho_{n,n-1} & 1 \end{pmatrix}$ is the (symmetric) correlation matrix.

Step 3 Transform the vector of correlated returns into a vector of ratings via $fcr_i = \arg\max_{cr} \left\{ \mathbf{1}\left[x_i^c \geq b^-_{cr|icr_i}\right] \times \mathbf{1}\left[x_i^c < b^+_{cr|icr_i}\right] \right\}$, where $1[\cdot]$ is an indicator function, equal to unity whenever the statement in brackets is true, and zero otherwise.

Step 4 Select, in each row of the matrix created in step 0, the entry (conditional forward value) corresponding to the rating simulated

in step 3. Compute the simulated (forward) portfolio value $SFV_P^{(1)}$ as the sum of these values, where the (1) indicates that this is the first simulation result.

Step 5 Repeat steps 1–4 many times and store the simulated portfolio values $SFV_P^{(i)}$.

Step 6 Sort the vector of simulated portfolio values in ascending order and compute summary statistics (*sim* is the number of iterations):

- $\overline{SFV_P} = \frac{1}{sim}\sum_{i=1}^{sim} SFV_P^{(i)}$,
- $EL^{(s)} = FV_P - \overline{SFV_P}$,
- $UL^{(s)} = \sqrt{\frac{1}{sim}\sum_{i=1}^{sim}\left(SFV_P^{(i)}\right)^2 - \left(\overline{SFV_P}\right)^2}$,
- $VaR = \overline{SFV_P} - SFV_P^{(a)}$, where $a = sim \times (1 - \text{confidence level})$ rounded to the nearest integer,
- $ES = \overline{SFV_P} - \frac{1}{a-1}\sum_{i=1}^{a-1} SFV_P^{(i)}$.

Each of these can also be expressed as a percentage of FV_P, the market value in one year time if ratings remain unchanged.

Step 7 Finally, the results may be used to fit a curve through the tail of the distribution to avoid extreme 'jumps' in time series of VaR or ES. This, however, is an art as much as science for which there is no 'one size fits all' approach.[6]

The creation of a matrix with all conditional forward values in step 0 makes the simulation very fast and efficient. Without such a matrix, it would be necessary to reprice every position within the loop (steps 1-4) and computing time would increase significantly. If programmed efficiently, even a simulation with 1,000,000 iterations need not take more than one minute on a modern computer. For many portfolios, this number of iterations is more than enough and fairly accurate estimates of credit risk are obtained, even at the highest confidence levels. In practice, therefore, importance sampling or other variance reduction techniques are not always needed.

[6] A possible strategy, depending on the composition of the portfolio, is to make use of a well-known result by Vasicek (1991), who found that the cumulative loss distribution of an infinitely granular portfolio in default mode (no recovery) is in the limit equal to $F(x) = N\left(\frac{\sqrt{1-\rho}N^{-1}(x) - N^{-1}(pd)}{\sqrt{\rho}}\right)$, where ρ is the (positive) asset correlation and $N(x)$ denotes the cumulative standard normal distribution (N^{-1} being its inverse) evaluated at x, representing the loss as a proportion of the portfolio market value, i.e. the negative of the portfolio return.

VaR and ES for credit risk are typically computed at higher confidence levels than for market risk. This is a common approach, despite increasing parameter uncertainty, also for commercial issuers aiming at very low probabilities of default to ensure a high credit rating. For instance, a 99.9 per cent confidence level of no default corresponds only to approximately an A rating. The Basel II formulas for the Internal Ratings Based (IRB) approach compute capital requirements for credit risk at the 99.9 per cent confidence level as well, whereas a 99 per cent confidence level is applied to determine the capital requirements for market risk (BCBS 2006b). Arguably, a central bank – with reputation as its main asset – should aim for high confidence levels, also in comparison with commercial institutions.

The discussion of the analytical and simulation approach has so far largely ignored the choice of parameters and data sources. There are, however, a number of additional complexities related to data and parameters, in particular for central banks and other conservative investors. The remainder of this section is therefore devoted to a discussion of the main parameters of the model, i.e. the probabilities of migration (including default), asset correlations and recovery rates. This discussion is not restricted to the ECB, but includes a comparison with other Eurosystem central banks, more details of which can be found in the paper by the Task Force of the Market Operations Committee of the European System of Central Banks (2007).

3.4 Probabilities of default/migration

Probabilities of default and migration can be obtained from one of the major rating agencies, which publish updated migration matrices frequently. These probabilities typically have a one-year horizon, i.e. equal to the standard risk horizon of the model. As the migration matrices of the rating agencies are normally fairly similar for any given market segment, the selection of a particular rating agency does not seem terribly important, although clearly, in order to generate meaningful time series, one should try to use the same source as much as possible. Also in practice there is not a clear preference among Eurosystem NCBs for any of the three major agencies, Standard & Poor's, Moody's and Fitch.

The methodologies used by the rating agencies for estimating default and migration probabilities rely on counting the number of migrations for a given rating within a calendar year. This number is divided by the total number of obligors with the initial rating and corrected for ratings that have

been withdrawn during the year ('cohort approach'). The approach is fairly straightforward and transparent, but there are several caveats, some of which are of particular relevance to central banks and other investors with high-quality, short-duration assets. The main caveats are related to the probabilities of default for the highest ratings, and the need to scale probabilities for periods shorter than one year. Ideally, these are addressed directly via the data.[7] If one only has access to the migration matrices, but not to a database of ratings, then other solutions are needed. A third caveat, not related to the methodology of estimating the migration matrix, is the distinction between sovereign and corporate ratings, and the limitations of migration probabilities for sovereign ratings. Each of these is discussed below.

3.4.1 Distinction between sovereign and corporate ratings

Most of the empirical work published by the rating agencies is based on corporate ratings (for which there is far more data), but central bank portfolios contain mostly government bonds. It is well known that default and migration probabilities for sovereign issuers are different from probabilities for corporate issuers. Comparing, for instance, the latest updates of migration probabilities by Standard & Poor's (2008a and 2008b) reveals that while, historically since 1981, some AA and A corporate issuers have defaulted over a one-year horizon (with frequencies equal to one and six basis points, respectively, see Standard & Poor's 2008a, table 6), not a single investment grade sovereign issuer has ever defaulted over a one-year horizon (based on observations since 1975, see Standard & Poor's 2008b, table 1. Even after 10 years, A or better-rated sovereigns did not default (Standard & Poor's 2008b, table 5).

The distinction between sovereign and corporate issuers is also reflected in, for instance, the risk classification used in Basel I and II (Standardized Approach). Except for ratings BB+ to BB− and ratings below B−, the risk weights for corporates are higher than for equally rated sovereigns (Table 3.2). While the absence of defaults is a comforting result, one should also be aware

[7] Instead of counting the number of defaults and migrations in a certain period of time, one could measure the *time* until default or migration, and derive a 'hazard rate' or 'default intensity'. With these, one can easily derive the expected time until default or downgrade for every rating class and, conversely, the probability of default or downgrade in any given time period. A related approach is to estimate (continuous time) generator matrices directly from the data (Lando and Skødeberg 2002), rather than via an approximation of a given discrete time migration matrix. The estimation of generator matrices takes into account the exact timing of each rating migration and therefore uses more information than traditional approaches.

Table 3.2 Risk-weighting of Standardized Approach under Basel II

	Sovereigns	Corporates
AAA to AA−	0%	20%
A+ to A−	20%	50%
BBB+ to BBB−	50%	100%
BB+ to BB−	100%	150%
B+ to B−		
Below B−	150%	

Source: BCBS (2006b).

that it is based on a limited number of observations. Hence, the statistical significance of the result is limited. Moreover, the rating agencies themselves acknowledge that the process of rating sovereigns is considerably more complex and subjective than for rating corporates. As a result, many investors use migration probabilities derived from corporate issuers, which leads to conservative, but probably more robust risk estimates. The same approach is adopted by the ECB and several other Eurosystem NCBs.

3.4.2 Probabilities of default for highest ratings

Defaults of AAA (or even AA) corporate (let alone sovereign) issuers rarely happen over the course of one year. As a result, the probabilities of default (PDs) for these ratings estimated by the rating agencies are (very close to) zero. Clearly, this does not necessarily imply that the actual default risk is zero, especially for non-AAA rated issuers. After all, it seems reasonable to assume that the default risk of a AA issuer is higher than the default risk of a AAA issuer. AAA issuers may default as well, even if none ever did. Determining default probabilities for the highest ratings is difficult and largely subjective, yet has a significant impact on the credit risk assessment of a central bank – or similar portfolio. Ramaswamy (2004a, exhibit 5.4) proposes positive default probabilities by first selecting, for each rating level, the highest empirical PD from either Standard & Poor's or Moody's. In addition he proposes some ad hoc approach to ensure that the ranking of ratings is respected, i.e. PD (AAA) $<$ PD (AA+), etc. In his example, such adjustments are needed down to rating A–. For instance, he proposes to set the PD for AAA issuers to one basis point, for AA– to four basis points and for A– to ten basis points (all at an annual horizon). Although the main purpose of the exercise is merely to be able to estimate default correlations, the proposed PDs seem not unreasonable, also for other purposes such as

stress testing. Another, statistically more robust approach has recently been proposed by Pluto and Tasche (2006). They propose estimating confidence intervals for each PD such that the probability of finding not more than the empirical number of defaults is very small. The PD is set equal to the upper bound of the confidence interval. Hence, this approach cannot be used to compute expected losses. However, it does ensure positive PDs, even if the empirical number of defaults is zero. Moreover, the PD decreases as the sample size of non-defaulted issuers increases, as it should. The methodology also respects the ranking of ratings. The approach seems not yet widely used in practice, however.

The ECB system uses the probabilities introduced in Ramaswamy (2004a) for certain analyses, thus manually revising the PDs for AAA and AA rated obligors upwards. In order to ensure that migration probabilities add up to 1, the probabilities that ratings remain unchanged (diagonal on the matrix) are reduced accordingly. Within the Eurosystem, several other central banks apply a similar approach, although some make smaller adjustments to sovereign issuers, or no adjustment at all. All respect the ranking of ratings in the sense that the PD of an issuer with a certain rating is higher than the PD of an issuer with a better rating.

3.4.3 Default probabilities for low duration assets

Central bank portfolios normally have low durations and, hence, a substantial proportion is invested in assets – bills, discount notes, deposits – with maturities less than one year, i.e. shorter than the risk horizon and shorter than the default probabilities obtained from the rating agencies. Clearly, probabilities of default increase with maturity, and therefore it may be necessary to scale annual PDs for assets with short maturities. Note that migration risk is irrelevant for instruments with a maturity less than one year, since the portfolio is assumed static during the year. In what follows, it is assumed that every position that matures before the end of the horizon is held in cash for the remainder of the period.

Scaling default probabilities to short horizons can be done in several ways. The easiest approach is to assume that the conditional probability of default (or 'hazard rate') is constant over time. The only information needed from the migration matrix is the last column which contains the annual probabilities of default. For each rating, the probability of default $pd(t)$ for maturity $t < 1$ follows directly from the one-year probability $pd(1)$ using the formula $pd(t) = 1 - [1 - pd(1)]^t$. This is approximately equal to $pd(1) \times t$, i.e. scaling the probabilities linearly with time.

Alternatively, one may wish to use all the information embedded in the migration matrix, taking into account that default probabilities are not constant over time, but increase as a result of downgrades. An approach advocated by the Task Force of the Market Operations Committee (2007) involves the computation of the root of the migration matrix from a decomposition in eigenvalues and eigenvectors.[8] This approach assumes that rating migrations are path-independent and the probabilities are constant over time. This is a very common assumption, despite empirical evidence to the contrary (see, for instance, Nickell *et al.* 2000).

Although theoretically appealing, finding and using the root of the migration matrix poses a number of problems in practice. These are fourfold:

- First, if one or more of the eigenvalues is negative (or even complex), then a real root of the migration matrix does not exist. The typical migration matrix is diagonally dominant – the largest probabilities in each row are on the diagonal – and therefore, in practice, its eigenvalues are real and positive, but this is not guaranteed.
- Second, the eigenvalues need not be unique. If this is the case, then the root of the migration matrix is not unique either. This situation raises the question which root and which short-duration PDs should be used. The choice can have a significant impact on the simulation results.
- There is a high likelihood that some of the eigenvectors have negative elements and, consequently, that the root matrix has negative elements as well. Clearly, in such cases, the root is no longer a valid migration matrix.
- Finally, even if, at a certain point in time, the root of the migration matrix exists, is unique and is a valid migration matrix, then it may still be of limited use if the main interest is in time series of credit risk measures.

Given these practical limitations, it seems better to use an approximation for the 'true' probability of default over short horizons. This can be done in several ways. One approach is to estimate a 'generator matrix' which, when extrapolated to a one-year horizon, approximates the original migration matrix as closely as possible (under some predefined criteria), while still

[8] Any $k \times k$ matrix has k (not necessarily distinct) eigenvalues and corresponding eigenvectors. If $\mathbf{C}$ is the matrix of eigenvectors and Λ is the matrix with eigenvalues on the diagonal and all other elements equal to zero, then any symmetric matrix $\mathbf{Y}$ (which has only real eigenvalues) can be written as $\mathbf{Y} = \mathbf{C}\Lambda\mathbf{C}^{-1}$ (where $\mathbf{C}^{-1}$ denotes the inverse of matrix $\mathbf{C}$). In special cases, a non-symmetric square matrix (such as a migration matrix) can be decomposed in the same way. The one-month migration matrix follows from $\mathbf{M} = \mathbf{Y}^{1/12} = \mathbf{C}\Lambda^{1/12}\mathbf{C}^{-1}$. The right column of $\mathbf{M}$ provides the monthly default probabilities. The matrix for other periods is found analogously.

respecting the conditions for a valid migration matrix (e.g. Israel *et al.* 2001; Kreinin and Sidelnikova 2001). An example of this is the approach adopted by one central bank in the peer group of Eurosystem central banks, which involves the computation of the 'closest three-month matrix generator' to the one-year matrix. It is calculated numerically by minimizing the sum of the squared differences between the original one-year migration probabilities and the one-year probabilities generated by raising the three-month matrix to the power of four. This three-month matrix provides plausible estimates of the short-term migration probabilities and also generates, in most situations, small but positive one-year default probabilities for highly rated issuers. Note, however, that also a numerical solution may not be unique or a global optimum.

Otherwise, within the peer group very different approaches are used to 'scale down' annual default probabilities. These range from scaling linearly with time to not scaling at all, i.e. applying annual default probabilities also to assets with shorter maturities, under the assumption that any position which matures before the end of the horizon, is rolled into a new position with the same obligor at all times. It is not uncommon to round the maturities of short duration positions upwards into multiples of one or three months.

The approach adopted by the ECB is based on the already discussed assumption that the conditional PD is constant over time. Hence, the PDs for maturities t are derived from the one-year probabilities only: $pd(t) = 1 - [1 - pd(1)]^t$. A limitation of this approach, as with any approach that ignores a large part of the migration matrix, is that it is impossible to differentiate between one-year positions held until maturity and shorter positions reinvested in assets of the same initial credit quality (which would be in a different name, if the original obligor had meanwhile been up- or downgraded). The implication is that the default risk of short positions is probably somewhat overstated. This conservative bias is however considered acceptable.[9]

For the actual implementation of this approach, the concept of multiple default states has been introduced. A default may occur in e.g. the first month of the simulation period, in the second month, and so on, leading to

[9] One justification for this bias is that a conservative investor like a central bank would normally sell a bond once it has been downgraded beyond a certain threshold. As this reduces risk in the actual portfolio, a buy-and-hold model of the portfolio will overestimate the credit risk. To some extent, the two simplifications offset each other. Note also that most, if not all, approximations used by the members of the peer group lead to conservative estimates of the 'true' short term probability of default.

different expected pay-offs, as some positions will have matured and coupons have been received if default occurs later in the year. Each one-year PD is broken down into probabilities for different sub-periods, and the last (default) column of the migration matrix is replaced by a number of columns with PDs for these sub-periods. This matrix is referred to as the 'augmented migration matrix'. The main benefit of this implementation is that long and short duration positions can be treated in a uniform way and, if needed, aggregated for individual names. Once the augmented migration matrix and the corresponding matrix of conditional forward values (step 0 of the simulation procedure) have been derived, it is not necessary to burden the program code with additional and inefficient if-when statements (e.g. to test whether a position has a maturity longer or shorter than the risk horizon).

An example may illustrate the concept of multiple default states. Consider again the migration probabilities from Table 3.1. The probability that a single-A issuer defaults over a one-year horizon (6 basis points) is broken down into probabilities for several sub-periods. Assume that the following sub-periods are distinguished: (0, 1m], (1m, 3m], (3m, 6m], (6m, 12m]. The choice of sub-periods is another 'art' and can be tailored to the needs of the user and restrictions implied by the portfolio; the example used here is reasonable, in particular for portfolios with a large share of one-month instruments (as the first portfolio considered in Section 4). The PD of the first sub-period is (conservatively) based on the upper boundary of its time interval (i.e. one month). The other PDs follow from

$$PD(t_1, t_2) = \underbrace{(1-p)^{t_1}}_{\substack{\text{probability}\\ \text{of survival}\\ \text{up to period } t_1}} \underbrace{\left[1-(1-p)^{t_2-t_1}\right]}_{\substack{\text{conditional probability}\\ \text{of default in period } t_1-t_2}} = (1-p)^{t_1}-(1-p)^{t_2}, \qquad 3.12$$

where p equals the one-year PD and t_1 and t_2 are the boundaries of the time interval. Note that these represent unconditional PDs. The augmented first row of Table 3.1 would look as shown in Table 3.3. Note that, as expected, the probabilities for the sub-periods are very close to the original one-year probability, scaled by the length of the time interval. The relationship is only approximate, as would become obvious if more decimals were shown (or if the original PD were larger). Note also that, by construction, the probabilities add up to unity. The corresponding standard normal boundaries are not shown in the table, as their derivation is the same as before.

Table 3.3 Original and augmented migration probabilities for bond with initial rating A

'State of the world'	D ≤ 1m	1m < D ≤ 3m	3m < D ≤ 6m	6m < D ≤ 1y	CCC-C	B	BB	BBB	A	AA	AAA
Original probability	0.06%				0.03%	0.17%	0.41%	5.63%	91.49%	2.17%	0.04%
Augmented probability	0.005%	0.010%	0.015%	0.030%	0.03%	0.17%	0.41%	5.63%	91.49%	2.17%	0.04%

Source: Standard & Poor's (2008a, Table 6 – adjusted for withdrawn ratings) and ECB calculations.

3.5 Recovery rates

The recovery rate measures the proportion of the principal value (and possibly accrued interest) that is recovered in the event of a default. The sensitivity of simulation results to this parameter depends on the relative contributions of default and migration to the overall risk estimate. Clearly, default is a fairly remote risk for highly rated portfolios. On the other hand, since migration risk increases almost linearly with duration, and central bank portfolios typically have low durations, the relative contribution of default risk to overall credit risk may still be substantial, in particular at very high confidence levels. The results of the simulation exercise in the next section confirm this intuition. Consequently, the choice of the recovery rate needs to be addressed carefully.

Among other things, the recovery on a bond or loan depends on its seniority and is likely to vary across industries. There is also evidence that recovery rates for bank loans are on average higher than for bonds (Moody's 2004). Some well-known empirical studies on recovery rates are from Asarnow and Edwards (1995), Altman and Kishore (1996), Carty and Lieberman (1996) and more recently Altman *et al.* (2005b). Rating agencies publish their own studies on recovery rates. Many studies conclude that recovery rates for senior loans and bonds are in the range 40–50 per cent. This is also the range that is used by most of the central banks in the ECB's peer group. When stochastic recovery rates are used, the mean is in the same range.

The ECB uses a uniform recovery rate, fixed at 40 per cent of the principal value. A uniform recovery rate is considered sufficiently accurate because the assets in the portfolio are fairly homogeneous, i.e. only senior debt and mostly government or government-related issuers.

3.6 Correlations

Correlation measures the extent to which companies default or migrate together. In the credit risk literature, the parameter often referred to is default correlation, formally defined as the correlation between default indicators (1 for default, 0 for non-default) over some period of time, typically one year. Default correlation can be either positive, for instance because firms in the same industry are exposed to the same suppliers or raw materials, or because firms in one country are exposed to the same exchange rate, but it can also be negative, when for example the elimination of a competitor increases another company's market share. Default correlation is difficult to estimate directly, simply because defaults, let alone correlated defaults, are rare events. Moreover, as illustrated by Lucas (2004), pair-wise default correlations are also insufficient to quantify credit risk in portfolios consisting of three assets or more. This is a consequence of the discrete nature of defaults. For these reasons, correlations of asset returns are used.

It is important to note that asset and default correlation are very different concepts. Default correlation is related non-linearly to asset correlation, and tends to be considerably lower (in absolute value).[10] While Basel II, for instance, proposes an asset correlation of up to 24 per cent,[11] default correlation is normally only a few per cent. Indeed, Lucas (2004) demonstrates that for default correlation the full range of -1 to $+1$ is only attainable under very special circumstances.

Other things being equal, risks become more concentrated as asset correlations increase, and the probability of multiple defaults or downgrades rises. With perfect correlation among all obligors, a portfolio behaves as a single bond. It should thus come as no surprise that the relationship between asset correlation and credit risk is positive (and non-linear). Figure 3.2 plots this relationship, using ES as risk measure, for a hypothetical portfolio.

Asset correlations are usually derived from equity returns. This is because asset returns cannot be observed directly, or only infrequently. In practice, it is neither possible nor necessary to estimate and use individual correlations

[10] The formal relationship between asset and default correlation depends on the joint distribution of the asset returns. For normally distributed asset returns, the relationship is given in, for instance, Gupton *et al.* (1997, equations 8.5 and 8.6).

[11] Under the Internal Ratings-Based Approach of Basel II, the formula for calculating risk-weighted assets is based on an asset correlation ρ equal to $\rho = w0.12 + (1 - w)0.24$, where $w = \frac{1-e^{-50pd}}{1-e^{-50}}$. Note that ρ decreases as pd increases.

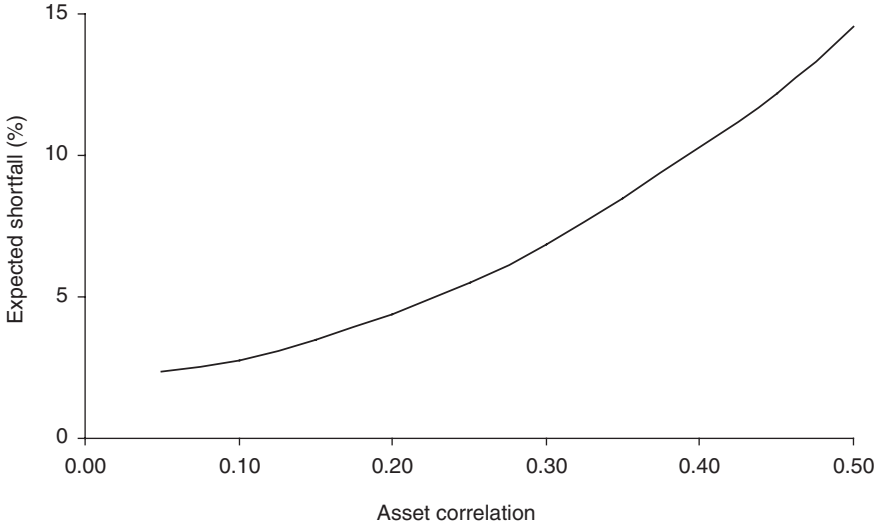

Figure 3.2 Impact of asset correlation on portfolio risk (hypothetical portfolio with 100 issuers rated AAA–A, confidence level 99.95%). *Source*: ECB's own calculations.

for each pair of obligors. First of all, scarcity of data limits the possibility of calculating large numbers of correlations ($n\,[n-1]\,/\,2$ for a portfolio of n obligors). Secondly, empirical evidence seems to indicate that sector concentration is more important than name concentration (see, for instance, BCBS 2006d). In order to capture the sector concentration, it is necessary to estimate intra-sector and inter-sector correlations, but it is not necessary to estimate each pair of intra-sector correlations individually. Inter-sector correlations can be estimated from equity indices using a factor model.

This approach has its limitations for central bank portfolios, which mainly consist of bonds issued by (unlisted) governments. Instead, the ECB model uses the 'Basel II level' of 24 per cent for all obligor pairs. Again, there is some variation in the choice of correlation levels among the peer-group members. For instance, some central banks prefer to use higher correlations, even up to 100 per cent, for seemingly closely related issuers, such as the US Treasury and Government Sponsored Enterprises (GSEs).

3.7 Credit spreads

The final parameter is the (forward) credit spread for different rating levels and for each currency in portfolio (USD, EUR and JPY). Spreads determine the mark-to-market loss (gain) in the event of a downgrade (upgrade). The ECB system loads spreads derived from Bloomberg data. Nelson and Siegel

(1987) curves are fitted in order to ensure a certain 'smoothness' in the spreads.

Finding reliable data can be a challenge, in particular for lower ratings and non-USD markets. For instance, Bloomberg provides USD 'fair market' credit curves with ratings down to B; for the EUR market, similar curves are only available for ratings BB or better. The same limitation applies to credit curves in Reuters. Fortunately, the sensitivity of output to these assumptions is minor, given the high credit quality of the portfolios and low one-year migration probabilities to non-investment grade. Still, in some cases, an 'expert opinion' on a reasonable level of spreads in certain non-investment grade markets is required.

4. Simulation results

The following sections present some empirical results for two very different portfolios. The first portfolio (in the following 'Portfolio I') is a subset of the ECB's USD portfolio, as it existed some time ago. The portfolio contains government bonds, bonds issued by the Bank for International Settlements (BIS), Government Sponsored Enterprises (GSEs) and supranational institutions – all rated AAA/Aaa – and short-term deposits with approximately thirty different counterparties, rated A or higher and with an assumed maturity of one month. Hence, the credit risk of the portfolio is expected to be low. The modified duration of the portfolio is low.

The other portfolio ('Portfolio II') is fictive. It contains more than sixty (mainly private) issuers, spread across regions, sectors, ratings as well as maturity. It is still relatively 'chunky' in the sense that the six largest issues make up almost 50 per cent of the portfolio, but otherwise more diversified than Portfolio I. It has a higher modified duration than Portfolio I. The lowest rating is B+/B1. Figures 3.3a and 3.3b compare the composition of the two portfolios, by rating as well as by sector (where the sector 'banking' includes positions in GSEs). From the distribution by rating, one would expect Portfolio II to be more risky.

These portfolios are identical to those analyzed in the paper published by the Task Force of the Market Operations Committee of the European System of Central Banks (2007), cited before. The analysis in that paper focused on a comparison of five different although similar credit risk systems, one of which was operated at the ECB. One of the findings of the original exercise was that different systems found fairly similar risk

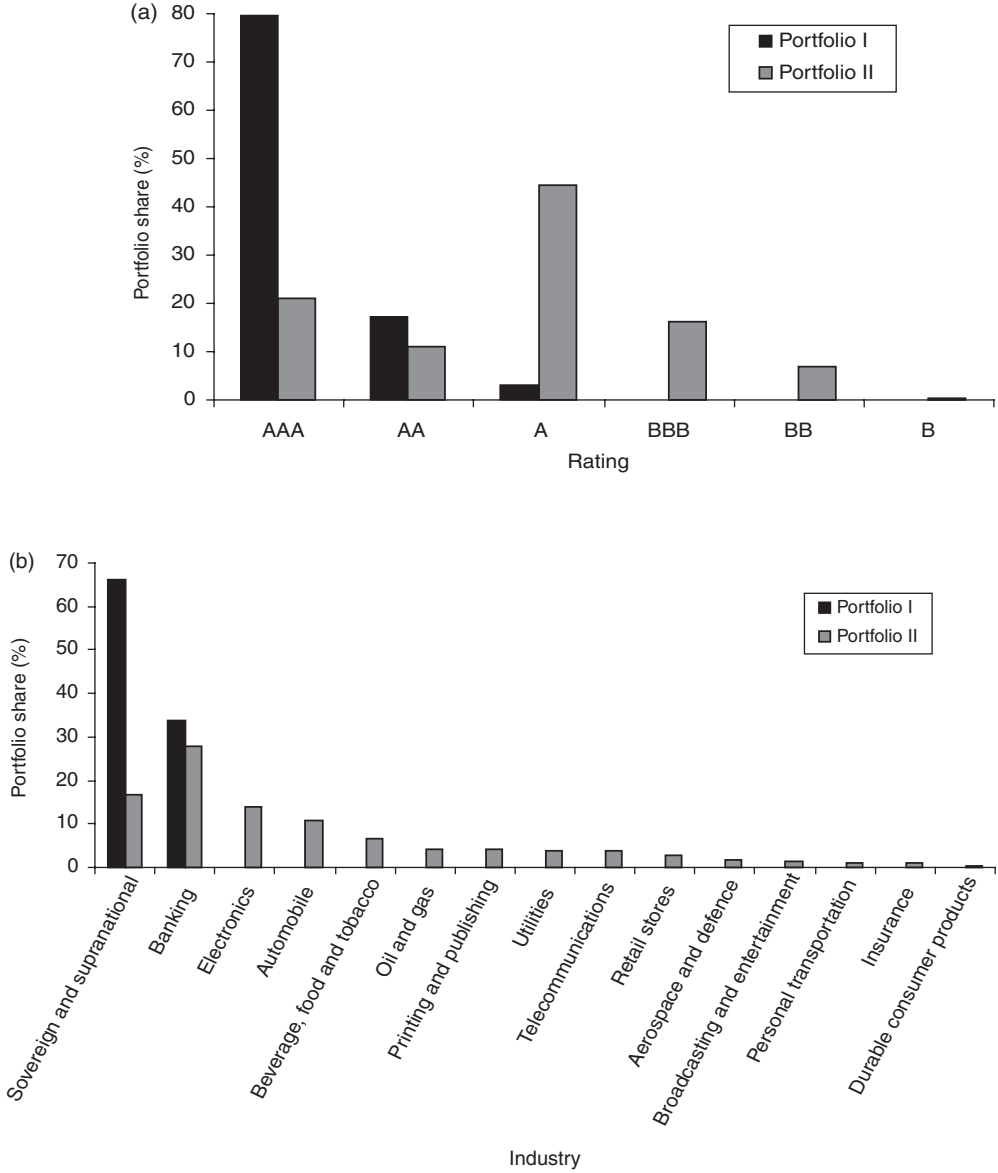

Figure 3.3 Comparison of portfolios by rating and by industry.

estimates, in particular at higher confidence levels which are the most relevant. In this book, the focus is on the ECB system only. The results presented in this section match the results in the paper that was published last year only approximately, because a new and improved system has since

been developed, taking into account some of the lessons from the earlier study. As in the paper, the results from a 'core' parameter set are compared with those obtained from various sensitivity analyses.

The simulation results include the following risk measures: expected loss, unexpected loss, VaR and ES, at various confidence levels and all for a one-year investment horizon, and the probability of at least one default. The inclusion of the latter is motivated by the belief that a default may have reputational consequences for a central bank invested in the defaulted company.

4.1 Portfolio I

This section presents the first simulation results for Portfolio I and introduces the common set of parameters, which are also used for Portfolio II (Section 4.2). The results provide a starting point for the scenario analysis in Section 4.3, and can also be used to analyse the impact of different modelling assumptions for parameters not prescribed by the parameter set, in particular short-horizon PDs. The common set includes a fixed recovery rate (40 per cent) and a uniform asset correlation (24 per cent). The credit migration matrix (Table 3.4) was obtained from Bucay and Rosen (1999), and is based on Standard & Poor's ratings, but with default probabilities for AAA and AA revised upwards (from zero) as in Ramaswamy (2004a) whereby the PD for AA has been set equal to the level of AA–. The augmented matrix (not shown) is derived from this matrix and is effectively of dimension 3×11: only the first three rows are needed because initial ratings are A or better. The number of columns is 11, because one default state is replaced by four sub-periods (those used in the example of Section 3.4). Spreads are derived from Nelson–Siegel curves (Nelson and Siegel 1987), where the zero-coupon rate $y^{cr}(t)$ for maturity t (in months) and credit rating cr is given by $y^{cr}(t) = \beta_1^{cr} + \left(\beta_2^{cr} + \beta_3^{cr}\right)\frac{1-e^{-\lambda t}}{\lambda t} - \beta_3^{cr} e^{-\lambda^{cr} t}$. The curve parameters are given in Table 3.5.

The main results are shown in Figure 3.4 and Table 3.6. The starting point for the analysis of the results is the validation of the models, using the analytical expressions for expected and unexpected loss given in equations (3.6) and (3.11), while keeping in mind that the results for Portfolio I are averages over different systems, based on different assumptions, in particular for the PD of short-duration assets. The analytical computations confirm the simulation results of Table 3.6: expected loss equals 1 basis point (i.e. the same as the simulated result); unexpected loss is around 27

Table 3.4 Common migration matrix (one-year migration probabilities)

To \ From	AAA	AA	A	BBB	BB	B	CCC/C	D
AAA	90.79%	8.30%	0.70%	0.10%	0.10%	–	–	0.01%
AA	0.70%	90.76%	7.70%	0.60%	0.10%	0.10%	–	0.04%
A	0.10%	2.40%	91.30%	5.20%	0.70%	0.20%	–	0.10%
BBB	–	0.30%	5.90%	87.40%	5.00%	1.10%	0.10%	0.20%
BB	–	0.10%	0.60%	7.70%	81.20%	8.40%	1.00%	1.00%
B	–	0.10%	0.20%	0.50%	6.90%	83.50%	3.90%	4.90%
CCC/C	0.20%	–	0.40%	1.20%	2.70%	11.70%	64.50%	19.30%
D	–	–	–	–	–	–	–	100.00%

Note: PD for AAA and AA adjusted as in Ramaswamy (2004a).
Source: Bucay and Rosen (1999).

Table 3.5 Parameters for Nelson–Siegel curves

	AAA	AA	A	BBB	BB	B	CCC/C
λ^{cr}	0.0600	0.0600	0.0600	0.0600	0.0600	0.0600	0.0600
β_1^{cr} (level)	0.0660	0.0663	0.0685	0.0718	0.0880	0.1015	0.1200
β_2^{cr} (slope)	−0.0176	−0.0142	−0.0149	−0.0158	−0.0242	−0.0254	−0.0274
β_3^{cr} (curvature)	−0.0038	−0.0052	−0.0061	−0.0069	−0.0139	−0.0130	−0.0080

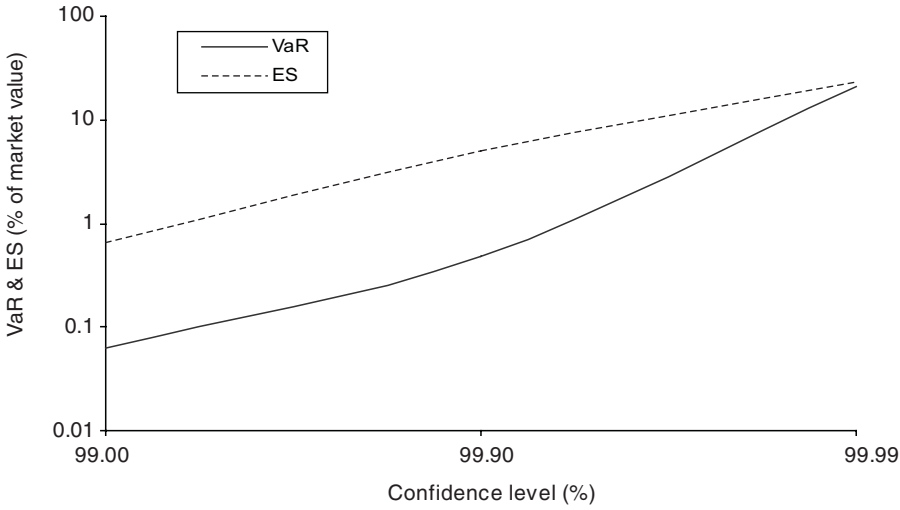

Figure 3.4 Simulation results for Portfolio I.

Table 3.6 Simulation results for Portfolio I

Expected loss		0.01%
Unexpected loss		0.28%
VaR	99.00%	0.06%
	99.90%	0.48%
	99.99%	21.27%
ES	99.00%	0.64%
	99.90%	4.98%
	99.99%	22.89%
Probability of at least one default		0.18%

basis points (vs. 28 basis points for the simulation). Further reassurance of the accuracy is obtained from the results in Table 3.7, which shows a decomposition of simulation results in the contributions of default and migration. This decomposition can be derived by running the model in 'default mode' with an adjusted migration matrix – setting all migration probabilities to zero, while increasing the probabilities that ratings remain unchanged and keeping PDs unchanged – and isolating the contribution of default. Nearly 50 per cent of expected loss, i.e. 0.5 basis point, can be attributed to default, which is easily and intuitively verified as follows: approximately 80 per cent of the portfolio is rated AAA, 17 per cent has a rating of AA and the remaining 3 per cent is rated A. Most AAA positions have a maturity of more than one year, while the (assumed) maturity of all AA and A positions is one month. If one multiplies these weights by the corresponding PDs (1, 4 and 10 basis points, respectively), scaled for shorter maturities and the loss given default (i.e. one minus recovery rate), then the expected loss in default mode and assuming a one-year maturity of deposits is approximately $(0.80 \times 0.0001 + 0.17 \times 0.0004 / 12 + 0.03 \times 0.0010 / 12) \times 0.6 = 0.5$ basis point.

The decomposition in Table 3.7 also shows that at lower confidence levels, migration is an important source of risk, but that default becomes more relevant as the confidence level increases. At 99.99 per cent, virtually all the risk comes from default.

From Table 3.6, a number of further interesting observations can be made. One of the first things that can be seen is that VaR and, to a lesser extent, ES are well contained until the 99.90 per cent level, but that these risk measures increase dramatically when the confidence level is raised to

Table 3.7 Decomposition of simulation results into default and migration

		Default	Migration
Expected loss		47.8%	52.2%
Unexpected loss		99.6%	0.4%
VaR	99.00%	–[a]	100.0%
	99.90%	52.6%	47.4%
	99.99%	100.0%	–
ES	99.00%	77.1%	22.9%
	99.90%	98.9%	1.1%
	99.99%	99.9%	0.1%

[a] At 99 per cent, there are no defaults. Recall that VaR has been defined as the tail loss *exceeding expected losses*. As a consequence, the model in default mode reports a negative VaR (i.e. a gain offsetting expected loss) at 99 per cent. For illustration, this result is shown in the table as a 0 per cent contribution from default (and, consequently, 100 per cent from migration).

99.99 per cent (which corresponds to the assumed probability of survival (non-default) of AAA-rated instruments, i.e. the majority of the portfolio). Evaluated at the 99.90 per cent confidence level, the CreditVaR is almost irrelevant when compared with the VaR for market risks (in particular currency and gold price risks). However, once the confidence level is raised to 99.99 per cent, credit risk becomes a significant source of risk too, with potential losses estimated in the region of 20 per cent of the portfolio. As confirmed by the results in Table 3.3, defaults have a significant impact on portfolio returns at this confidence level.

In order to determine the statistical significance of (differences in) simulation results, standard errors for the VaR estimates can be calculated. Standard errors are based on the observation that the number of scenarios with losses exceeding the VaR is a random variable which follows a binomial distribution with mean $n(1-\alpha)$ and standard deviation $\sqrt{n\alpha(1-\alpha)}$, where n equals the number of draws in the simulation and α corresponds to the confidence level of the VaR. For example, if the 99.99 per cent VaR is estimated from 100,000 simulations, then the expected number of scenarios with losses exceeding this VaR is $100{,}000 \times (1-0.9999) = 10$; the corresponding standard deviation equals 3.16. The simulation results above and below the VaR are recorded (using interpolation when the standard deviation is not an integer number) and the difference, expressed as a percentage

of the forward portfolio value (FV_p) and divided by two, is reported as the standard error. For very large samples, it is reasonable to approximate the distribution of the number of losses exceeding the VaR by a normal distribution, and conclude there is a 68 per cent probability that the 'true' VaR falls within one standard deviation around the estimated VaR. Note that the standard deviation of the binomial distribution increases with the number of iterations n, but that this value represents only the index of observations. As the number of iterations increases, individual simulation results are less dispersed. As a result, the standard error of the VaR estimates is expected to decrease as the number of iterations increases.

The reported standard errors indicate that the estimates of the 99.00 per cent and 99.90 per cent VaR are very accurate. After 100,000 iterations, the reported standard errors are practically 0. However, the uncertainty surrounding the VaR increases substantially as the confidence level rises to 99.99 per cent: after 1,000,000 iterations and without variance reduction techniques, the standard error is nearly 3 per cent. Increasing the number of iterations brings it down only very gradually. Not surprisingly, given the lack of data, simulation results at very high confidence levels should be treated with care.

For Portfolio I, with its large share of short duration assets, the probability of at least one default depends strongly on how one-year default probabilities are converted into shorter equivalents. For example, if the (very conservative) assumption had been made that the PDs of short-duration assets equal the one-year PDs, then the calculation would have been as follows. Portfolio I consists of six obligors rated AAA, twenty-two rated AA and eight rated A. If, for simplicity and illustration purposes, it is assumed that defaults occur independently, then it is easy to see that the probability of at least one default would be equal to $1-(1-0.01\%)^6\times(1-0.04\%)^{22}\times(1-0.10\%)^8 = 1.73\%$. However, under the more realistic assumption that the PD of all thirty AA and A obligors and two of the six AAA obligors (one-month deposits) equals only 1/12th of the annual probability, then the probability of at least one default reduces to $1-(1-0.01\%)^4\times(1-0.01\%\ /\ 12)^2\times(1-0.04\%\ /\ 12)^{22}\times(1-0.10\%\ /\ 12)^8 = 0.18\%$ only, in line with the results reported in Table 3.6.

The calculations in the previous paragraph are based on assumed default independence. Since these computations are concerned with default only, it is useful to discuss the impact of default correlation. Consider a very simple although rather extreme example of a portfolio composed of two issuers

A and *B*, each with a PD equal to 50 per cent.[12] If the two issuers default independently, then the probability of at least one default equals $1-(1-50\%)^2 = 75\%$. If, however, defaults are perfectly correlated, then the portfolio behaves as a single bond and the probability of at least one default is simply equal to 50 per cent. On the other hand, if there is perfect negative correlation of defaults, then if one issuer defaults, the other does not, and vice versa. Either *A* or *B* defaults and the probability of at least one default equals 100 per cent. It is a general result that the probability of at least one default decreases (non-linearly) as the default correlation increases. Note that this corresponds to a well-known result in structured finance, whereby the holder of the equity tranche of an asset pool, who suffers from the first default(s), is said to be 'long correlation'. Given the complexity of the computations with multiple issuers, it suffices to conclude that one should expect simulated probabilities of at least one default to be somewhat lower than the analytical equivalents based on zero correlation, but that, more importantly, the assumptions for short-duration assets can have a dramatic impact on this probability.

4.2 Portfolio II

Portfolio II has been designed in such a way as to reflect a portfolio for which credit risk is more relevant than for Portfolio I. It is therefore to be expected that risks are higher than in the previous section (see also Figures 3.3a and 3.3b). The simulation exercise is repeated for Portfolio II and to some extent similar observations can be made as for Portfolio I. For completeness, Table 3.8 summarizes the main results. It shows, among other things, that the contribution of default to overall risk is substantially lower than for Portfolio I, mainly because the duration of Portfolio II is higher. A second and less important reason is that credit spreads between A (average rating of Portfolio II) and BBB are somewhat larger than between AAA (bulk of Portfolio I) and AA. Note also that at 99.99 per cent, the contribution of migrations to VaR and ES is non-negligible (and higher than at 99.9 per cent). The most interesting part comes when the risk measures – in particular VaR and ES – are compared for the two portfolios. This is illustrated graphically in Figures 3.5a and 3.5b.

[12] This rather extreme probability of default is chosen for illustration purposes only, because perfect negative correlation is only possible with a probability of default equal to 50 per cent. The conclusions are still valid with other probabilities of default, but the example would be more complex. See also Lucas 2004.

Table 3.8 Simulation results for Portfolio II, including decomposition

			Breakdown	
		Loss	Default	Migration
Expected loss		0.18%	49.4%	50.6%
Unexpected loss		0.62%	79.2%	20.8%
VaR	99.00%	2.28%	85.5%	14.5%
	99.90%	8.60%	99.9%	0.1%
	99.99%	11.56%	92.9%	7.1%
ES	99.00%	4.20%	87.3%	12.7%
	99.90%	9.86%	94.4%	5.6%
	99.99%	13.47%	92.7%	7.3%
Probability of at least one default		12.0%		

Figures 3.5a and 3.5b show that while VaR and ES are higher than for Portfolio I at the 99.00 per cent and 99.90 per cent confidence levels (as expected), the numbers are actually *lower* at the 99.99 per cent confidence level. Note that, because of the logarithmic scale of the vertical axis, the difference at 99.99 per cent is actually substantial and much larger than it may visually seem. The explanation for this possibly surprising result is the same as for the steep rise in VaR and ES at the 99.99 per cent confidence level for Portfolio I: concentration. At very high confidence levels, credit risk is not driven by average ratings or credit quality, but by concentration. Even with low probabilities of default, at certain confidence levels defaults will happen, and when they do, the impact is more severe if the obligor has a large weight in the portfolio. Since Portfolio I is more concentrated in terms of the number as well as the share of individual obligors, its VaR and ES can indeed be higher than the risk of a portfolio with lower average ratings, such as Portfolio II. In other words, a high credit quality portfolio is not necessarily the least risky. Diversification matters, in particular at high confidence levels. This result is also discussed in Mausser and Rosen (2007). Another consequence of the better diversification is that the risk estimates are much more precise than for Portfolio I. For instance, after 1,000,000 iterations, the standard error of the 99.99 per cent VaR is only 8 basis points.

Figure 3.6 compares the concentration of Portfolios I and II. Lorenz curves plot the cumulative proportion of assets as a function of the cumulative proportion of obligors. An equally weighted, infinitely granular

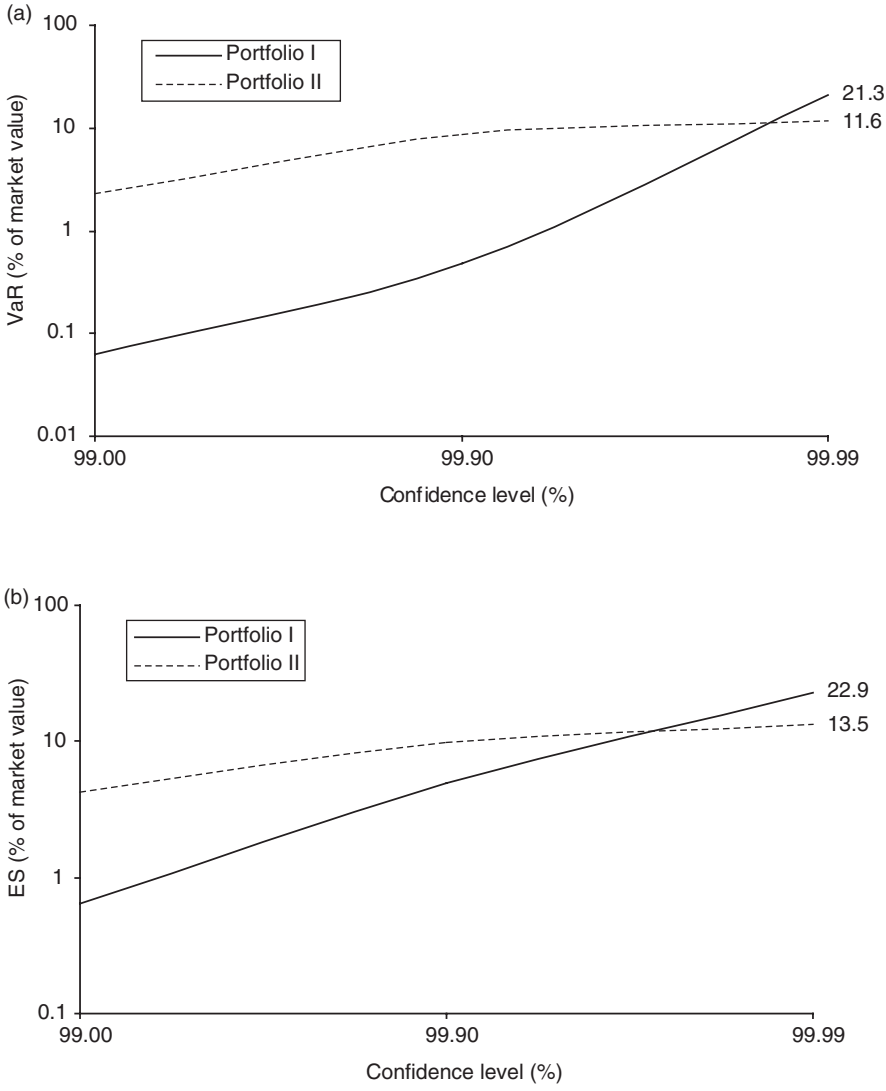

Figure 3.5 Comparison of simulation results for Portfolios I and II.

portfolio is represented by a straight diagonal line (note that such a portfolio may still be poorly diversified, as the obligors could all be concentrated in one sector, for instance); at the other extreme, a portfolio concentrated in a single obligor is represented by a horizontal line followed by an almost vertical line. The greater the disparity between the curve and the diagonal, the more the portfolio is concentrated. The figure confirms that Portfolio I is more concentrated than Portfolio II, which itself is also fairly concentrated.

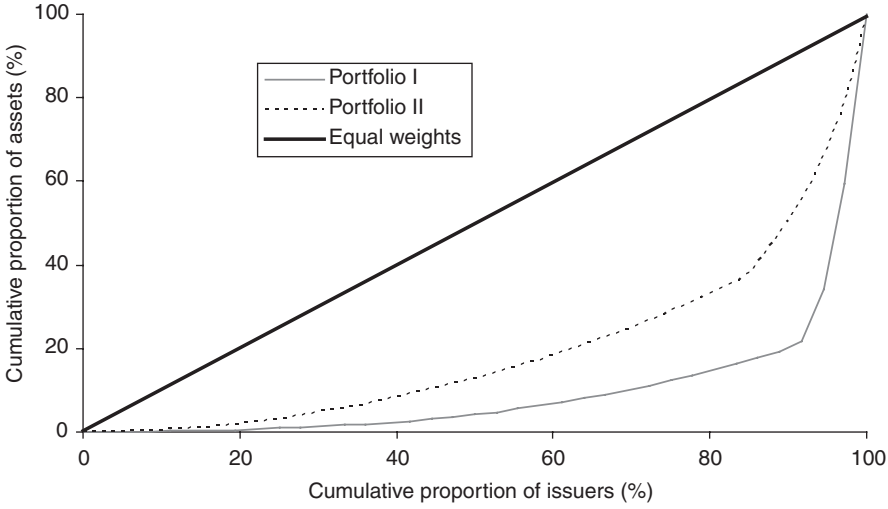

Figure 3.6 Lorenz curves for Portfolios I and II.

Note that the relative size of individual obligors does not affect the probability of at least one default, which is much higher for Portfolio II than for Portfolio I and rises to a level – around 12 per cent – that may concern investors who fear reputational consequences from a default in their portfolio. Statistically, this result is trivial: the larger the number of (independent) issuers in the portfolio, the larger the probability that at least one of them defaults. The probability of at least one default in a portfolio of n independent obligors, each with identical default probability pd, equals $1-(1-pd)^n$. For small n and pd, this probability can be approximated by $n \times pd$, and so rises almost linearly with the number of independent obligors. Clearly, increasing the number of independent obligors improves the diversification of the portfolio, reducing VaR and ES. It follows that financial risks (as measured by the VaR and ES) and reputational consequences (if these are related to the probability of at least one default) move in opposite directions as the number of obligors rises.

4.3 Sensitivity analysis

It is instructive to repeat the Monte Carlo simulations under alternative parameter assumptions in order to analyze the sensitivity of risk measures. This section discusses the impact of changes in the parameters that are to some extent discretionary, i.e. the probability of default for AAA issuers, the

Table 3.9 Sensitivity analysis for Portfolio I

	Base	PD(AAA) = 0.5 bp	PD(AAA) = 0	Correlation = 48%	Recovery = 20%
EL	0.01%	0.01%	0.01%	0.01%	0.01%
UL	0.28%	0.19%	0.03%	0.28%	0.36%
ES 99.9%	4.98%	3.01%	0.63%	5.10%	6.58%
Probability of at least one default	0.18%	0.16%	0.14%	0.15%	0.18%

recovery rate in case of default and the asset correlation between different obligors. The portfolio of attention is Portfolio I, which is particularly sensitive to one of the parameters. The following scenarios are considered:

- First, a reduction of the PD for AAA issuers from 1 basis point to 0.5 basis point per year.
- Second, it is assumed that all AAA issuers (mainly sovereigns and other public issuers) are default risk-free. Note that this does not mean these are considered completely credit risk-free, as downgrades and therefore marked-to-market losses over a one-year horizon might still occur.
- An increase in the asset correlation; as an (arguably) extreme case, the correlation doubles from 24 per cent (the maximum in Basel II) to 48 per cent. Note, however, that most of the issuers in Portfolio I are closely related – for instance the US government and Government Sponsored Enterprises – so that a somewhat higher average correlation than in the portfolios of other market participants could be justified.
- A reduction in the recovery rate from 40 to 20 per cent.

Note that while the last two bullets can be considered as stress scenarios, the first ones actually reduce risk estimates. This is because a PD equal to 1 basis point for AAA issuers is considered a stress scenario in itself that does not justify further increases in the PD assumptions. As already discussed in Section 3.4, the PD for AAA issuers is one of the key parameters of the model; analyzing the sensitivity of results to this parameter is essential. Lowering the PD in the sensitivity analyses is considered the most realistic way of doing so. The results are summarized in Table 3.9 and Figure 3.7.

From the results, a number of interesting conclusions can be drawn. The main, but hardly surprising observation is that ES (and similarly VaR) change dramatically if government bonds and other AAA issuers are assumed default risk-free. Other parameter variations have a much smaller impact on the results, although obviously each of these matters individually.

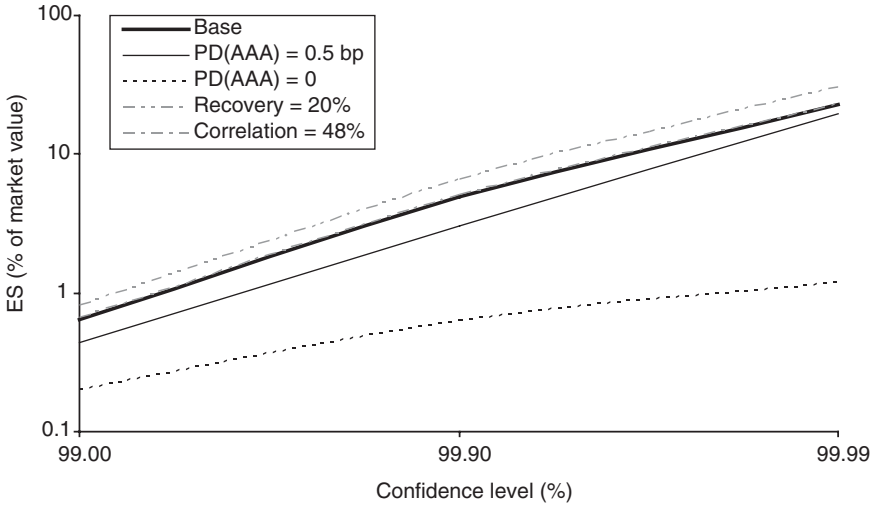

Figure 3.7 Sensitivity analysis for Portfolio I.

A change in the assumed recovery rate can have a significant, although not dramatic, impact on risk measures such as ES, but the influence of changes in correlations is very small; in fact, even when the correlation is doubled, a change in ES is hardly visible in Figure 3.7.

A similar analysis can be done for Portfolio II, but it does not add new insights. As Portfolio II contains only a minor share of government and other AAA bonds, the impact of alternative PD assumptions is much smaller than for Portfolio I.

5. Conclusions

Credit risk is gaining in importance within the central banking community and in many other public and private investors. From surveys of central bank reserves management practices that are published regularly, it is clear that many central banks are expanding into non-traditional assets, often implying more credit risk taking. Although growing, the proportion invested in credit instruments is likely to remain small, and the typical central bank portfolio continues to be concentrated in a small number of issuers only. If the assumption is made that these issuers carry some default risk, then it can be shown that the credit value at risk and expected shortfall of these portfolios may be higher than initially anticipated, once the confidence level is increased to very high levels.

Naturally, this observation crucially depends on the quality of the parameter assumptions. Empirically, defaults of AAA or AA rated issuers in one year time are (virtually) non-existent. For government issuers, these are even rarer. Assumptions of a positive probability of default are typically made rather ad hoc and not based on empirical evidence. Another issue is how to scale annual default probabilities for assets with shorter maturities, which make up a large proportion of most central bank portfolios. Both assumptions can have a substantial impact on the credit risk assessment, and are probably more critical for central banks than for commercial investors.

A portfolio credit risk model is a necessary but not a sufficient instrument for assessing overall risks, because defaults may have non-financial consequences as well. In particular, for a central bank or other public investor whose 'core' performance is not measured in financial return, a default by one of the issuers in the portfolio may damage the investor's reputation. Reputation risk is even harder to quantify than credit risk, but it may be tempting to use the probability of at least one default in the portfolio as a proxy. Paradoxically, perhaps, this probability rises as the portfolio is diversified into more issuers. One may argue that reputation is damaged more by a default of a large issuer than the default of a small issuer (in which case positive returns on the rest of the portfolio might more than compensate the loss), but the example illustrates that central banks sometimes need to make a trade-off between financial and reputation risk. Since reputation is their main asset, investment styles are likely to remain conservative (as they should).

All in all, though, the special characteristics of credit return distributions, the importance of credit risk models in commercial banks and the growth of the credit derivatives market are arguments for building up expertise in credit markets and credit risk modelling, also because there are clear spin-offs to other areas of the central bank, in particular those responsible for regulation and financial stability. The best way do to so is by being active in these markets with positions and by using a portfolio credit risk model. Central banks have substantial experience with market risk models already; credit risk modelling is a natural next step that also allows making these risks more comparable and, ultimately, integrating them. After all, central bank balance sheets and, hence, financial risks are largely dictated by their mandate, or reflect past developments. Currency and gold price risks can only be hedged to a limited extent, if at all, but it is important to know how these risks interact with other risks in the balance sheet.

4 Risk control, compliance monitoring and reporting

Andres Manzanares and Henrik Schwartzlose

1. Introduction

The aim of a risk control framework for a central bank's investment operations is to correctly measure and mitigate – or set an upper bound to – the financial risks arising from the investment activities of the bank and in particular from the holding of domestic and foreign currency investment portfolios.[1] The risk control framework of a central bank is ideally formulated and enforced by an organizationally independent unit that is separated from other business units and in particular from investment managers and bank supervisors.[2] Indeed, if staff involved in portfolio management report on risk and credit exposures, their unbiased measurement is not ensured. Similarly, eligibility and credit limits for investment operations should not be misinterpreted as giving indications of non-public knowledge about a certain counterparty derived from banking supervisory tasks. Chapter 1 dealt extensively with the considerations and specificities of public investors which are key inputs into the definition of a risk management framework for this type of investor. The present chapter deals with how these considerations are mapped into risk management policies and how these policies are made operational in terms of concrete risk management methodologies and infrastructure. The primary components of a sound risk management framework are the following: a comprehensive and independent risk measurement approach; a detailed structure of limits, guidelines and other parameters used to govern risk taking; and strong information systems for controlling, monitoring and reporting risks. The risk control framework defines what these risks are, how they are measured and how wide the admissible range for position taking is. This generally

[1] Note that risks associated with the conduct of monetary policy open market operations are handled in Chapter 8.
[2] See, e.g. (a) BCBS 2006a, (b) Counterparty Risk Management Policy Group II 2005, (c) BCBS 1998a.

implies that a risk control framework be rule-based, in order to objectively define risk and to set a transparent bound to the amount of voluntary risk taken.

The aims of the risk control framework are twofold:

(1) To map the actual risk–return preferences of the central bank senior management in the most accurate way into a set of rules that can be followed by risk-taking business units. This is mainly achieved by defining absolute limits (notably eligibility criteria and exposure bounds) and the benchmark investment portfolios.

(2) To define the leeway granted to investment managers for position taking (relative credit and market limits with respect to the benchmark portfolios) and ensure adherence to it. This step should be guided by senior management's confidence in investment managers' ability to achieve consistently better returns through active trading and by the institution's tolerance to financial losses.

Once appropriate benchmark portfolios have been defined, in a way that they are a feasible solution in view of the risk–return preferences, they can serve as a yardstick for the performance of the actual portfolios.[3] Profits or losses stemming from tactical investment decisions can be assessed using standard performance measurement tools and a short time horizon. The performance of the strategic currency and asset allocation may be discussed via a different analysis that takes into account the constraints imposed by policy requirements and a longer time horizon.

In order to illustrate the general perspective given in this chapter, many examples will refer to the ECB's risk management setup. For a better comprehension of these references, the ECB portfolio management structure is briefly introduced in the Section 2 of this chapter. Section 3 covers limit setting, introducing the different types of limits, some principles for deriving them and standard internal procedures for maintaining them. Section 4 is devoted to a number of tasks that are associated to compliance monitoring, reporting and related areas where, for the sake of transparency and independence, the risk management function should play an oversight role. Finally, sections on reporting (Section 5) and systems issues (Section 6) complete the picture of the risk control framework for investment operations.

[3] Performance is used throughout this book as relative return of the actual investment portfolio with respect to the benchmark portfolio.

2. Overview of the distribution of portfolio management tasks within the Eurosystem

The remaining sections of this chapter have a general central bank perspective. Each topic is then illustrated or contrasted by examples drawn from the ECB's risk management setup. In order to set the scene for these illustrations and to avoid unnecessary repetition, this section provides a brief overview of the ECB's portfolio management setup as of mid 2008. The ECB owns and manages two investment portfolios:[4]

- Foreign reserves. A portfolio of approximately EUR 35 billion, invested in liquid, high credit quality USD- and JPY-denominated fixed-income instruments. Managed by portfolio managers in the Eurosystem National Central Banks (NCBs).
- Own funds. A portfolio of approximately EUR 9 billion, invested in high credit quality EUR-denominated fixed-income instruments. Managed by portfolio managers located at the ECB in Frankfurt.

The Eurosystem comprises the ECB and the fifteen national central banks of the sovereign states that have agreed to transfer their monetary policy and adopt the euro as common single currency. The Eurosystem is governed by the decision-making bodies of the ECB, namely the Governing Council[5] and the Executive Board[6]. The foreign reserves of the ECB are balanced by euro-denominated liabilities vis-à-vis NCBs stemming from the original transfer of a part of their foreign reserves. As argued in Rogers 2004, this leads to foreign exchange risks being very significant, since the buffer generally provided by domestic currency denominated assets is unusually small, the

[4] The ECB holds a gold portfolio worth around EUR 10 billion which is not invested. The only activities related to this portfolio are periodic sales in the framework of the Central Bank Gold Agreement (CBGA). A fourth portfolio, namely the ECB staff pension fund, is managed by an external manager. These portfolios are not discussed further in this chapter.

[5] The Governing Council (GC) is the main decision-making body of the ECB. It consists of the six members of the Executive Board, plus the governors of the national central banks from the fifteen euro area countries. The main responsibilities are: 1) to adopt the guidelines and take the decisions necessary to ensure the performance of the tasks entrusted to the Eurosystem; and 2) to formulate monetary policy for the euro area (including decisions relating to monetary objectives, key interest rates, the supply of reserves in the Eurosystem, and the establishment of guidelines for the implementation of those decisions).

[6] The Executive Board (EB) consists of the President and Vice-President of the ECB and four additional members. All members are appointed by common accord of the Heads of State or Government of the euro area countries. The EB's main responsibilities are: 1) to prepare Governing Council meetings; 2) to implement monetary policy for the euro area in accordance with the guidelines specified and decisions taken by the Governing Council – in so doing, it gives the necessary instructions to the euro area NCBs; 3) to manage the day-to-day business of the ECB; and 4) to exercise certain powers delegated to it by the Governing Council – these include some of a regulatory nature.

ECB not being directly responsible for providing credit to the banking system nor for the issuance of banknotes. The ECB plays the role of decision maker and coordinator in the management of its foreign reserves. The investment framework and benchmarks (both strategic and tactical) are set centrally by the ECB, whereas the actual day-to-day portfolio management is carried out by portfolio managers located in twelve of the Euro Area National Central Banks (NCBs). Each NCB manages a portfolio of a size which generally corresponds to the proportion of the total ECB foreign reserves contributed by the country.[7]

From the outset in 1999 all NCBs managed both a USD and JPY portfolio. Following a rationalization exercise in early 2006, however, currently six NCBs manage only a USD portfolio, four manage only a JPY portfolio and two NCB's manage both a USD and a JPY portfolio. The currency distribution is fixed; in other words the NCB's managing both a USD and a JPY portfolio are not permitted to reallocate funds between the two portfolios.

The incentive structure applied vis-à-vis portfolio managers is limited to the regular reporting on return and performance and an associated 'league table', submitted regularly for information to the ECB decision-making bodies.

A three-tier benchmark structure applies to the management of each of the USD and JPY portfolios. In-house defined and maintained strategic benchmarks are reviewed annually (with a one-year investment horizon), tactical benchmarks monthly (with a three-month investment horizon) and day-to-day revisions of the actual portfolios take place as part of active management. The strategic benchmarks are prepared by the ECB's Risk Management Division and approved by the Executive Board (for the ECB's own funds) and by the Governing Council (for the foreign reserves). The tactical benchmarks for the foreign reserves are reviewed by the ECB's Investment Committee, where tactical positions are proposed among investment experts. While practically identical eligibility criteria apply for the benchmarks and actual portfolios, relative VaR tolerance bands permit the tactical benchmarks to deviate from the strategic benchmarks and the actual portfolios to deviate from the tactical benchmarks. Most portfolio managers tend to stay fairly close to the benchmarks; still, the setup ensures a certain level of diversification of portfolio management style, due to the

[7] Exceptions exist for some NCBs of countries that were not part of the Euro area from the outset which, for efficiency and cost reasons, chose to have their contributions managed by another NCB (as well as those NCBs that have received such a mandate). In particular, no portfolio management tasks related to the ECB's foreign reserves are conducted by the central banks of Malta, Cyprus and Slovenia.

full autonomy (within the investment guidelines and associated limits) given to portfolio managers. Settlement of transactions and collateral management is carried out by NCB back offices, in respect of accounts legally owned by the ECB.

The investment process for the ECB's own funds management is fully contained within the ECB. It has a two-tier structure. The benchmark is set internally, portfolio managers are ECB staff located in Frankfurt and transactions are settled by the ECB back office.

All risk management in relation to the foreign reserves and own funds is carried out centrally by the ECB's Risk Management Division (RMA). Supporting this overall setup is the ECB's portfolio management system located centrally at the ECB, but made available through a private wide area network to the portfolio managers located at NCBs (and at the ECB, for the own funds). This system permits all actual and benchmark positions, risk figures, limits etc. to be available on-line to NCB portfolio managers as well as investment and risk management staff located at the ECB. While information related to all portfolios is available to centrally located staff, information related to individual NCB portfolios is not shared among NCB portfolio managers. The system is also used for the management of the ECB's own-funds portfolio, but not for the individual reserves of Eurosystem central banks.[8]

3. Limits

We will use the usual breakdown of risks in an investment portfolio, although a central bank is a very idiosyncratic investor in this aspect. Risks faced by most commercial banks are roughly equally spread between credit and market risks (Nugée 2000). In contrast, central banks' financial concerns regarding investment portfolios are typically concentrated in the form of market risk, which generally is so large that it dwarfs credit risks.

3.1 Defining limits

Limits are arguably the main quantitative parameters contained in the risk management framework. They will determine, together with the benchmark

[8] Some of these NCBs use the same system from the same vendor for the management of their own reserves. However, this is run as a separate instance of the system, at the location of the NCB.

portfolio, the choice of the institution in the risk–return trade-off. Setting most risk limits requires a substantial amount of judgement from decision makers. The key task for risk managers in this respect should be to ensure that decision makers are well aware of the implications by providing them with an approximate feeling of both the impact of random events and their likelihood. It follows that risk managers should take the pulse of higher management's risk aversion periodically and in times of turbulence in financial markets, update the limits when needed. Risk management is responsible for:

- establishing risk policies, methodologies and procedures consistent with institution-wide policies;
- reviewing and approving (possibly also defining) models used for pricing and risk measurement;
- measuring financial risk across the organization as well as monitoring exposures to risk factors and movement in risk factors;
- enforcing limits with traders;
- communicating risk management results to senior management.

It goes without saying that market and credit risk methodologies must be solid and constrain risks adequately and leave no types of risk uncovered. On the other hand it is also for operational reasons important that the framework is coherent and without too many exceptions. If the framework is too complex, it will be difficult to communicate to management and portfolio managers and in the end most likely only a few people will fully understand it. A very complex framework also implies unnecessary operational risks and the cost of supporting it in terms of both ongoing staff resources and IT solutions becomes unnecessarily high. It is therefore important to find an appropriate trade-off between complexity (and equal treatment of risk types and risk-taking entities) and the ease with which the framework can be implemented and supported. The following subsections briefly describe the ECB's setup for each risk type.

3.2 Market risk limits

Market risk is defined, following the Basel II accord, as the risk of losses in on- and off-balance-sheet positions arising from movements in market prices. Market risk notably comprises foreign exchange and interest rate risks.[9] For multi-currency portfolios where tactical flows between currencies

[9] Gold, which still plays an important role in central banks' asset allocation schemes, may be considered as a currency or as a commodity. Potential losses due to changes in its market price are also considered market risk.

are admissible from a policy viewpoint, a way to put into work the principles outlined above is to define a benchmark that sets the currency composition and the structure of each currency sub-portfolio. Active portfolio managers may then take not only tactical curve and credit positions within each currency sub-portfolio but also foreign exchange positions. Alternatively, the currency composition may be set as fixed, thus reducing the leeway granted to portfolio managers to curve and credit positions with respect to the benchmark of each currency sub-portfolio. The latter option is preferred by the Eurosystem.

Being able to account for all types of market risks including foreign exchange risk is a major feature of VaR versus previously popular risk measures, which is especially important for most central banks where the latter represents the bulk of total financial risks. Box 4.1 elaborates further on this comparison.

Box 4.1. Modified duration versus VaR

Traditionally, modified duration and convexity have been used for measuring the sensitivity of fixed-income portfolios to yield-curve changes. An approach to the risk control of the foreign currency denominated portfolios of central banks would thus be to measure separately their interest rate sensitivity and combine it with some kind of volatility measure for exchange rates. Such an approach has two important drawbacks. First, a stepwise measurement of risks (first, interest rate and then foreign exchange risks) does not allow to summarize properly the total potential losses and to account for correlations between shocks to the yield curve and shocks to the exchange rate. Second, the variety of yield-curve positions that can be taken for tactical investment positions is not captured by one single risk measure and hence complicates the framework for relative risk control.

A clear trend in risk management theory over recent years has therefore been to aim at a comprehensive coverage of risks, such as to capture the total effects and inter-linkages among potential adverse events. In this respect, VaR allows to incorporate such linkages by means of estimated variance-covariance matrices that are then used as an indication, in terms of likelihood, of how the effect of different risk factors on a portfolio can either cumulate or offset each other. In contrast to modified duration, where one single risk factor (sovereign yield-curve parallel shifts) is the source of all the risk accounted for, VaR can encompass in principle all risk factors such as term structure and credit spreads, as well as exchange rate changes. Modified duration fails to reflect the risk effect of yield-curve positions with no duration impact and may overestimate the risk involved in duration positions by possibly ignoring other risk-offsetting effects. Moreover, limits defined by means of VaR have the advantage of automatically becoming more restrictive in terms of duration whenever bond markets become more volatile. Indeed, VaR has become the industry standard for measurement of market risk, despite some limitations and disadvantages. VaR is by definition a transparent and easy way to grasp the importance of

Box 4.1. (cont.)

risks incurred by holding a portfolio. Since this measure has become the industry benchmark, it has the important property of comparability.

All these arguments tend to suggest the use of relative VaR for measuring relative risks of the actual vis-à-vis the benchmark portfolio. Relative VaR is simply the VaR of the rescaled difference portfolio. In other words, the benchmark portfolio is rescaled to have the same market value as the actual portfolio, and then the difference portfolio is created, i.e. long the actual portfolio and short the rescaled benchmark portfolio. Finally, the absolute VaR of the latter difference portfolio is obtained. The VaR may also be expressed as a percentage of the actual portfolio's market value. For the kind of portfolio that most central banks hold, with the bulk placed on money market and fixed rate instruments, parametric delta-normal VaR, estimated through exponentially weighted moving averages (EWMA) as introduced by JPMorgan in 1993 may be a good choice regarding the estimation method for VaR. This method assumes joint normality of returns for a grid of fixed-term interest and exchange rates (the risk factors), which allows the variance of the whole portfolio losses to be estimated (generally assuming zero mean future returns) at a small computational cost. However, if the use of derivatives is more extensive and, in particular, if options are eligible, there may be a case for computing VaR through Monte Carlo simulations. In order to test the appropriateness of the Gaussian assumptions, a number of statistical tests can be used, which have been developed to monitor the reliability of reported VaR figures for regulatory purposes. See Campbell (2006) for a review and references.

Alexander (1999) or Jorion (2006) provide good overviews on market risk measurement. Market risks are nowadays commonly measured by VaR, which is simply an estimate of a given quantile, e.g. 95 per cent, of the probability distribution of losses in a portfolio due to price developments.[10] Specifically, if we denote with F_h the cumulative probability distribution function of the h-period portfolio return, conditional on all the information on portfolio composition at a given time, the h-period VaR at confidence level α is defined as: $VaR_t^h(\alpha) = -F_h^{-1}(\alpha)$. The negative sign is a normalization in order for losses to be positive numbers. The VaR of a portfolio is sometimes called *absolute* VaR to stress the fact that the probability distribution considered refers to total returns. If, on the contrary, the main concern is the size of potential losses of the portfolio considered in comparison to a benchmark portfolio of the same current market value serving as a yardstick for performance, the relevant risk measure is called *relative* VaR.[11] Relative

[10] An excellent account of the rise of VaR as industry standard and a general overview of market risk measurement is given in Dowd (2005).

[11] Relative VaR is a measure of the risk of losses with respect to the benchmark result and is defined as the VaR of the difference portfolio (i.e. actual minus the market-value-scaled benchmark portfolio). Relative VaR is sometimes called differential VaR. See Mina and Xiao (2001) for details.

VaR allows the measurement of the aggregate risk born, relative to the benchmark, in the domestic currency.

In the case of the ECB the lion's share of the market risk faced is due to the potential losses incurred on the foreign reserve portfolios in case of appreciation of the euro.[12] The ECB is, as the first line of defense of the Eurosystem in case of intervention in the foreign exchange market (article 30 of the ESCB Statute), constrained to hold large amounts of liquid foreign currency assets with no currency risk hedge. The currency choice and distribution of these reserves are determined on the basis of policy considerations only secondarily concerned with financial risk control. The latter concern is reflected in the periodic adjustments of the foreign reserves currency composition, which consider, among other things, risk–return aspects in the allocation proposals. Furthermore, foreign exchange risk is buffered in accounting terms through revaluation accounts and through an additional general risk provision. Once the strategic benchmark portfolio characteristics are set for a whole year,[13] active market risk management is mainly confined to monitoring and controlling the additional market risk induced by positions taken vis-à-vis the strategic and tactical benchmarks.

3.2.1 The ECB's market risk control framework

First, risk from *foreign exchange rates* (FX risk) is not actively managed in the sense that it is accepted as resulting from policy tasks. The allocation between JPY and USD is however also influenced by risk considerations. Second, the *strategic benchmarks* are set to comply with a no-loss constraint over an annual horizon at a certain confidence level (i.e. also a kind of VaR constraint; see Chapter 2). Third, market risk limits for the *active layers* of ECB investment operations are set in the form of relative VaR limits, namely ten basis points for the relative VaR of the USD and JPY tactical benchmarks vis-à-vis their respective strategic benchmarks and to five basis points for the USD and JPY actual portfolios vis-à-vis their respective tactical benchmarks. As regards the ECB's own-funds portfolio denominated in EUR, portfolio managers may deviate by up to five basis points from the strategic benchmark.[14] All these relative VaR limits refer to a daily horizon and a 99 per cent confidence level. The value of the relative VaR limits is

[12] VaR due to foreign exchange risk is much higher than VaR stemming from interest rate and spread changes, by a factor of around fifteen.

[13] The task of defining the strategic benchmark portfolios in each currency on an annual basis that satisfy the risk–return preferences of the decision-making bodies is described in detail in Section 7 of this chapter.

[14] There is no tactical benchmark for the ECB's own-funds portfolio.

approved by the Executive Board for the own funds and the Governing Council for the foreign reserves based on proposals prepared by the Risk Management Division. The level of relative VaR limits may be reviewed at any time should the ECB consider it necessary; however in practice the limits change only infrequently.

The implementation of market risk limits based on relative VaR requires appropriate risk measurement IT systems. The ECB uses market data provided by RiskMetrics, which is a widely recognized data and software provider for risk measurement purposes. The decay factor is set to 1 (no decay) and a relatively long period for estimation of the variance–covariance matrix is applied (two years). The latter parameter choices lead to a very stable estimate of VaR over time. This has the advantage of smoothing away high frequency noise and the disadvantage of possibly disregarding meaningful peaks and troughs in volatility which are relevant for risks.

3.3 Credit risk limits

Credit risk is the potential that a counterparty or debt borrower will fail to meet its obligations in accordance with contractual terms. Counterparties of central banks are generally highly rated, high credit quality banks and therefore the credit risk of a reserves portfolio is quite different from that of a commercial bank with a large part of its assets in the form of loans to unrated borrowers.

In the case of central banks, credit risk incurred for *investment* operations can be largely mitigated through a comprehensive credit risk framework aiming to minimize financial losses deriving from exposures to insolvency. In fact, there is a priori no policy constraint on the investment universe or the treatment of counterparty risk, and thus a risk-averse central bank can generally aim at holding a high credit quality portfolio and low counterparty risk.

The nature of the credit control framework is typically prudential and its aim twofold: first, it sets a minimum standard for creditworthiness among debtors and counterparties and second, it forces a minimal dispersion of credit risks in the portfolio in order to avoid high exposure concentration (in terms of sectoral or geographical location, taking into account the likelihood of idiosyncratic financial market shocks). Also, when defining eligibility conditions, reputation risks play a role: a loss realized by a default of a single name may trigger public reactions incommensurate with the size of the actual financial cost.

The first of these aims is typically laid down in the conditions for eligibility of instruments, issuers, issues and counterparties. Central banks typically restrict the eligible investment universe according to both market depth and credit quality considerations. Putnam (2004) argues that, while credit quality constraints are necessary if the concern is to avoid a highly negatively skewed return distribution, they are not per se a protection against absence of liquidity in a crisis situation. However, institutions are reasonably vague as to the exact degree of risk aversion, while asserting their preference for prudent management (Ramaswamy 2004b).

The second aim is achieved by adding, on top of these conditions, numerical limits for credit exposures to countries, issuers, and counterparties and procedures to calculate these exposures and ensuring compliance both in the actual and in the benchmark portfolio.[15] These criteria are a result of mapping the banks' perceived tolerance to credit (and associated reputational) risk into a working scheme for managing credit risk concentrations.

3.3.1 Credit quality and size as key inputs to limit setting formulas

A simple rule-of-thumb threshold system for setting credit exposure limits to counterparties can be easily defined. In essence, it consists of selecting a 'limit' function $L(Q, S)$, whereby Q is the chosen measure of credit quality, while S is a size measure, such as capital. Limits are non-decreasing in both input variables. The size measure typically aims at avoiding to build up disproportionate exposure to some counterparties, issuers, countries or markets.

The importance of credit quality in determining eligibility and setting limits is obvious. Chapter 3 introduced the relevant concepts, as they are also needed for credit portfolio modelling. Typically, credit quality is understood to mean mainly probability of default for classical debt instruments, while it also incorporates tail measures for covered bonds and structured instruments.

As is always the case when the probability distribution a random variable is summarized by a one-dimensional statistic, possibly critical information is bound to be lost. In the case of credit risks, which can be assumed to have very skewed distributions, using probabilities of default disregards conditional

[15] As a general principle, an instrument's exposure should always be calculated as its mark-to-market value (in the case of derivatives, its replacement cost). The calculation of market values in a separate risk management system may be data and time consuming. This is why a tight integration of the systems used in the front, middle and back office greatly simplifies the oversight of compliance with credit risk limits.

loss sizes. Using systematically expected loss fails to provide an adequate solution, since it does not reflect the dispersion of losses.[16] On the other hand, there is a need to map credit rating into limits and a yes/no decision on eligibility, and it is definitely more intuitive to do the latter from a one-dimensional credit scale.

How to measure credit quality for an eligibility and limit-setting framework? The most evident way, followed probably by all institutional investors, is to rely on ratings of major international rating agencies. These have large coverage, are cheaply available for investors, and investors may trust that rating agencies maintain a high quality of their ratings as their business model relies to a large extent on their brand name capital. Whether this is all what is needed for an institutional investor who has no ambitions to enter the world of non-rated companies is another question. External ratings are normally intended to be valid for a long period of time ('through the cycle') and generally do not react to small movements in the risk profile of the institution. Ratings are usually changed when it is unlikely that they will be reversed in a short period of time.[17] For this reason, credit risk ratings normally lag market developments and may experience significant swings in times of financial crisis (as has been observed several times in the past). As a result, ratings may not be the most efficient short-term predictors of default or changes in credit quality.[18] Furthermore, the application of a through-the-cycle approach and efforts to avoid rating reversals by the rating agencies lead to ratings which are relatively stable but show serial correlation in rating changes. What can thus be done to complement reliance on external credit ratings?

First, one may invest resources into **understanding rating methodologies**, in order to spot possible weaknesses that could be considered in the credit risk control framework. Although having been criticized heavily for

[16] Assume we had two issuers, one with a 10 bp probability of losing 10 per cent of the investment, the other with a 1 bp probability of losing 100 per cent. The expected loss, and hence the rating, would be the same, but the risk would definitely not be the same.

[17] 'Rating agencies state that they take a rating action only when it is unlikely to be reversed shortly afterward. Based on a formal representation of the rating process, it has been shown that such a policy provides a good explanation for the empirical evidence: Rating changes occur relatively seldom, exhibit serial dependence, and lag changes in the issuer' default risk.' (Löffler 2005)

[18] 'Rating stability has facilitated the use of ratings in the market for a variety of applications. As a result, rating changes can have substantial economic consequences for a wide variety of debt issuers and investors. Changes in ratings should therefore be made only when an issuer's relative fundamental creditworthiness has changed and the change is unlikely to be reversed within a short period of time. By introducing a second objective, rating stability, into rating system management, some accuracy with respect to short-term default prediction may be sacrificed.' (Moody's 2003)

their part in the US sub-prime crisis of 2007 (having inflated the use of high ratings for structured products that later on exhibited large losses), rating agencies tend to publish extensively on their rating methodologies and on defaults statistics. Even if these publications do not answer all questions, mostly the wealth of information provided is already more than what can be digested by a smaller risk management unit. Understanding what is behind ratings allows to better analyse the appropriate level of the rating threshold, and may also be useful for an efficient aggregation of ratings.

Second, one may aim at **understanding main factors driving the relevant industries**. For instance, it is necessary to have a fair degree of understanding of what is going on in the banking system, in covered bonds, in structured finance, in corporates, or in MBSs if one is invested in those markets. This is a pre-condition to be able to react quickly in case credit issues arise.

Third, one may **monitor market measures of credit risk** such as bond or Credit Default Swaps (CDS) spreads. These aggregate views of market participants in a rather efficient way, and obviously may react much earlier than credit ratings. Of course, by nature, monitoring those also does not put an investor in front of the curve as the information will then already be priced in. Still, it is better to be slightly behind the curve than not even to be aware of market developments. It is not obvious to incorporate market risk indicators directly in a limit-setting formula, but they can at least be used to trigger discussions which can lead to an exclusion of a counterparty or issuer or to a lowering of a limit.

Finally, one can set up an **internal credit rating system** on the basis of public information on companies, such as balance sheet information, complemented by high-frequency news on the company (see e.g. Tabakis and Vinci 2002). Such a monitoring requires substantial expertise and is therefore costly. It will probably make sense only for larger and somewhat lower rated investments. Ramaswamy (2004b) indicates that, due to the availability of rating scores issued by the major rating agencies for most, if not all, of a central banks' counterparties, the development of an internal credit rating system is generally too cost intensive compared to its marginal benefits.

If relying on ratings by several rating agencies, it is also crucial to **aggregate ratings** of the selected major rating agencies in the most efficient way such as to obtain a good aggregate rating index. The investor (in this case the central bank) has an idea of what minimum credit quality it would like to accept, expressed in its preferred underlying risk measure. By considering the methodological differences in producing ratings, the rating

scales used by the different rating agencies can in principle be mapped into the preferred underlying measure of credit quality. In other words, the investor's focus on the preferred underlying risk measure requires translating ratings from the scale they were formulated in by the rating agency to the investor's scale. This may be formulated as estimating an 'interpretation bias'. Concretely, one may assume that the preferred credit quality measure can be represented in scale from one to ten. In the master scale of the central bank, a rating of ten would correspond to the highest credit quality AAA, a nine to the next highest one, etc, and 1 to the lowest investment grade rating. Rating agencies may use similar scales, a rating by a certain rating agency of e.g. 'nine' may in fact correspond in the central bank's master scale to the ratings 8 and 9, and could thus be interpreted to mean an 8.5 rating in this master scale. Ratings are noisy estimates of the preferred credit quality in the sense that they are thought (in our simplistic approach) to be, for $i = 1, 2, \ldots, n$, $j = 1, \ldots, m$:

$$R_{j,i} = R_j + b_i + \varepsilon_{j,i}$$

where $i =$ the counter for the rating agency; $j =$ the counter for the counterparty; $n =$ the number of eligible rating agencies; $m =$ the number of names (obligors or securities); $R_{j,i} =$ the estimated credit quality by agency i of the rated counterparty j expressed in the 1 to 10 rating scale of this agency; $R_j =$ the preferred credit quality of the rated counterparty j, as expressed in the central bank's master scale; $b_i =$ the constant, additive 'bias' of rating agency i, in the sense that if the rating agency i provides a rating of e.g. 'seven', this could mean in terms of PDs of the central bank's master scale a 'six', such that the bias of the rating would be '+1'; $\varepsilon_i =$ are independent random variables distributed with cumulative distribution functions F_i, respectively.

A rating aggregation rule in the context of establishing eligibility is to be understood as follows. First, it is assumed that the central bank would like to make eligible all names having an expected preferred rating measure of above a certain threshold T. For instance, T could correspond to a 'six' in its master scale. Then, an aggregation rule is simply a rule that defines a composite rating out of the available ratings of rating agencies, and makes the name eligible if and only if the composite exceeds the threshold. The composite can also be bound to be an integer, but does not have to. Generally, a rating aggregation rule C is a function from the available ratings to a real number in [1,10], whereby the non-existence of a rating is

considered to be a specific input to be considered explicitly in the mapping. There are two main desirable characteristics of rating aggregation rules:

- Unbiasedness, i.e. the composite rating should be an unbiased estimator of the preferred rating measure, $E\lfloor C(R_{j,1}, \ldots, R_{j,n})\rfloor = R_j$;
- Efficiency: among all unbiased statistics, an efficient one should minimize the variance, i.e. C minimizes $E\lfloor (C(R_{j,1}, \ldots, R_{j,n})] - R_j)^2\rfloor$ subject to $E\lfloor C(R_{j,1}, \ldots, R_{j,n})\rfloor = R_j$.

If there were no rounding issues, and one knew the bias and standard deviations of the individual rating agencies, then the optimal rating aggregation rule would obviously be[19]

$$C_j* = \sum_{i=1}^{n} (R_{j,i} - b_i) \frac{1/\sigma_i^2}{\sum_{k=1}^{n} 1/\sigma_k^2}$$

Despite its theoretical optimality, this rule is rarely used in practice. Why? First, there may be a lack of knowledge on biases or on diverging standard errors of ratings. Second, rounding obviously creates complications. Third, complexities arise due to the need to take strong assumptions for averaging essentially qualitative 'opinions'. Fourth, the term bias in this model setup may be wrongly interpreted as a rating bias by an agency, rather than a correction for the fact that different risk measures are being compared, thus reducing the transparency of the process.[20] Alternative aggregation rules may be classified in the most basic way as follows:

(i) *Discrimination or not between rating rules: (A)* Rules discriminating between rating agencies: through re-mapping (i.e. recognizing non-zero values of the b_is, through different weights (justified by different variances of error terms), through not being content with the ratings of only one specific agency, etc.); (B) Rules not discriminating between agencies.

(ii) *Aggregation technique: (A)* Averaging rules: weighted or unweighted averages; (B) n-th best rules: first best, second best, third best, worst.

[19] This can be derived by considering, for every counterparty j, the linear regression of the vector of size n, $R_{j,i} - b_i$ (as the dependent variable) to the constant vector of size n, 1 as regressor. By assumption, the variance–covariance matrix of the error terms is a diagonal matrix with elements σ_i^2. The best linear unbiased estimator (BLUE) is then given by C_j^*.

[20] It is difficult to imagine that a 'true' rating, if it existed and could be properly expressed in one dimension, would be purposely missed by a rating agency, normally conscious to maintain its brand name. Another issue is whether rating agencies have enough information available to estimate reliably such measures as probability of defaults, given that the latter are very rare events. Even the most reasonable assumptions can turn out to be wrong in such a context.

(iii) *Require a minimum number of ratings or not. (A)* Do not care about number of ratings; (B) Require a minimum number of ratings, or at least require a better e.g. average if the number of ratings is low.

As it is difficult to find analytical solutions for establishing which of the rules above are biased or inefficient and to what extent, it is easiest to simulate the properties of the rules by looking at how they behave under various assumptions in terms of number and coverage of eligible rating agencies, relative biases between rating agencies, possible assumptions about the extent of noise in the different ratings, etc. Simulations conducted in the ECB suggest that the second-best rating rule performs well under realistic assumptions, and is also rather robust to changes in the rating environment.

3.3.2 Exposure calculation

Monitoring compliance with limits requires calculating the exposures consuming these limits. While this is trivial for some instruments, it is complex for others, and requires substantial analysis of many practical details. As a basic principle, all exposures should be calculated on a mark-to-market basis.[21] The exposure of counterparty limits is impacted by a range of transactions (securities and FX outright transactions, deposit and repurchase agreements, futures and swaps). For outright operations (purchases or sales of securities), replacement values should affect exposure to counterparties before settlement of the operation. For reverse repo operations, exposures arise if temporarily the value of the collateral falls short of the cash leg (which is very rarely the case, at least in case of haircuts and low trigger-levels for margin calls). Moreover, one may consider that concentration is an issue for repo operations. This can be addressed for instance by setting an (artificial) exposure coefficient for reverse repos and repos. For instance, one may set that 5 per cent of the notional repo and reverse repo values are deemed to be exposure to the repo counterparty, and consume the existing limit to the counterparty.[22] For reverse repos, the question also arises how *issuer* limits (i.e. the limit of the issuer of the collateral used) are affected if non-Government bonds are accepted as collateral. It is generally difficult to quantify the credit risk of accepting non-Government collateral, since 'double default' (of the issuer as well as the repo counterparty) risk depends on the default correlation between the two parties, for which there

[21] This implies demanding technical requirements on the portfolio management system, which needs to be able to compute exposures dynamically taking into account real-time transactions and market movements.

[22] An alternative to address concentration risk from repo and reverse repo operations is to set maximum total volumes towards each counterparty.

is limited data. Formally, the joint probability of default for the counterparty as well as issuer is $PD(cpy \cap issuer)$ is given by[23]

$$\begin{aligned} PD(cpy \cap issuer) &= PD(cpy) \times PD(issuer) + \rho \times \sigma(cpy) \times (issuer) \\ &= PD(cpy) \times PD(issuer) + \rho \times \sqrt{PD(cpy) \times (1 - PD(cpy))} \\ &\quad \times \sqrt{PD(issuer) \times (1 - PD(issuer))} \\ &\approx \rho \times \sqrt{PD(cpy) \times PD(issuer)} \end{aligned}$$

if $PD(cpy)$ and $PD(issuer)$ are both small (a natural assumption, given the eligibility criteria for counterparties and issuers) and where ρ is the default correlation. Note that the joint default probability is almost a linear function of a rather uncertain parameter ρ, and that this joint probability is likely to be very small, for every reasonable level of the univariate PDs.

Finally, for OTC derivatives like interest rate swaps, the choice has to be made between considering only actual mark-to-market values to affect exposures, or also potential market value (at some horizon and some confidence level). For interest rate swaps, the ECB has opted to only consider actual market values because, beyond a certain value, collateralization is required.

3.3.3 The ECB's credit limits

The ECB has adopted the following *size measures* in its limit-setting formulas:

(i) **Counterparties**: Equity[24] (functional form: linear scheme with a kink at the level of the median capital of counterparties);
(ii) **Issuers:** Minimum size of outstanding debt of a certain type as eligibility criteria (previously, also different sizes of outstanding debt lead to different limits);
(iii) **Countries:** GDP;
(iv) **Types of debt instruments:**[25] Market capitalization.

[23] See for instance Lucas (2004).

[24] For counterparties, the absolute figure of equity could be used as an indicator of size but it is not a clear indicator of the risk profile of an institution. Considering the range of innovative capital instruments issued by banks the amount of equity reported by financial institutions alone cannot be used as a meaningful indicator of the risk assumed with a counterparty without additional analysis. The use of Tier I capital for all counterparties would be more consistent, if this were universally available. Tier I capital is normally lower than total equity since some equity instruments may not meet all requirements to qualify as Tier I resources.

[25] Instrument-type limits are applied for instance in the ECB's own funds portfolio to non-Government instruments; covered bonds; unsecured bank bonds.

Table 4.1 Rating scales, numerical equivalents of ratings and correction factors for counterparty limits

Fitch long-term investment scale	Moody's long-term investment scale	Standard & Poor's long-term investment scale	Numerical equivalent	Rating factor
AAA	Aaa	AAA	1	1.0
AA+	Aa1	AA+	2	0.9
AA	Aa2	AA	3	0.8
AA−	Aa3	AA−	4	0.7
A+	A1	A+	5	0.6
A	A2	A	6	0.5
A−	A3	A−	7	0.4
BBB+	Baa1	BBB+	8	0.3
BBB	Baa2	BBB	9	0.2
BBB−	Baa3	BBB−	10	0.1

In terms of rating agencies, the ECB has used so far in its investment operations Fitch Ratings, Moody's and Standard & Poor's and is currently considering to add DBRS. With regard to rating aggregation, the ECB has so far used in case of multiple ratings, the *second-best* rating.[26] The minimum rating for deposits is A, and for delivery-versus-payments (DvP) operations BBB. For non-Government debt instruments the minimum rating requirement is AA−. Numerical equivalents of ratings and rating factors for counterparty limit-setting are shown in Table 4.1.

This simplistic linear scheme was tested in the ECB against somewhat more theoretical alternatives. For instance limits can be set such that regardless of the credit quality, the expected loss associated with the maximum exposure should be the same (so limits being in principle inversely proportional to probabilities of default); or limits can be set such that the sum of the expected and unexpected loss would be made independent of the credit quality ('unexpected loss' being credit risk jargon for the standard deviation of the credit losses). Since the differences between these theoretical approaches and a simple linear scheme are however moderate, the simplicity of the linear approach was considered more important.

The overall counterparty limit is proportional, with a kink however, to the capital of the given counterparty and to a rating factor which evolves as described in Table 4.1.

[26] Based on the Basel Committee's proposal in the Basel Accord II to use the second best rating in case of multiple ratings.

If capital < (median capital[27]): limit = capital * 35% * rating factor

If capital > (median capital): limit = [(median capital) * 35% + (capital – (median capital)) * 17.5%] * rating factor

In fact deposits, which account for almost all of the daily counterparty exposure, cannot have maturities longer than one month.[28] Limits are also in place for groups of connected counterparties (e.g. separate legal entities belonging to the same banking group), to control for the most obvious default correlations, and for overall country exposures[29]. Following the general principles laid down by the Basel II Committee, limits applied to issuers and counterparties[30] are clearly documented and defined in relation to capital, total assets or, where adequate measures exist, overall risk level.

The current methodology has clear advantages since it is simple and easy to implement. Within this framework, the ECB follows a similar approach to those small- or middle-size asset managers which have not developed internal credit assessment systems due to lack of resources. The combination of the two basic inputs, equity and long-term rating, defines a cost-efficient credit risk model. At the same time, the system is flexible in the sense that the addition of new counterparties can be easily considered when requested by portfolio managers.

In the case of the ECB, two distinct credit risk methodologies are maintained in parallel, namely one for the ECB's foreign reserves and another for the ECB's own-funds portfolio. While they are based on the same principles and at the outset were quite similar, different eligible asset classes, different stakeholders and a significant number of exceptions means that they have become somewhat divergent. Two somewhat diverging methodologies also mean that procedures and supporting IT systems are rather complex.

Another contributor to complexity is the fact that the ECB foreign reserves are managed by twelve NCBs, with widely differing portfolio sizes. The implementation of the credit risk methodology (CRM) for the foreign

[27] The median capital is obtained by first ordering the foreign reserves counterparties according to their capital size. For an uneven number of counterparties, the median capital is the capital of the counterparty in the middle. If the total number of foreign reserves counterparties is even, then the median is the mean of the capital of the two counterparties in the middle.

[28] Other specific instrument types are also subject to eligibility constraints, in particular on the allowed maturity.

[29] Country risk encompasses the entire spectrum of risks arising from the economic, political and social environments that may have consequences for investments in that country. Country exposure may be defined as the total exposure to entities located in the country. Eligibility criteria based on external ratings are applied at the ECB.

[30] Counterparty limits are limits established for credit exposures due to foreign exchange transactions, deposits, repos and derivatives. In sum, they encompass credit risk arising from short-term placements and settlement exposures.

reserves therefore has a number of built-in features to ensure an even playing field among portfolio managers. As mentioned above overall total limits are calculated and then distributed to NCB portfolio managers according to portfolio size. An additional level of complexity is added by a feature permitting counterparty limits to be re-allocated between NCBs, according to portfolio manager's wishes.

3.4 Liquidity limits

Liquidity risk is typically the risk that a sudden increase in the need for cash cannot be met due to balance sheet duration mismatches, inability to liquidate assets without incurring large losses or a drying up of the usual sources of refinancing (credit lines, issuance of money market instruments, etc.). A simple model of liquidity risk is presented in Chapter 7 Section 3.1 (or see Freixas and Rochet 1997, 227–9).

In the case of central banks, one needs to differentiate between local and foreign currency, in so far as a central bank knows no funding risks in the currency it issues. Despite being long-term investors and not having to match fixed date liabilities to their assets, trading liquidity risk *is* an issue for central banks, although with a very particular meaning. Many central banks need to be able to eventually liquidate some of its foreign denominated assets for the purposes of defending their currency.[31] The risk may be split into two components: the first being unable to intervene in the foreign exchange market with sufficient promptness; the second is the possibility that the central bank is forced to urgently liquidate assets at a price disadvantage due to a lack of depth in the market.

Although unfunded FX operations are feasible (by means of forward sales or spot sale of assets borrowed through FX swaps or repos), their usefulness is disputed, since sooner or later they have to be reversed. To address this issue, many central banks have enshrined in their investment guidelines the maintenance of a liquid sub-portfolio made up of a few types of instruments.[32] Putnam (2004) recommends, for this purpose, to regard securities exclusively with respect to their ease of liquidation and not to assume that a good credit rating is sufficient to guarantee the latter. Stability

[31] This need is more acute in a context of currency pegs or currency board, but also exists when the currency regime is more or less a managed float.

[32] An analysis of liquidity in markets can be found in the study by the CGFS (1999).

of the shape of the return distribution in times of market distress is therefore viewed as the overriding priority, especially since most VaR models are based on normality assumptions. Return distributions with a large potential for negative skewness are to be avoided, such as complex securities with embedded options, while liquid markets such as major currency pairs and some commodities should be considered more warranted to make a part of the liquid sub-portfolio than is usually common practice.

In view of the expanding instrument universe considered by central banks, a simple but effective way to articulate an overall policy as regards portfolio liquidity for the reserves is to ensure, by means of a limit, that at least a minimum amount is invested in highly liquid assets. When assessing what is a highly liquid instrument, special care should be given to rare stress scenarios, which might be relevant in times of foreign exchange market turmoil. Liquidity in the context of the foreign currency reserves demands that funds can be mobilized quickly and without large transaction costs to fund FX interventions. FX trades settle on a T+2 basis, thus funds must be available on a T+2 basis. Since the reserve currency of intervention will have large reference sovereign debt markets, at first sight a policy of accepting government debt and very short-term deposits as highly liquid seems justified. Indeed, many assets can be considered liquid in normal circumstances, but a flight to quality is still largely a flight into Treasuries.

The size of the liquidity limit and the definition of which instruments are considered highly liquid is obviously interrelated, in particular if the size of the limit cannot be determined 'scientifically', but necessarily must be defined to a large extent based on the judgement of decision makers. This amount should also depend on the timing of intervention, i.e. given a certain intervention size one can distinguish between one that is concentrated on a single day, and one that is spread over several days. In the latter case, portfolio managers have time to sell instruments considered less liquid (e.g. agencies and instruments issued by the Bank for International Settlements (BIS)) and (some) time deposits will mature. In principle, past operations should give an indication of the amount invested in liquid instruments that would be needed to fund a future intervention, although an appropriate buffer should be added and also the continuous growth in the size of foreign exchange markets should be taken into account. An alternative possibility is to define an adjustable liquidity limit, based on a dynamic indicator of the likelihood of an intervention. However, the definition of such an indicator is both technically and politically delicate.

The ECB has opted for a simple fixed-limit approach amounting to USD 10 billion. Highly liquid investments are defined as: (i) cash on bank accounts; (ii) US treasury bills, notes and bonds held outright; (iii) collateralized and uncollateralized deposits with a remaining time to maturity equal to or less than two business days.

3.5 Maintenance of risk limits

The set of eligible issuers, counterparties and countries as well as the inputs that go into determining limits are not static. There is therefore a need for a set of processes to cater for the maintenance of the list and the limits associated with the list of eligible entities. In the case of the ECB, the ECB Risk Management Division (RMA) monitors factors affecting the creditworthiness of countries, issuers and counterparties on a daily basis. If deemed necessary, limits for countries, issuers and counterparties are adjusted in the portfolio management system as soon as the new information becomes available. This is particularly the case for changes of credit ratings. In contrast, the ECB updates size parameters only annually, as changes to those are typically limited.

The ECB RMA subscribes to credit-related information from the credit rating agencies it relies upon. On a daily basis the division receives credit rating actions and news from these agencies related to countries, issuers and counterparties. An additional source of potential early warning information, regarding the possible deteriorating credit quality of counterparties, is daily information received from the ECB's Internal Finance Division, indicating new and outstanding settlement failures. The information received is checked for credit events relevant for the ECB's foreign reserves and own funds. In the case of a relevant event, the associated limits are adjusted in the portfolio management system and any front offices affected are contacted. If a country, issuer or counterparty is no longer eligible or if its credit limits are reduced, NCB front offices must take the necessary actions to eliminate or reduce the associated exposures within a time limit defined by the ECB RMA (considering the exposure, time to maturity and other relevant parameters). In the case of the default of a counterparty, an elaborate set of pre-defined closeout procedures is initiated. In addition, to the monitoring of credit-related events, the ECB RMA summarizes news and relevant credit actions and analysis in a weekly credit newsletter, sent by e-mail to ECB staff and management.

4. Portfolio management oversight tasks

In addition to the classical risk/performance reporting and limit compliance monitoring that a central bank risk management unit typically takes care of, this section also addresses some of the more mundane tasks that do not normally draw much attention as long as they work well.

4.1 Limit compliance monitoring

The purpose of the monitoring of limits compliance is to ensure that limits are complied with at all times and in the case of non-compliance, that the appropriate procedures are applied in order to bring exposures back within limits, as soon as possible. The main challenge when implementing risk limits is to ensure that exposures are recalculated with sufficient frequency and portfolio managers can check, before a trade is committed, what its impact on each exposure is. The monitoring of limits compliance thus heavily relies on IT systems and it is especially critical to adequately integrate the portfolio management system used by the front office with the risk systems in place.

In the case of the ECB, the portfolio management system, Wallstreet Suite (formerly FinanceKit), is configured with all market (including the liquidity limit) and credit risk limits and permits portfolio managers to check limit compliance prior to the completion of transactions. The system checks compliance with limits as transactions are entered and at regular intervals during the business days. The latter is necessary as limit utilization not only changes when positions change, but also as a consequence of changing markets, for example due to changes in market prices or exchange rates. When exposures change, the system generates associated log entries, containing limit and exposure information. The ECB RMA uses this log to check limit compliance.

During the morning of each business day ECB RMA checks that all limit exposures at the end of the previous business day (defined as 19:00) do not exceed the associated limits. If one or more limits are exceeded, the reasons for the breaches are determined. Two types of breaches are defined: technical and real breaches. Technical breaches typically relate to events outside the control of the portfolio manager (for example a change of a foreign exchange rate or a technical problem in the portfolio management system) and do generally not necessitate any follow-up. Real breaches, typically

caused by an action of the portfolio manager, require the head of the relevant front office[33] to submit a formal report to Risk Management. The report details the circumstances that led to the breach and actions taken by the front office to restore the limit.[34] Upon receipt of the report RMA validates the explanation against its findings. If deemed satisfactory RMA formally confirms the findings and the proposed course of action to the involved front office. Otherwise the front office is contacted via phone in order to arrive at a mutually agreeable solution. In case of a disagreement regarding the most appropriate way forward, the Head of ECB RMA will contact the head of the involved front office to resolve matters. Breaches exceeding limits by more than 10 per cent that cannot be restored within three business days or breaches which exceed the limit by more than 25 per cent, irrespective of the time needed to restore the limit, are reported to the Executive Board, via an email to the board member to which RMA reports (Internal Audit and the Chairman of the ECB Investment Committee are also informed). A similar procedure is applied in case an NCB front office does not follow the agreed course of action or if a limit has not been restored within agreed timescales.

All breaches are recorded by the ECB RMA[35] and are reported in a daily risk management report sent out by email and in monthly credit limit breach reports sent to NCBs. Real breaches are reported in the annual reports and to the Compliance Coordination Group, which is group of experts drawn from a representative set of ECB business areas, which assists the Executive Board and the Governing Council in assessing the actual compliance of the Eurosystem and the economic agents residing in the euro area with the ECB statute and legal acts. The group, chaired by the ECB Legal Division, submits to the Executive Board a biannual Compliance Report on the status of compliance.

A special set of counterparty exposure limits is defined in the context of a securities-lending programme run for the ECB's own-funds portfolio. The programme is managed by an external agent of behalf of the ECB. On a daily basis the agent provides exposures data on an ftp site that RMA has access to. This data is downloaded automatically and is stored in the RMA risk data warehouse. A report is run daily and checks the compliance against

[33] It is recalled the decentralized setup for ECB foreign reserves management, means that (currently) twelve front offices are involved in the management of the reserves. The ECB's own funds are managed by the so-called 'own-funds management unit' of the ECB's Investment Division.

[34] 'Restoring the limit' in this context means bringing exposure back within limits.

[35] Including the storage of hard-copies of relevant documentation printed from relevant system.

limits (agreed with the agent) and configured in the system. In the case of a breach, the ECB's account manager with the agent is contacted and requested to provide a written explanation for the breach and to ensure that the exposure is brought back within limits as soon a possible. The reporting of breaches in the context of the automated securities lending programme follows the same procedures as for other types of limit breaches.

4.2 Valuation – validation of end of day prices

For valuation purposes it is best practice where possible to mark to market (or model, if market prices are not available) investment positions at least on a daily basis. This is important not only for return and performance calculations, but also to ensure the accuracy of risk and limit utilization figures, as well as for accounting purposes. While portfolio management systems typically are linked with market data services and hence are able to reflect market moves during the trading day in valuations and risk figures, it is sometimes the case that the sources configured do not provide accurate prices and yields or at least do not do so consistently. It is therefore common practice to define a point in time at (or near) the end of the trading day at which portfolios are evaluated by means of a special set of quality controlled prices and yields. The ECB follows this practice.

On a daily basis, the ECB RMA validates the day's closing prices for all relevant financial instruments as defined in the portfolio management system. The purpose is to ensure that the end-of-day prices recorded correspond to observed market prices at (or near) closing time (defined as 14:15 ECB time for foreign exchange rates and 17:00 ECB time for all other prices and rates). This is important as these prices/yields are used to mark-to-market the foreign reserves and own-funds portfolios and the associated benchmarks at the end of the business day. These prices are also fed into the daily accounting revaluation procedures. Hence, these prices/yields impact directly the return and performance of the portfolios as reported daily, monthly and annually and also impacts the ECB's profit and loss at year end, when potential unrealized accounting losses are transferred to the profit and loss account.

The ECB portfolio management system contains a database of daily prices for all instruments and curves defined in the system. The mechanism used for the daily freezing of prices uses the Reuters feed to the portfolio management system as a starting basis. Prices in the system are continuously updated during the day. At freezing time, a batch server program copies the prices observed and stores them as the frozen prices of the day. RMA,

simultaneously, but independently, freezes the prices for all instruments by its own means. The latter consist mainly of an in-house application that conducts the freezing, i.e. connects to wire services and saves the prices obtained from them at a time selected by the user. Immediately after freezing time, the portfolio management system prices are compared to the prices frozen by RMA. Corrections are only processed if two market sources converge and both differ from the FinanceKit price by more than a threshold amount set a priori for each asset type. In such cases, the most reliable market source is taken for correction of bid-and-ask quotes.

RMA has access to two wire services – Reuters and Bloomberg – that generally are regarded as good quality sources of market quotes, for the majority of market sectors. Their prices are filtered by proprietary rules that aim at ensuring their quality (e.g. Bloomberg Generic Prices or Reuters composite prices). However, these two wire services may not provide tradable quotes, since RMA does not have access to tradable quotes from, for instance, neither the Bloomberg trade platform nor Reuters' link to the GovPX quote database, which is the largest electronic trading system for the US bond market.[36] Despite these limitations, Bloomberg still allows to select a concrete entity as quote source, if the general Bloomberg quote is considered not reliable. To a lesser extent, Reuters may also make several data feeds which can be tapped, simply by adding a source code to the RIC, after the '=' sign. When none of these options is available, the suffix RRPS may be used. RRPS stands for Reuters Pricing Service and is the Reuters equivalent of Bloomberg Fair Value, i.e. a synthetic price calculated from the curve by Reuters for less liquid instruments.[37] The quality of sources has been tested empirically and has improved over time thanks to discussions with and suggestions received from portfolio managers.

4.3 Validation of prices transacted at

Often deals are completed on the phone or through a system that is not linked to an institution's portfolio management system. Hence, the data associated with a trade often needs to be re-keyed into the portfolio management system.[38] To ensure accuracy it is common practice, at least in

[36] Access to GovPX prices to Reuters was considered too expensive if exclusively used as price source.

[37] As illustration, the thirty-year Treasury note with ISIN code US912810FT08 (as of mid November 2006) has the Reuters code 912810FT0=RRPS and the Bloomberg code T 4.5 02/15/36 Govt.

[38] In order to be able to ensure limits compliance prior to agreeing a trade a portfolio manager would need to enter the deal (at least tentatively) in the portfolio management system prior to agreeing it with the counterpart or submitting

the central banking world, to apply the four-eyes principle to deal-entry. A 'workflow' is defined where one trader agrees the deal with a counterparty and enters it in the portfolio management system and another trader is responsible for verifying the details of the deal and then finalizing it vis-à-vis the portfolio management system. After the commitment the deal will typically 'flow' to the back-office for settlement. Even if the four-eyes principle is applied an additional control may be considered, namely the validation of the price transacted at (and entered into the portfolio management system) against the prices prevailing in the market at the time the transaction was concluded. There may be two reasons for applying such additional checks: 1) it reduces operational risk further, in the sense that it may catch a keying mistake, that went unnoticed through the first validation control and 2) it could potentially act as a guard against fraud where a trader intentionally trades at an off-market price disadvantageous to their own institution.

In the case of the ECB the four-eyes principle is applied to all deals entered and, the portfolio management system additionally performs a check of the prices/yields of all transactions entered. Moreover, instrument class specific and maturity-dependent tolerance bands, determining permissible deviations from market price at the time a transaction is entered have been configured in the ECB's portfolio management system. Whenever a transaction is committed the price/yield of the transaction is compared with that observed in the market (i.e. to the latest price/yield for the instrument as seen by the portfolio management system through its market data services interface). If the price discrepancy exceeds the pre-defined threshold, the system warns the trader and a log-entry is generated. If the price is correct the trader may choose to go ahead anyway. On each business day, RMA inspects the log-entries for the previous business day and compares (using market data services) the prices/yields transacted at, with those in the market at the time the transaction was carried out. If the price can be confirmed, a note is made against the log-entry with an explanation (stored in the risk data warehouse); otherwise the trader who carried out the transaction is contacted in order to obtain an explanation and/or a proof of a contemporaneous market quote, which then, if plausible, is noted against the trade.

If no plausible explanation can be obtained, a procedure analogous to that for limit breaches is initiated. The Head of RMA formally requests the

it to an electronic trading system. ECB rules stipulate this to be the case and a deal must be fully finalized in the portfolio management system within a maximum of 15 minutes after the deal was agreed.

head of the NCB Front Office to submit a formal report to RMA on the business day following the day that the report was requested. The report should detail why the trade was concluded at this particular price. Should the explanation given not be satisfactory, procedures analogous to those defined for limit breaches are followed (see Section 4.1). Checks are performed against mid market rates (price/yield). In case a yield is entered for an instrument whose reasonability check is price based, the traded yield is first converted into a price and then the reasonability check is applied against the price corresponding to the traded yield. Tolerance bands are symmetric and reasonability warnings may result from both trades concluded above and below market prices. The tolerance bands are calculated separately for each market (EUR, USD and JPY), for each instrument class and for each maturity bucket (where applicable). This separation of instruments reflects that price volatility depends on time to maturity.

The market rate prevailing when the deal is entered into FinanceKit is used as the benchmark rate for the comparison. Updates or modifications made to the deal later do not change this. For instruments with real-time and reliable market data in FinanceKit, an hourly data frequency is implicitly assumed as the basis for the tolerance band estimation. The hourly frequency instead of a lower frequency was chosen in order to account for the time lag that persists between the time a transaction is made and the time the transaction is input and compared against the market price. It also takes into account that independently of any time lags, transaction prices may deviate slightly from the quoted market prices. All instrument classes without a reliable intra-day data feed are checked against the frozen 17:00 CET prices of the previous day. The methodology applied to calculate the rate reasonability tolerance bands is described in Box 4.2.

Box 4.2. Calculation of rate reasonability tolerance bands at the ECB

The tolerance bands are always calculated on the basis of historical market data with a daily frequency. Hourly tolerance bands are obtained by down-scaling the daily tolerance bands by the square root of eight. It is assumed that there are 8 trading hours a day, and that normally, the period between the recording of the price from the price source and the transaction would not be longer than one hour.

The tolerance bands are re-estimated annually according to the statistical properties of each instrument class during the last year. The portfolio management system does not support defining tolerance bands for individual instruments and consequently the instruments are divided into instrument classes and the same band is applied for each instrument within

Box 4.2. (cont.)

the same class. The tolerance band for an instrument class (or maturity bucket) is calculated as the maximum of all the tolerance bands for the instruments within the class.

The tolerance band is defined in such a way that breaches should occur only in up to 1 per cent of trades, when the trades are completed according to the prevailing market yield[39] and assuming that there is a one-hour time-lag between the time the deal was completed and the time the current market yield was recorded (a lag of one day is assumed for instruments for which frozen 17:00 CET prices are used). Estimations of the actual probability of breaches suggest that in reality the occurrence of breaches is over 5 times rarer than the 99 per cent confidence level would indicate.

The tolerance bands for single instruments are calculated based on the assumption that the logarithmic yield changes between the trading yields and the recorded market yields are normally distributed. The logarithmic yield change for instrument i at time t is defined as

$$r_{i,t} = \ln\left(\frac{Y_{i,t}}{Y_{i,t-1}}\right)$$

where $Y_{i,t}$ is the observed yield of instrument *i* at time *t*. The unit of time is either an hour or a day. The volatility of logarithmic yield changes σ_i for instrument *i* is estimated by calculating the standard deviation of *daily* logarithmic yield changes during the last year. Hourly tolerance bands are obtained by dividing the daily volatility by the square root of eight, assuming that there are eight trading hours a day. To achieve a 99 per cent confidence that a given correct yield is within the tolerance band, the lower and upper bounds for the band are defined as the 0.05 and 99.5 percentiles of the distribution of $r_{i,t}$, respectively. Since the logarithmic yield change is assumed to be normally distributed, $r_{i,t} \sim N(0, \sigma_i^2)$, the tolerance band for instrument *i* can be expressed, using the percentiles of the standard normal distribution, as

$$TB = (\Phi^{-1}(0.005)\sigma_i, \Phi^{-1}(0.995)\sigma_i] \cong [-2.58\sigma_i, 2.58\sigma_i]$$

where Φ denotes the standard normal cumulative distribution function. The tolerance band for instruments within instrument class *J* and maturity bucket $m=[l_m, u_m[$ is then calculated as

$$TB_m^J = \max\{TB_i : i \in J, l_m \leq mat_i < u_m\}$$

where mat_i is the time to maturity of instrument *i*.

[39] Yield is used in the remainder of this section, even if it might actually refer to price for some instruments. Yield is the preferred quantity for the reasonability check, since the volatility of logarithmic changes is more stable for yields than for prices in the long end of the maturity spectrum. Prices are only used for comparison if no reliable yield data is available.

4.4 Dealing with backdated transactions

Market values, return, performance and risk figures obviously depend on the positions maintained in the portfolio management system. However, occasionally it turns out, for example at the time of trade confirmation (with the counterparty), that some details of a transaction were incorrect. It may also happen that a transaction was not entered into the system in a timely manner or for some reason was not finalized. In such cases it may be necessary to introduce transactions into the system or alter transactions already entered in the system retrospectively. Such changes may impact portfolio valuation, limit compliance and other important factors back in time, and therefore pose a problem for compliance monitoring and for reporting already completed. It is therefore important to have processes in place that capture such events and ensure that 1) they cannot be used to circumvent compliance controls and 2) any report impacted by such changes is assessed and, if deemed necessary, re-issued. The ECB defines backdated deals as transactions entered (or changed) in the portfolio management system one or more days after their opening date.[40] Such changes may invalidate performance and risk figures in the portfolio management system and risk data warehouse. Backdated deals could also cause limit breaches to be missed, as the portfolio management system only keeps track of and checks the present limit utilization against limits.[41]

On the morning of each business day, the ECB RMA checks all transactions with an opening date prior to the previous business day which were entered or changed on the previous business day. The changes are assessed, and if necessary system processes rerun in order to cause the update of the relevant systems.[42] Special procedures associated with end of period (month, quarter and year) reporting ensure that only very rarely is it necessary to re-issue monthly, quarterly or annual reports on risk and performance. The same applies to end-of-period accounting entries. Similar procedures apply in case of changes due to incorrect static data or systems problems, which also

[40] The opening date ordinarily being the date the transaction is entered into the system, but in the case of backdated transactions may be the date the transaction would have been entered into the system, had it not accidentally been omitted for one reason or another.

[41] A backdated transaction could mask a limit breach on day T, if an offsetting transaction was entered on T+1 and the original (otherwise limit breaching transaction) was entered only at day T+1, backdated to day T.

[42] The ECB's portfolio management system does not automatically update for example return and performance figures that change due to transactions outside a five-business-day 'window'. A similar 'window' applies to the risk data warehouse which in normal circumstances is only updated on an overnight basis. Changes that go back less than five business days do not in general necessitate any action in terms of systems maintenance.

occasionally necessitate the re-evaluation of reporting that has already taken place as well as the potential rerunning of system activities for significant periods of time.

4.5 Maintenance and regular checks of static and semi-static data

Substantial amounts of static or (semi static) data are required to support a trading operation. Of particular interest to risk management are details related to financial instruments, issuers and counterparties as well as market risk factors. Furthermore there may be a need to maintain data related to fixings (e.g. to cater for floating rate notes) or deliverable baskets for futures. All this data needs to be set up and maintained. The incomplete or erroneous set-up of financial instruments could lead to wrong valuations and risk exposure calculations and hence have serious consequences.

The distribution of these tasks across the organization varies, from institution to institution. In some, the back office would be the main 'owner' of the majority of this type of data and would be responsible for its maintenance. In other organizations the responsibility may be split between front office, risk management and back office according to the type of data and the speed with which it may need to be set up or changed and some organizations may have a unit specifically responsible for this type of data maintenance.

At the ECB, for the purpose of efficiency, a separate organizational unit (Market Operations Systems Division) is responsible for the maintenance of the static data in the portfolio management system. This is the same unit which is responsible for the running and support of the portfolio management system. To facilitate this setup, the ECB has established the notion of 'data ownership'. In other words each different type of static data is owned by one organizational unit. For example data related to the configuration of financial instruments are owned by the ECB RMA, whereas details related to settlement instructions are owned by the ECB Back Office Division. A set of procedures govern the maintenance of the data; with changes typically originating from data-owning business areas being forwarded by formal channels to the unit responsible for the actual update of the data and updates taking place by means of the four-eyes principle. For some data e.g. the maintenance of cheapest-to-deliver baskets for futures contracts,[43] processes

[43] The cheapest-to-deliver (CTD) bond determines the modified duration contribution of a bond future contract in the portfolio management system. Whenever a new bond is included in the deliverable basket, it may become the CTD. If the basket in the portfolio management is incomplete, a wrong CTD bond may be selected and as a result the contract's impact on risk figures may be wrong.

have been defined in the data updating which ensures that the data is maintained on a regular basis. The ECB RMA checks the integrity of the data at regular intervals.

4.6 Maintenance of strategic benchmarks

For financial institutions using in-house benchmarks there is a need to ensure that benchmark properties, in particular asset allocation and risk characteristics (duration and liquidity) stay relatively stable over time. For example, due to the passing of time the duration of instruments held by the benchmark shortens from month to month, and instruments may mature or roll from one maturity bucket to the next, thus impacting asset allocation. In order to permit portfolio managers (and the staff maintaining potential tactical benchmarks) to anticipate future benchmark changes, it is best practice to agree on a set of rebalancing rules. As an example, one such rule could state that 'US bills in the zero–one year maturity bucket in the USD strategic benchmark are held until maturity and a maturing bill is always replaced by the latest issued bill.'

Apart from the need to maintain risk and asset allocation properties, rebalancing transactions need to be replicable by the actual portfolios. If the portfolio is large, this may occasionally limit the investment choices of the strategic benchmark due to liquidity considerations and there may also be considerations related to the availability of issuer limits.[44] For performance measurement reasons, it is important that the benchmark trades at realistic market prices with realistic bid–ask spreads.

These represent typical areas of contention between portfolio and risk managers due to their differing objectives. For example portfolio managers may argue that they need to be able to replicate the benchmark 100 per cent to ensure that any positions that they have vis-à-vis the benchmark are directly tractable to a decision to deviate from the benchmark. This may for example lead portfolio managers to argue that the benchmark should not increase exposure to a certain issuer, if the actual portfolio is already very close to its limit due to a long position and hence cannot follow the benchmark, thus forcing the actual portfolio to take another (if similar) position in another issuer. Another typical topic for discussion is how closely the benchmark should strive to follow its asset allocation. In other

[44] Obviously the benchmark should respect the same constraints vis-à-vis issuer and other limits as the actual portfolio.

words what amounts of deviation from the asset allocation approved by the organization's decision-making bodies in the context of the setting of the benchmark, are acceptable. Following the asset allocation very strictly may lead to transactions that portfolio managers would characterize as superfluous (and costly, if the actual portfolio replicates the transactions). The ECB applies monthly rebalancing to its strategic benchmarks for both the own funds as well as for the foreign reserves according to rebalancing procedures defined by the ECB RMA. For the ECB's foreign reserves this rebalancing coincides with monthly meetings of the ECB's Investment Committee and associated changes to the tactical benchmarks. Virtual transactions in the benchmark are studied in a template spreadsheet and once deemed appropriate, simulated in the ECB's portfolio management system thereby permitting an evaluation of their market risk characteristics and the viewing of the new benchmarks by portfolio managers, prior to their implementation. In addition to the monthly rebalancing the ECB RMA re-invests on a daily basis the cash holdings of the strategic benchmarks. These holdings are generally rolled over at money market tom-next rates.

5. Reporting on risk and performance

In commercial as well as central banks, risk and performance reporting typically occurs at least at three levels: the overall organizational level, the department (or business unit) level, and the individual portfolio manager or trading desk level. Often risk management will design risk and performance reports to suit the specific needs of each organizational level.

For an efficient reporting process, it is important to take into account the audience and its specific needs. **Board members** tend to focus on the return and performance over longer periods, market risk concentrations, and possibly the results from regular stress tests. Board members typically appreciate brief and concise reports without too much detail at a relatively high frequency (e.g. daily) and in-depth reports with more detail and analysis at a low frequency (monthly or even less frequently, depending on the type of institution and its lines of business). **Business area managers** are likely to be more interested in returns, large exposures and aggregate risk positions. They have a need for daily or even more frequent reports. **Portfolio managers** are interested in detailed return and risk summaries, possibly marginal risk analysis, and individual risk positions. They have a need for high frequency reporting available ad hoc and on-line.

In addition to internal management reports, financial institutions may be subject to regulatory risk reporting. There is also a trend toward greater voluntary disclosure of risks to the general public, as exemplified by the fact that a number of financial institutions reveal VaR figures in their annual reports. Central banks are following in the footsteps of their commercial counterparts at varying pace. Some put a very high emphasis on transparency, while others weigh quite carefully which information to disclose.

5.1 Characteristics of a good reporting framework

A good performance and risk reporting framework should ensure that reports satisfy a number of properties. They should be:

(i) **Timely**. Reports must be delivered in time and reflect current positions and market conditions. What is timely depends on the purpose a report serves and its audience.

(ii) **Accurate** (to the level possible) and internally consistent. Within the time (and resource) constraints given reports should be as accurate as possible, while internal inconsistencies must be avoided. Occasionally there may be a need to sacrifice some accuracy for timeliness.

(iii) **On target** – i.e. fulfill the needs of their intended audience precisely. Reports should be created with the end-user in mind. They should as far as possible contain precisely the type of information that is needed by the recipient presented in a way that maximizes the recipient's utility.

(iv) **Concise**. The level of detail should be appropriate. Most management and staff do not have the time to dig into large amounts of detail. Reporting serves its purpose best, when it is it to the point and does not contain superfluous elements.

(v) **Objective and fair**. Numbers tend to speak for themselves. However, analysis which may pass direct or indirect judgements on the performance of business units must be objective and unbiased.

(vi) **Available on demand** (to the extent possible). With the level of sophistication of today's IT solutions, it is possible to provide many risk and performance reports on-line. These types of reports may be designed and vetted by risk management, but may be made available to other parts of the organization who may run the reports on demand. A selection of such reports can in many cases fulfill immediate and ad hoc needs of department heads and portfolio managers.

Furthermore, a good reporting framework also permits analysis of unforeseen risk and performance issues. For this purpose the availability of flexible and user-friendly systems, which allow the integration of data from many sources and allow risk management users to create ad hoc reports themselves are obviously beneficial.

Once produced, reports need to reach their audience and ideally feedback flows the other way enabling the risk management function to adapt and improve its output. Another consideration is ensuring confidentiality of report contents, if necessary. Excluding reports that are available directly to end users in portfolio management or other line-of-business systems, reports may reach their audience in the following ways:

(i) **Surface mail**. Reports may still occasionally be distributed in hard-copy using traditional internal (or external) surface mail. However, with the technological solutions available these days this is reserved for reports of particular importance, life-span or perhaps confidentiality.
(ii) **Email attachment**. The most common mechanism for distributing reports is as an attachment to an email.
(iii) **Intranet based**. With the advent of web-based solutions such as for example Microsoft SharePoint or Business Objects Web Intelligence, which make it easy for non-technical users to easily manage (parts of) a website, a third option for distributing reports is by making the reports available on an intranet or internet site. This may be combined with 'advertising emails' containing links to newly issued reports or a system permitting end users to subscribe to notifications about reports issued.

The ECB RMA presently distributes most reports using email, however the intention is to make more use of intranet-based distribution.

5.2 Making sure the necessary data is available

Significant investment in information systems infrastructure typically goes into establishing a solid risk and performance reporting framework. For measuring risk and performance, at least the following two basic types of information are necessary:

(i) **Position data**. Risk reporting systems require position information for all portfolio positions. Given the sheer number of different financial instruments and transactions, the task of gathering these positions is complex. Coupling this with the fact that positions in the typical

financial institutions may be held in a number of different systems, each with their own data models and interfaces, and perhaps located in various time zones, does not make the job any easier.

(ii) **Market data**. Market data consists of time series of market prices and rates, index levels, benchmark yield curves, spreads, etc. The data must be clean, complete and available in time. Often, such data must be collected from several sources and then processed to eliminate obvious outliers. Furthermore, to ensure a consistent picture the full dataset should ideally be captured at the same time of the day. Obtaining, cleaning and organizing appropriate market data is a challenge for most financial institutions; however, with the advent of third-party providers that (for a significant fee) take on these tasks, there is also the possibility to outsource significant parts of this work, in particular if one operates in the more liquid markets.

Given the associated problems, in practice, for global financial institutions, it may not be possible to ensure 100 per cent accurate aggregate reporting of risk (and performance) at a specific point in time. Still, with advances in technology the problem is no longer insurmountable and global banks will typically be able to compile a reasonably accurate risk overview, on a daily or even more frequent basis, thanks to large investments in IT infrastructure.[45]

Central bank risk managers often have a somewhat easier life in these respects. Typically they only need to retrieve and consolidate position data from one or at most a few systems and physical locations, and the markets they operate in are often quite liquid, hence access to data is less of a problem. Finally, due to the instruments typically traded by central banks, the scope of datasets required to value investments and calculate risk figures also tends to be more manageable.

In the case of the ECB, the RMA has established a data warehouse which integrates data from the ECB's portfolio management system as well as from the agent running the automated securities lending programme of the bank. This provides the division with all relevant data to perform its risk management tasks vis-à-vis the ECB's foreign reserves and own funds. Given the few sources of data, most risk calculations take place using the infrastructure of the portfolio management system, thus ensuring full correspondence between the data seen by portfolio managers directly in the portfolio management system with data stored in and reported from the data warehouse.

[45] See also Section 6.2.

5.3 Reporting for ECB investment operations

The ECB RMA reports regularly on risk, return and performance of both the ECB foreign reserves and own-funds portfolios as well as the associated strategic and tactical benchmarks. Reporting takes place at daily, weekly, monthly, quarterly and annual frequency. The contents and frequency of reports is determined by RMA, in consultation with the receiving business areas. Efforts are devoted to ensure that reporting fulfils the criteria set out in Section 5.1. Reports sent to NCBs are restricted to information of common interest and efforts are made to discontinue non-critical reporting that is used little or only used by a small number of recipients.

The figures reported are calculated by the ECB's portfolio management system. Return and performance figures are based on time-weighted rate of return (TWR) and mark-to-market pricing. Exogenous in- and outflows are excluded from the calculations. Performance for actual portfolios is calculated vis-à-vis the associated tactical benchmarks and performance for the tactical benchmarks is calculated relative to the respective strategic benchmarks. All reporting is trade date based, i.e. transactions affect portfolios (and all relevant risk factors) with the full market value as soon as the transaction has been completed.

Box 4.3 sets out the regular reports regarding risk and performance produced and circulated by the ECB RMA with respect to the ECB's own funds and foreign reserves. The higher frequency reports are typically quite terse and contain high-level figures and other standalone pieces of information, but little commentary. Lower frequency reports tend to provide more background information and more in-depth analysis. The following list provides a brief description of the regular reports produced by the RMA.

Box 4.3. ECB Risk Management – Regular reports

Daily report – foreign reserves (for Executive Board members, business area management and staff):

- Absolute and relative VaR for the aggregate actual and for the benchmark portfolios
- Cumulated return and performance figures from the start of the year
- Large duration positions (i.e. positions exceeding 0.095 duration year)
- Liquidity figure indicating the amount held by the ECB in investments considered particularly liquid (such as US treasuries)
- Limit breaches (technical as well as real, including risk management's assessment)

Box 4.3. (cont.)

- Market values of aggregate portfolios
- Large credit exposures (exposures exceeding EUR 200 million).

Daily report – own funds (for Executive Board members, business area management and staff):

- Absolute and relative VaR for the actual and benchmark portfolios
- Cumulated return and performance figures from the start of the year
- Exposure figures for automated securities lending programme
- Limit breaches (technical and real, including risk management's assessment)
- Market values of actual portfolio
- Large credit exposures (exposures exceeding EUR 200 million).

Weekly report – foreign reserves (for NCBs' Front offices, ECB's Investment Division):

- Absolute/relative VaR and spread duration for all NCB portfolios
- Note: this information is also available on-line to NCB front-offices, however only related to their own portfolios).

Monthly performance report – foreign reserves (for NCBs' Front offices, ECB's Investment Division):

- Monthly and year-to-date return and performance for benchmarks and actual portfolios for FX reserves
- League table of monthly and year-to-date returns
- Daily returns and modified duration positions for all portfolios
- Real limit breaches for the month.

Monthly report to the Investment Committee:

- Cumulative return and performance figures to date from the start of the year and from the date of the previous investment committee meeting, for benchmarks and aggregate USD and JPY portfolios
- Modified duration, spread duration and absolute VaR of benchmarks and aggregate portfolios
- Relative VaR of tactical benchmarks and aggregate portfolios
- Market values of aggregate portfolios
- Commentary of main market developments over period since the previous Investment Committee (ICO) meeting
- Recap of tactical benchmark positions agreed at the previous ICO
- Analysis of how tactical benchmark positions and market developments resulted in the observed performance
- Tactical benchmark positions (in terms of modified duration) broken down by instrument class and time-bucket

Box 4.3. (cont.)

- Evolution of return and performance since the last ICO
- Performance attribution (allocation/selection) for the tactical versus strategic benchmarks
- Return and performance of benchmarks, aggregate portfolios and individual NCB portfolios, since the last ICO, during the last calendar month and since the beginning of the year
- Rebalancing proposal for the strategic benchmarks for the next period
- Evaluation of investment decisions taken by the ICO three months before (hit or miss).

Quarterly report to Asset and Liability Committee:

- Evolution of market value, by currency (and gold) – in EUR and local currency
- Return and performance by currency portfolio (in local currency) since start of year
- Commentary relating evolution of market value and performance to market movements and positions taken.
- League table of best-performing NCBs in USD and JPY portfolios
- Evolution of currency distribution and factors that have impacted it
- 95 per cent monthly VaR, with breakdown of market risks by risk factor (for foreign reserves)
- Utilization of market risk limits (average and maximum utilization)
- Aggregate duration and spread-duration positions (averages and maximum).

Annual reports to decision-making bodies on Performance, Market and Credit Risk – foreign reserves:

- Main changes to the management of the reserves over the year (e.g. benchmark revisions, new eligible instruments)
- Market developments and their impact on market value of portfolios
- Return and performance of benchmarks and actual portfolios and associated analysis
- League table of best performing NCBs in USD and JPY portfolios
- Historical returns and performance of aggregate portfolios since inception
- Utilization of market risk leeway (and relative VaR) by tactical benchmarks and actual portfolios
- Evolution of market risk in terms of modified duration and VaR over the year for actual portfolios and benchmarks
- Risk-adjusted return, information ratios for actual portfolios and tactical benchmarks
- Limit breaches over the year
- Credit risk: counterparty usage, large exposures, measure of concentration (collateralized and un-collateralized deposits)
- Instrument usage in terms of average holdings
- Analysis of credit risk (using CreditVaR methodology)
- Assessment of the strategic benchmarks (comparison with risk-free investments and market indices with similar composition).

Box 4.3. (cont.)

Annual reports to decision-making bodies on Performance, Market and Credit Risk – own funds:

- Main changes to the management of the own-funds portfolio over the year (e.g. benchmark revision, credit risk methodology changes, new issuers)
- Market developments and their impact on market value of the portfolio
- Return and performance of benchmark and actual portfolio and associated analysis
- Analysis of portfolio performance (performance attribution)
- Historical return and performance of benchmark and actual portfolio
- Utilization of market risk leeway (and relative VaR) by actual portfolio
- Evolution of market risk in terms of modified duration and VaR over the year for the actual portfolio and benchmark
- Risk-adjusted return, information ratios for actual portfolios and tactical benchmarks
- Credit risk: counterparty usage, large exposures, measure of concentration (collateralized and un-collateralized deposits), average holdings by credit rating
- Securities lending volumes and counterparty usage (also automated securities lending programme)
- Analysis of credit risk (using CreditVaR methodology)
- Assessment of the strategic benchmarks (comparison with risk free investments and market indices with representative compositions).

6. IT and risk management

It is virtually impossible to carry out financial risk management without sophisticated IT systems and staff to support them. Systems need to gather the necessary data, generate risk information, perform analysis, calculate and enforce limits and so forth. The overall IT architecture should be flexible and allow for easy integration of new applications and platforms, so that risk management can easily adapt to new requirements and keep implementation times reasonable for the introduction of new types of business. At the same time it is important that sufficient controls are in place to ensure that systems are reliable and changes are well managed.

This section addresses some of the issues and considerations that typically affect the organization of IT for risk management. As in previous sections a general discussion is contrasted with concrete details of the set-up applied by the ECB's RMA.

6.1 IT architecture and standards

An IT architecture is a set of standards and guidelines, ideally derived from business principles, that an organization's staff and projects should adhere to, when making IT-related decisions. The intention is to ensure that all technology decisions are guided by the same overall principles. This in turn increases the likelihood that the paradigms and technology on which solutions are built are the same or similar across the organization and leads to maximization of benefits in terms of interoperability as well as economies of scale for example in relation to maintenance and support.

Risk management systems are part of the organization's overall IT infrastructure and as such must comply with the same standards and guidelines as other systems. However, due to the vast amount of information from various business areas and systems that the typical risk management function needs in order to operate, it has a particular interest in the IT infrastructure being homogenous across the organization.

6.2 The integrated risk management system

Ideally all risk management reporting and compliance monitoring activities can be achieved in real time from one single system, thus integrating all data sources in one place and requiring risk management staff to only fully master one system. In reality this is rarely the case, not because it is not technically feasible, but rather because costs and continuous change make it very difficult to achieve in practice.

The classical integrated risk management system architecture comprises the following four components. First, an underlying enterprise-wide **data transfer infrastructure**, by which relevant position and risk data can be obtained from other systems. This could for example be based on a number of point-to-point connections between a central risk management database and transaction-originating systems across the organization. Best practice would suggest some type of middleware-based approach, as for example based on the concept of an Enterprise Service Bus (ESB[46]). The ESB

[46] An ESB provides services for transforming and routing messages, as well as the ability to centrally administer the overall system. Whatever infrastructure is in place, it is necessary that it permits the integration of new as well as old (legacy) systems. Literature (and vendors) cite the following key benefits, when compared to more traditional system-interfacing technologies: faster and cheaper accommodation of existing systems; increased flexibility; scales from point-to-point solutions to enterprise-wide deployment; emphasizes configuration rather than integration development; incremental changes can be applied with zero down-time. However, establishing an ESB can represent

provides services for transforming and routing messages, and can be centrally administered. Whatever infrastructure is in place, it is necessary that it permits the integration of new as well as old (legacy) systems. Second, an integrated risk management system comprises a **risk data warehouse** where the relevant information for risk (and return/performance) analysis is (replicated and) stored together with the data calculated by the enterprise risk system (see below). The risk data warehouse will typically be updated on a daily basis with transaction and market data. This data is sourced via the data transfer infrastructure. Analysis data as calculated by the enterprise risk system will typically also be stored in the risk data warehouse. The third element of a risk management IT architecture is an **enterprise risk system**, which centrally carries out all relevant risk (and return/performance calculations) and stores the results in the risk data warehouse. As a special case, when the organization is in the fortunate situation of only having one integrated system, this system may be integrated with front- and back-office systems, so that portfolio management data (trading volumes and prices) all reside in the same system and are entered almost in real time. The main advantage of this, as regards risk management, is that traders can simulate the risk impact of an envisaged transaction using their main portfolio management tool. Finally, the risk management architecture comprises a **reporting infrastructure**, permitting analysis and reporting based on the data stored in the risk data warehouse. The reporting infrastructure could range from a basic SQL-based reporting tool, to an elaborate set of reporting tools and systems, permitting reporting through a number of different channels, such as direct data access, through email or through a web-based intranet solution.

Integrity and security requirements for risk management systems must fulfill the same high standards as other line-of-business systems. Hence the management and monitoring of such systems is likely best kept with an infrastructure group, for example as part of a central IT department.

In addition to the above-mentioned integrated risk management solution, risk managers will have access to the normal desktop computing services, such as email, word processors etc. found in any enterprise. This is usually supplemented by access to services of market data service providers (such as Reuters and Bloomberg) and rating agencies (such as for example Fitch, Standard & Poor's and Moody's). In addition, risk managers may use specialized statistical or development tools supplemented by off-the-shelf components and libraries for the development of financial models.

a very significant investment and it requires a mature IT governance model and enterprise-wide IT strategy to already be in place. See Chappell (2004) for further details.

What is constant for a risk management function is change. The business always wants to expand into new markets, new investment types etc. It is therefore critical that the ability is in place to permit risk management systems to respond easily and cost-efficiently to new and changing requirements.

The main systems and applications used by the ECB RMA are described in Box 4.4.

6.3 The risk management IT team

An essential part of a risk management team is a group of people who combine thorough risk management business knowledge with a good functional understanding of the functionality of the systems in place and (at least) a basic understanding of IT principles. These staff will typically have the knowledge to adapt systems, when requirements change and would take the lead when functional specifications are prepared in relation to project work. They are also essential for the support of less IT-literate staff, and may develop small tactical solutions or carry out minor projects that either could not find resources centrally or for which requirements not yet are crystallized enough to form the basis for a formal project.

Such staff are often found as an integral part of a bank's risk management unit or part of a specialized group supporting trading and risk management systems. Ideally these individuals are integrated into the bank's risk management unit and have a thorough understanding of the rationale, criticality and operational risk of each task conducted for risk management purposes. In some organizations, it is generally the view that staff with IT expertise should work in and belong to a central IT department. However, this does not take the needs of the risk management function adequately into consideration. The reality is that risk management is very IT heavy and that it would be difficult to run a successful risk management function without staff with significant IT expertise. Having IT-literate staff in risk management also facilitates the communication and understanding between risk management and central IT.

6.4 Systems support and operations

A risk management function requires IT systems support on least at three levels:

- **Infrastructure**. As any other user of corporate IT services, the risk management function has a need for support in relation to the provision

of the general infrastructure, including desktop applications as well as support in relation to the infrastructure running dedicated risk management solutions. The latter services tend to require intricate knowledge of the architecture and operations of the systems supported. This in turn necessitates the establishment and retention of a group of people who have the necessary skills, which may be expensive and tends to conflict with central IT's inclination to organize itself along lines of technology expertise (for example Windows, network services, UNIX, database support etc.). Without significant attention, the latter may result in the need to involve a significant number of people in problem detection and resolution and may make it difficult to provide efficient and timely support.

- **Applications**. The main applications used by risk management are usually large complex systems such as for example a portfolio management system with front-, middle- and back-office functionality or dedicated risk management systems. There is a need for support in relation to these systems, for example to answer questions regarding functionality or to solve configuration issues. This type of support is either best provided internally in the risk management function, if related to systems exclusively used by risk management, or provided by a team dedicated for example to the functional support of front office systems. Such teams will have the dedication and focus enabling them to build up and maintain the necessary expertise.
- **Development support**. A risk management function will typically be developing its own models, reports and even small applications. These should be developed using tools and development environments sanctioned in the organizations IT architecture. However, risk management staff may not be familiar with all aspects of the technology utilized and may need advice and concrete assistance from relevant IT experts. Only rarely will this type of support be available in a formalized way[47].

With respect to support of systems from either central IT or other business areas within the organization it is best practice to establish a service level agreement, which essentially is a contract between business areas which determines which services are being provided and by whom, and also stipulates expected and guaranteed worst-case response times.

[47] At the ECB, a small unit has been established in the IT department, which is responsible for the development (and subsequently for the support) of small IT solutions in collaboration with business areas. However, the emphasis in the ECB case is still on the development (and maintenance) of systems by IT staff (or consultants managed by IT staff).

Box 4.4. The systems used by the ECB Risk Management Division (RMA)

Wall Street Systems – Wallstreet Suite[48]

Wallstreet Suite is the integrated portfolio management system used by the ECB for the management of the ECB's own funds and foreign reserves. It is configured with instrument, counterparty, country and issuer data, settlement instructions and limit data. The system together with the static and semi-static data is maintained by a separate ECB division. Positions and risk-figures are calculated on-line, based on real-time market data. The system calculates all return and most risk figures; limits are monitored on-line. The system's risk management related semi-static data includes credit risk and market risk limits as well as tolerance bands (for the rate reasonability check). Variance–covariance figures for VaR calculation purposes as provided by RiskMetrics are downloaded automatically by the system. The system is a key system from a risk management perspective, as almost all position, return and risk information used by the ECB RMA originates from this system. Although the system also incorporates back-office functionality, only the front- and middle-office features are utilized fully in the ECB's installation.

Risk Engine

The Risk Engine is the ECB RMA risk data warehouse. The system stores position, risk, compliance and performance data related to both the ECB foreign reserves management and the own funds. It is the main system used by the ECB RMA for compliance monitoring and reporting purposes. The system, which is based on technology from Business Objects, was built in-house.[49] It draws most of its data from TremaSuite, as well as the agent organization that is running the automated securities lending programme for the ECB's own funds.

Foreign Exchange Counterparty Database (FXCD)

This system implements the ECB's credit risk methodologies for the foreign reserves and own funds. The system stores information about eligible counterparties, counterparty groups, issuers and countries. Based on credit ratings obtained from the three major rating agencies as well as on the basis of balance-sheet and GDP figures (relevant for some sovereign issuers) the system calculates limits according to the credit risk methodologies. The system also implements an algorithm permitting the reallocation of counterparty limits among NCBs and facilitates the reconciliation of counterparties, counterparty groups, issuers and countries and associated limits with TremaSuite. The bespoke system has been built by a third-party systems developer, to specifications developed by the ECB. It is based on a thin-client architecture. For reporting purposes it utilizes the same Business Objects reporting infrastructure as the Risk Engine system.

[48] The system was previously called Trema FinanceKit, but was renamed Wallstreet Suite following the 'merger' in August 2006 of Wallstreet Systems and Trema AB, after both companies had been acquired by a team of financial services executives backed by Warburg Pincus.

[49] This was more as an integration exercise that a classical systems development project. The main technology components of the system comprise Business Objects Reporting Tools, Business Objects Data Integrator and well as an Oracle relational database.

Box 4.4. (cont.)

Matlab

Matlab is a powerful general purpose high-level language and interactive calculation and development environment that enable users to perform and develop solutions to computationally intensive tasks faster than with traditional programming languages such as C and C++. The general environment may be supplemented with a range of libraries (called toolboxes in Matlab terminology) some of which address problems such as financial modelling and optimization. The ECB RMA deems Matlab to be a very productive environment for the development of financial models, and makes substantial use of it in areas such as strategic asset allocation, performance attribution and credit risk modelling.

Spreadsheets

Spreadsheets are the foundation upon which many new financial products (and the associated risk management) have been prototyped and built. However, spreadsheets are also an IT support, management and regulatory nightmare as they quickly move from being an ad hoc risk manager (or trader) tool to become a complex and business critical application that is extremely difficult for IT areas to support. The ECB RMA was prior to the introduction of its risk data warehouse rather dependent on Microsoft Excel and macros written in Excel VBA (Visual Basic for Applications) as well as an associated automation tool, which permitted the automatic execution of other systems. Most of the regular risk management processes for investment operations were automated using these tools. Data storage was based on large spreadsheets and reporting used links from source workbooks into database workbooks. Needless to say, maintaining this 'architecture' was quite a challenge. However, it emerged, as of last resort, as central resources were not available to address the requirements of RMA in the early years of the ECB and the tools used were those that central IT were willing to let business areas use. After the introduction of a risk data warehouse the usage of Excel has returned to a more acceptable level, where Excel is used for analysis purposes externally to the risk data warehouse. In addition an add-in has been constructed, which permits the automatic import of data from the risk data warehouse into Excel for further manipulation. This represents a happy compromise between flexibility and control.

RiskMetrics RiskManager

RiskManager is a system that integrates market data services and risk analytics from RiskMetrics. It supports parametric, historical and simulation-based VaR calculations, what-if scenario analysis, stress testing and has a number of interactive reporting and charting features. The system is used by RMA as a supplement to the relatively limited VaR calculations and simulation capabilities offered by WallStreet Suite. RiskManager is generally delivered to RiskMetrics clients as an ASP solution. However, for confidentiality reasons the ECB has elected to maintain a local installation which is loaded with the relevant position information from WallStreet Suite on a daily basis thorough an interface available with WallStreet Suite.

6.5 Projects

As mentioned above, one of the constant features of a risk management function is change. Hence, the involvement in projects in various roles is the order of the day for a risk management function. Projects vary in scope, size and criticality. Some may be entirely contained within and thus be under full control of the risk management function, some may involve staff from other business areas and in others risk management may only have a supporting role.

In most organizations there is excess demand for IT resources and hence processes are established that govern which initiatives get the go-ahead and which do not. Due to their importance these decisions often ultimately need the involvement of board-level decision makers. Hence, also in this context is it important that risk management has a high enough profile to ensure that its arguments are heard, so that key risk management projects get the priority and resources required.

For cross-organizational projects it is common practice to establish a steering group composed of managers from the involved organizational units and have the project team report to this group. It is also quite common to establish a central project monitoring function to which projects and steering groups report regularly. However, for small projects such structures and the associated bureaucracy is a significant overhead; hence there tends to be a threshold, based on cost or resource requirements, below which projects may be run in a leaner fashion, without following the normal bureaucracy or a leaner version thereof. With respect to the setup and running of projects the following general remarks may be made. They are based on a combination of best practice from literature and hard-earned experience from the ECB. First, before a project starts it is crucial that its **objectives are clear** and agreed among its stakeholders. Otherwise it will be difficult for the project team to focus and it will be difficult ex post to assess whether the effort was worthwhile. Early in the project it is also important that the scope of the project is clearly identified. The depth to which these elements need to be documented depends on the size and criticality of the project. Second, if at all possible one should strive to keep the **scope of projects small and the timeline short**. If necessary the overall scope should be broken down into a number of subprojects, to ensure a short development/feedback cycle. Long running projects tend to lose focus, staff turnover impacts progress and scope creep kicks in. Third, establish a **small, self-contained and focused team**. A few people with the right skills and given the right conditions can move mountains. In a small and co-located team, communication is easy and the associated

overheads and risks (of misunderstandings) are minimized. A team that is spread across locations or needs to interact with a large number of other teams or departments for support is not likely to be particularly efficient, as communication is hindered and other teams have their own priorities. Part-time team members should be kept to a minimum and time-slicing of key team members should be avoided altogether. Fourth, **expectations** should be **managed**. The project should not be over-sold and it should be ensured that management are well aware of potential risks and complications both from the outset and throughout the development and testing stages. If possible, one should under-promise and over-deliver. Finally, **management support** needs to be ensured. For large and complex projects it is crucial for success in the long run that management is convinced of the project's necessity. Projects of non-negligible size will likely run into one or more problems during their lifetime. With the support of fully committed local management such issues are much easier to overcome. For larger projects, ideally a project 'sponsor' in senior management should be identified, who could prove very useful if any roadblocks need to be moved out of the way during the course of the project.

When central IT resources are scarce there may be an argument for also trying out concepts locally in a setting where risk management staff prototype solutions themselves. This avoids the initial bureaucracy and communications hurdles and permits quick progress. Also the resource consumption can be adjusted according to the demands of other tasks which may temporarily take priority. Such an approach means that central IT staff is only involved at a point where the concept is reasonably mature and hence saves the IT resources until such time when they really are needed.

6.6 Build or buy

One is often faced with the choice between developing a solution in-house and acquiring it from a vendor instead. Even if in recent years, when the trend has moved more and more towards purchasing systems, there are no hard-and-fast rules that determine the optimal solution, and hence in most cases a careful analysis must be done, prior to making a decision. Also, all but the most simple implementations are neither a full off-the-shelf solution nor a full in-house build. Typically an in-house build will rely on a number of finished components acquired from vendors and an off-the-shelf solution will, if nothing else, typically be reliant on the development or configuration of bespoke interfaces to existing in-house systems. So it is often not a question of one or the other, but where in the spectrum between a full in-house development and a complete off-the-shelf solution a project lies. In this context, the following should be given consideration.

- **Importance of match with existing business processes**. Often a system acquired from a vendor comes with its own assumptions and paradigms. Sometimes these represent best practice and it may be worthwhile to consider whether to adjust in-house practices accordingly. However, in other cases they just represent yet another alternative and may therefore not warrant a change to in-house practices and there may be other overriding reasons for not making such changes.
- **Fit with requirements**. An in-house development can be tailored to fit the requirements of the business exactly. This is obviously a strength, if the business area knows clearly what it wants. However, at the same time it can also constitute a risk, as it may turn out that some requirements were not quite as fully thought out as originally envisaged and hence an element of trial-and-error may creep into the development process. A third-party solution on the other hand may not offer quite as good a fit; however, as it is likely not to be the first installation, and the system may have gone through a number of revisions and improvements it will probably offer a more stable and robust solution. It is also likely to cover a wider functional scope, which may not be required initially, but may make it easier to extend the solution in the future.
- **Fit with IT standards**. A bespoke system can obviously be tailored to fit exactly with the IT standards set by the organization while it may not be possible to find a third-party system which adheres to the IT standards set. However, if the latter is the case then either the system must be very specialized or perhaps the standards set may be too constraining and might need a revision. If there are no other good reasons for attempting an in-house solution then perhaps it may be a case of the tail wagging the dog, and hence one should push for a revision of IT standards.
- **Availability of in-house expertise and ability to retain it**. An in-house build to a large extent necessitates the availability of the required skills in-house. For risk systems such skills comprise detailed financial knowledge (for example related to the calculation of performance, valuation of instruments), in-depth IT systems knowledge such as expertise with relational databases, programming languages, systems architecture etc. While some recourse may be taken to consultancy, it is crucial that experienced staff with such skills are available in-house and that the organization will be available to retain these staff in the future.
- **Risks and costs**. Everything else being equal, the complexity and hence the risks for time delays, cost overrun and insufficient quality are significantly higher with an in-house build than the purchase of an off-the-shelf solution. Costs include costs of staff, data, software licensing

and consultancy. As a central bank's core competency obviously does not lie with systems development it would be surprising if it could compete with systems vendors on cost. Furthermore, system vendors will have the advantage of economies of scale, as they may sell the same or very similar solutions to many institutions.

6.7 Complete outsourcing of IT systems – application service provider solutions

Historically, financial institutions have often been developing their own trading and risk systems. Today, trading systems are traditionally bought more or less off the shelf and an increasing number of companies also outsource risk technology development to take advantage of the work that has already been put into risk analytics by third parties and to potentially benefit from the associated reduction in costs.

Application service providers (ASPs) provide computer-based services to customers over a network, typically the internet. Applications offered using the ASP model are also sometimes called on-demand software or software as a service (SaaS). Vendors provide access to a particular application or a suite of applications (such as a risk management or performance attribution solution) using a standard protocol such as HTTP, which typically enables customer organizations to deploy the application to end users by means of an internet connection and a browser.

Often ASP solutions have low up-front investment costs as they require little technology investment by the customer and often are provided as pay-as-you-go solutions. This makes such solutions particularly attractive to small- to medium-sized businesses. In addition, ASP solutions to a large extent eliminate issues related to software upgrades, by placing the onus on the ASP to maintain up-to-date services, 24×7 technical support, as well as physical and electronic security and in-built support for business continuity.

While such services are often an attractive proposition they also pose some problems, such as potential confidentiality issues, as the client must trust the ASP with data which may be highly confidential and this data also has to be transferred reliably and without loss of confidentiality across for example the internet. Other potential problems relate to the timing of upgrades, which no longer is under the control of the customer, and what might happen if the ASP provider runs into financial or other trouble. The latter could cause an abrupt discontinuation of a service which the client in the meantime may have become highly reliant upon. While such problems may also arise with traditional suppliers, providing systems are installed on-site, these problems are less imminent, as the services are fully under the control of the client organization.

5 Performance measurement

Hervé Bourquin and Roman Marton

1. Introduction

Performance analysis can be considered the final stage in the portfolio management process as it provides an overall evaluation of the success of the investment management in reaching its expected performance objective. Furthermore, it identifies the individual contributions of each of its components and underlying strategies to the overall performance result. The term 'performance analysis' covers all the techniques that are implemented to study the financial results obtained in the portfolio management process, ranging from simple performance measurement to performance attribution. This chapter deals with performance measurement, which in turn can be loosely defined as the analytical framework underlying the calculation and assessment of investment returns. Chapter 6 introduces performance attribution as the second leg of a performance analysis.

Where Markowitz (1952) is often considered to be the founder of modern portfolio theory (the analysis of rational portfolio choices based on the efficient use of risk), Dietz (1966) may be seen as the father of investment performance analysis. The theoretical foundations of performance analysis can be traced back to classic economic equilibrium and corporate finance theory. Over the years, numerous new concepts that describe the interdependencies between return (ex ante and ex post) and risk measures by the application of specific factor models have been incorporated into the evaluation of investment performance (e.g. Treynor 1965; Sharpe 1966; Jensen 1968). Most of these models can be implemented directly into the evaluation framework, whereby the choice of a method should match the investment style of the portfolio management.

A critical component of any performance analysis framework is given by the definition of a benchmark portfolio. A benchmark portfolio is a reference portfolio that the portfolio manager will try to outperform by taking

'active' positions against it. These active positions are the expression of an investment strategy of the portfolio managers, who – depending on their expectations of changes in market prices and on their risk aversion – decides to deviate from the risk factor exposures of the benchmark. In contrast, purely passive strategies mean that portfolio managers simply aim at replicating the chosen benchmark, focusing e.g. on transaction cost issues. Central banks usually run a rather passive management of their investment portfolios, although still some elements of active portfolio management are adopted and out-performance versus benchmarks is sought without exposing the portfolio to a significantly higher market risk than that of the benchmark (sometimes called 'semi-passive portfolio management'). Passive investment strategies may be viewed as being directly derived from equilibrium concepts and the Capital Asset Pricing Model (which is described in Section 3.1). The practical applications of this investment approach in the world of performance analysis are covered in Section 2 of this chapter.

While the literature on the key concepts of performance measurement is wide, such as e.g. Spaulding (1997), Wittrock (2000), or Feibel (2003), it does not concentrate on the specific requirements of a public investor, like a central bank, which typically conducts semi-passive management of fixed-income portfolios with limited spread and credit risk. The first part of this chapter presents general techniques to properly determine investment returns in practice, also with respect to leveraged instruments. The subsequent section then focuses on appropriate risk-adjusted performance measures, also extending them to asymmetric financial products which are in some central banks part of the eligible instrument set. The concluding Section 4 presents the way of performance measurement at the ECB.

2. Rules for return calculation

The 'global investment performance standards' (GIPS) are a set of recommendations and requirements used to evaluate investment management practice. It allows the comparison of investment performance internationally and provides a 'best practice' standard as regards transparency for the recipients of performance reports. The GIPS were developed by the 'Investment Performance Council' and were adopted by the 'CFA Institute Board of Governors' – the latest version is of 2005 (see Chartered Financial Analyst Institute 2006).

2.1 Basic formulae

It is a requirement of the GIPS that the calculation of the single-period return must be done using the *time-weighted rate of return* (TWRR) method in order to neutralize the effect of cash flows (see e.g. Cotterill 1996 for details on the TWWR). The term 'time-weighted rate of return' was chosen to illustrate a measure of the compound rate of growth in a portfolio. Both at instrument level and aggregate level (e.g. portfolio level) the TWRR based on discrete points in time $\Delta t = [t-1;\ t]$ is determined as follows:

$$TWRR_{disc,\Delta t} = \frac{MV_t - CF_{in,t} + CF_{out,t} - MV_{t-1}}{MV_{t-1}} = \frac{MV_t - CF_{in,t} + CF_{out,t}}{MV_{t-1}} - 1 \tag{5.1}$$

where $TWRR_{disc,\Delta t}$ is the time-weighted return in period Δt; MV_{t-1} is the market value (including accrued interest) at the end of time $t-1$; MV_t is the market value (including accrued interest) at the end of time t; $CF_{in,t}$ and $CF_{out,t}$ are the cash inflows and outflows during period Δt.

The cash flows adjustment is done both at instrument and portfolio level. At instrument level, $CF_{in,t}$ and $CF_{out,t}$ occur within the portfolio and are induced by trades (e.g. bond buy or sale) or by holdings (e.g. coupon or maturity payments). At portfolio level $CF_{in,t}$ represent flows into the portfolio and $CF_{out,t}$ are flows out of the portfolio. As this method eliminates the distorting effects created by exogenous cash flows,[1] it is used to compare the returns of investment managers.

The TWRR formula (5.1) assumes that the cash flows occur at the end of time t. Following a different approach, the cash flow occurrence can be treated as per beginning of t.[2] The alternative $TWRR^*_{disc,\Delta t}$ is then

$$TWRR^*_{disc,\Delta t} = \frac{MV_t - MV_{t-1}}{MV_{t-1} + CF_{in,t} - CF_{out,t}} \tag{5.2}$$

[1] The following example illustrates the neutralization of the cash flows by the TWRR. Assume a market value of 100 on both days t–1 and t; therefore, the return should be zero at the end of day t. If a negative cash flow of −10 occurs on day t, the market value will be 100 – 10 = 90 and the corresponding TWRR will be

$$\frac{MV_{end}}{MV_{begin}} - 1 = \frac{100 - 10 + 10}{100} - 1 = \frac{100}{100} - 1 = 0.$$

[2] In practice, the end-of-period rule is more often used than the start-of-period approach. A compromise would be weighting each cash flow at a specific point during the period Δt as the original and the modified Dietz methods do – see Dietz (1966).

If the return calculation is done based on specific (finite) points in time (e.g. on a daily basis) then the resulting returns are called discrete-time returns – as shown in equations (5.1) and (5.2). For ex post performance measurement, in practice discrete-time returns are the appropriate instrument to determine the growth of value from one relevant time unit (e.g. day $t-1$) to another (e.g. day t). A more theoretical method is the concept of continuously compounded returns (also called continuous-time returns). Discrete-time returns are converted into their continuous-time equivalents $TWRR_{cont,\Delta t}$ by applying equation (5.3):

$$TWRR_{cont,\Delta t} = \ln\left(1 + TWRR_{disc,\Delta t}\right) \tag{5.3}$$

As its name already indicates, the continuously compounded return is the product of the geometrically linked (factorized) returns of every infinitesimal small time unit within the analysis period. The main disadvantage of this approach is that underlying market values are not intuitively understandable and, hence, it is rather difficult to explain those performance results to a broad audience, since they assume that the rate of growth of a portfolio can be compounded continuously – over every infinitesimal (and therefore theoretical) time interval. Practitioners can better size the sense of a return when it captures the actual evolution of the market value and cash flows at end versus the start of the observation day. Even if logarithmic returns possess some convenient mathematical attributes (e.g. linkage over time can be processed additively in contrast to discrete returns for which the compounding must be done multiplicatively), discrete-time time-weighted returns are usually favoured by practitioners. *The returns in this chapter will thus from now on represent discrete-time time-weighted returns.*

Once one has generated the TWRR for every given instrument for each given single period, the return on a portfolio for the total period is calculated in two steps. First, for the specified discrete time unit Δt (e.g. the course of a day), the returns on the different instruments that compose a portfolio are arithmetically aggregated by the respective market value weights, i.e.

$$R_{P,\Delta t} = \sum_{i=1}^{N} \left(R_{P,i,\Delta t} \times w_{P,i,t-1}\right) \tag{5.4}$$

where $R_{P,\Delta t}$ is the portfolio return and $R_{P,i,\Delta t}$ is the return on the i-th component of portfolio P in period Δt, respectively; $w_{P,i,t-1}$ is the market value weight of the i-th component of portfolio P as of day $t-1$; and N is the number of components within portfolio P.

In a second step, the portfolio return is quantified for the whole period observed. The return $R_{P,\Delta T}$ for the entire period ΔT is obtained by geometrically linking the interim (i.e. single-period) returns (this linkage method is also a requirement from the GIPS):

$$R_{P,\Delta T} = \prod_{\forall \Delta t \in \Delta T} (1 + R_{P,\Delta t}) - 1 \tag{5.5}$$

So far, the case of the determination of return figures in a *single*-currency world has been presented. But the foreign reserves portfolios of central banks consist by definition of assets denominated in multiple currencies. The conversion of currency-local returns into *returns in base currency* is given by

$$R_{Base,\Delta t} = (1 + R_{local,\Delta t}) \times (1 + R_{xch-rate,\Delta t}) - 1 \tag{5.6}$$

where for period Δt: $R_{Base,\Delta t}$ is the total return in base currency (e.g. EUR); $R_{local,\Delta t}$ is the total return in local currency; and $R_{xch\text{-}rate,\Delta t}$ is the change of the exchange rate of the base currency versus the local currency. As it can be seen in Section 4.1 of Chapter 6 on performance attribution, this multiplicative relationship leads to intra-temporal interaction effects in additive attribution models.

2.2 Trade-date versus value-date approach

The general requirement for a sound performance measurement is that the mark-to-market values are applied, i.e. all future cash flow must be properly discounted. For example, when a bond is bought, the market value of the position related to this transaction at trade date is the net of the market value of the bond at trade date and the discounted value of the settlement amount. The general description of this approach to performance measurement is commonly referred to as the *trade-date approach*. It is often compared with the so-called *value-date approach*, where in the case of a purchase of a bond, this position would only appear in the market value of the portfolio at value (i.e. settlement) date. Accordingly, the price of the bond would not influence the market value of the portfolio in the period between trade and value date. This is not satisfactory since market movements in the position from trade date (when the actual investment decision is taken) are not taken into account as opposed to the trade-date approach and as all modern standards for performance measurement recommend.

The trade-date approach has actually three main advantages. Firstly, portfolio injections and divestments impact properly the portfolio market value at trade date. Secondly, from trade date on it calculates correctly the return on an instrument consecutive to its purchase or sale. Finally, payments (e.g. credit interest, fees, etc.) are properly reflected in the performance figures as of trading date.

2.3 Actual versus all-cash basis

In principle there are two ways for a portfolio to outperform a corresponding benchmark. First, by selecting bonds, i.e. by being over- (under-) exposed relative to the benchmark in those bonds with a better (worse) performance than the benchmark. This covers yield-curve and duration position taking as well as pure bond selection. Second, by using leveraged instruments, i.e. by using instrument with a different payoff structure than the normal spot bonds. Leveraged instruments include forwards, futures, options etc. In order to separate the two out-performance factors, the GIPS require that the performance be measured both on *actual basis* and *all-cash basis*. These concepts can be defined as follows. Actual basis measures the growth of the actual invested capital; i.e. it is a combination of both fixed-income instrument picking and used leverage. This is the conventional method of measuring (active) returns by looking at the growth in value of the funds invested. All-cash basis tries to eliminate the effect of the used leverage by restating the position into an equivalent spot position having the same market exposure. The return is then stated under the following form: $(MV_{end} - Interest_{margin}) \,/\, MV_{start}$ where $Interest_{margin}$ corresponds to the daily margin.[3] This removes the effect of the leverage on the return. The all-cash basis (active) return is consequently the (active) return measured on the restated cash equivalent positions. The comparison of the actual and all-cash basis returns allows calculating the return at the level of the leveraged instruments or leveraged instruments types (e.g. daily return for a given bond future or for all bond futures included in a portfolio).

[3] After entering a futures contract the investor will have a contract with the clearer, while the clearer will have a contract with the clearing house. The clearer requires the investor to deposit funds (known as initial margin) in a margin account. Each day the futures contract is marked-to-market and the margin account is adjusted to reflect the investor's gain or loss. This adjustment corresponds to a daily margin that is noted $Interest_{margin}$ in the formula of the text above. At the close of trading, the exchange on which the futures contract trades, establishes a settlement price. This settlement price is used to compute the gains or losses on the futures contract for that day.

When managers use leverage in their portfolio, then the GIPS require that the returns be calculated on both the actual and the all-cash basis. Since the benchmarks of most central banks portfolios are normally un-leveraged, the comparison between benchmark and all-cash basis returns shows the instrument selection ability of the fund manager, whereas the difference between the actual and the all-cash basis returns indicates how efficient the use of leverage in the fund management was, i.e. $MV_{end} / MV_{start} - (MV_{end} - Interest_{margin}) / MV_{start} = Interest_{margin} / MV_{start}$.

3. Two-dimensional analysis: risk-adjusted performance measures

3.1 Capital Asset Pricing Model as a basis

Investors have a given level of risk aversion, and the expected return on their investment depends on the level of risk they are ready to bear. Therefore, considering the return dimension in performance measurement would only be half of the truth. The insertion of the risk dimension into the performance analysis has been formalized by the *Capital Asset Pricing Model (CAPM)* and its diverse modifications and applications (Treynor 1962; Sharpe 1964; Lintner 1965; Mossin 1966). The capital market line which represents the ray connecting the profiles of the risk-free asset and the market portfolio M in a risk–return diagram is given by

$$R_P = R_F + \left(\frac{R_M - R_F}{\sigma(R_M)}\right) \cdot \sigma(R_P) \tag{5.7}$$

where R_F is the risk-free rate; R_P is the return on investment portfolio P; R_M is the return on market portfolio M; $\sigma(R_P)$ is the standard deviation of historical returns on investment portfolio P; and $\sigma(R_M)$ is the standard deviation of historical returns on market portfolio M.

This relationship implies that in equilibrium the rate of return on every asset is equal to the rate of return on the risk-free asset plus a risk premium. The premium is equal to the price of the risk multiplied by the quantity of risk, where the price of risk is the difference between the return on the market portfolio and the return on the risk-free asset. The systematic risk, i.e. the beta, is defined by

$$\beta_P = \frac{\sigma(R_P, R_M)}{\sigma^2(R_M)} \tag{5.8}$$

where β_P is the beta of investment portfolio P with respect to market portfolio M; $\sigma(R_P,R_M)$ is the covariance between the historical returns on investment portfolio P and market portfolio M; and $\sigma^2(R_M)$ is the variance of historical returns on market portfolio M.

By using the beta expression, the CAPM relationship can be written as follows:

$$R_P = R_F + \beta_P \cdot (R_M - R_F) \tag{5.9}$$

In the following paragraphs of this section two groups of selected risk-adjusted performance ratios are presented. The first one is applied to the absolute return, comprising the Sharpe and Treynor ratios, while the second group consists of extended versions, i.e. reward-to-VaR and information ratios, that are focusing on relative return (i.e. performance). All of these measures are considered to be relevant for the return and performance analysis of a central bank's portfolios. The second group identifies the risk-adjusted performance due to the portfolio management relative to a benchmark – this allows us to determine how successful the management activity (passive, semi-passive or active) has been in the performance generation process.

3.2 Total performance: Sharpe ratio

The Sharpe ratio SR_P of portfolio P (or reward-to-variability ratio as it was originally named by Sharpe) is defined (in its ex post version) as follows (see Sharpe 1966):

$$SR_P = \frac{R_P - R_F}{\sigma(R_P)} \tag{5.10}$$

Comparing with formula (5.7) reveals the intuition behind this measure: if the ratio of the excess return and the total risk of a portfolio lies above (beneath) the capital market line, it will represent a positive (negative) risk-adjusted performance versus the market portfolio. Since central banks are naturally risk averse and manage their portfolios in a conservative manner by taking limited active leeway against the benchmark, the core of the return is generated by the benchmark, while the performance of the managed portfolio against its benchmark represents a small fraction of the overall return. An appropriate performance/risk ratio could therefore provide information regarding the 'efficiency' of the reference benchmark. The major problem by using the Sharpe ratio as a performance measure in

a central bank is that its interpretation can be difficult. To use it for an assessment, one would ideally have to compare it with the Sharpe ratio of a market index having exactly the same characteristics in terms of credit risk, market risk and liquidity profile as the benchmark. However, by doing so, one would obtain a reference Sharpe ratio that should be almost identical to the one of the measured portfolio. For that reason, the ECB calculates the Sharpe ratio of its benchmarks, but considers this ratio more as an indicative ex post measure rather than an efficient performance indicator.

There is an alternative pragmatic approximation which allows making use of the Sharpe ratio under these circumstances. Assuming that the market index, which should be used for the comparison, took the same total risk as the benchmark, the implied return for the market index could be calculated via the capital market line of the market index. When comparing the returns on the portfolio and the hypothetical market index (with the same risk exposure as the benchmark) the risk-adjusted out-/underperformance, the portfolio alpha α_P, can be determined:

$$\alpha_P = R_P - (R_F + SR_{Index} \cdot \sigma(R_B)) = SR_P \cdot \sigma(R_P) - SR_{Index} \cdot \sigma(R_B) \quad (5.11)$$

where SR_{Index} is the Sharpe ratio of the market index (any representative market index in terms of asset classes and weightings) and $\sigma(R_B)$ is the standard deviation of historical returns on benchmark B.

To be able to rank different portfolios with different risk levels (i.e. to compare the risk-adjusted out- or underperformances), it is in addition necessary to normalize the alphas, i.e. to set them to the same total risk unit level by dividing by the corresponding portfolio standard deviation. The resulting performance measure is called the normalized portfolio alpha $\alpha_{norm,P}$ (see Akeda 2003):[4]

$$\alpha_{norm,P} = \frac{\alpha_P}{\sigma_P} \quad (5.12)$$

3.3 Passive performance: Treynor ratio

The ex post Treynor ratio TR_P of portfolio P is given by (see Treynor 1965)

$$TR_P = \frac{R_P - R_F}{\beta_P} \quad (5.13)$$

[4] See also Treynor and Black (1973) for adjusting the Jensen alpha by the beta factor.

This performance indicator measures the relationship of the portfolio return in excess to the risk-free rate and the systematic risk – the beta – of the portfolio. The ratio can be directly derived from the CAPM (with the benchmark portfolio B replacing the market portfolio M):

$$\frac{R_P - R_F}{\beta_P} = R_B - R_F \tag{5.14}$$

The left-hand-side term is the Treynor ratio of portfolio P and the expression on the right-hand side can be seen as the Treynor ratio for the benchmark B, because the beta against the benchmark itself is one. The Treynor ratio is a ranking measure (in analogy to the Sharpe ratio). Therefore, for a similar level of risk (e.g. if two portfolios replicate exactly the benchmark and thus are managed passively against that benchmark) the portfolio that has the higher Treynor ratio is also the one that generates the highest return of the two.

For the purpose of measuring and ranking risk-adjusted performances of well-diversified portfolios, the Treynor ratio would be a better measure than the Sharpe ratio, because it only takes into account the systematic risk which cannot be eliminated by diversification. For its calculation, a reference benchmark must be chosen upon which the beta factor can be determined. In case of skewed return distributions (that are mainly the case for low modified duration portfolios) a distorted beta and Treynor ratio can occur (see e.g. Bookstaber and Clarke 1984 on incorrect performance indicators based on skewed distributions). The majority of the central bank currency reserves portfolios and their representative benchmarks normally do not consist of instruments with embedded optionalities (i.e. uncertain future cash flows), and so the empirical return distributions should not deviate in a significant manner from the normal distribution in terms of skewness and curvature.

3.4 Extension to Value-at-Risk: reward-to-VaR ratio

The quantile-based VaR has evolved rapidly to one of the most popular and widespread tools in financial risk measurement (see e.g. Jorion 2006 and Holton 2003),[5] not at least because of its enshrinement in capital adequacy

[5] See also Marton, R. 1997, 'Value at Risk – Risikomanagement gemäß der Basler Eigenkapitalvereinbarung zur Einbeziehung der Marktrisiken', unpublished diploma thesis, University of Vienna.

rules. If the VaR concept is used for risk control, it could also be incorporated into the risk-adjusted performance analysis. This could be realized by applying the reward-to-VaR ratio proposed by Alexander and Baptista (2003), which is based on the Sharpe ratio (i.e. reward-to-variability ratio). The reward-to-VaR ratio measures the impact on ex post portfolio return of an increase by one percentage of the VaR of the portfolio, by moving a fraction of wealth from the risk-free security to that portfolio. The calculation process depends on the assumption whether asset returns are considered as being normally distributed or not. In the first case the reward-to-VaR ratio RV_P of portfolio P is given by

$$RV_P = \frac{SR_P}{t^* - SR_P} \tag{5.15}$$

whereof SR_P is the Sharpe ratio of portfolio P and

$$t^* = -\Phi^{-1}(1 - \alpha) \tag{5.16}$$

where $\Phi^{-1}(.)$ is the inverse cumulative standard normal distribution function and $(1–\alpha)$ is the VaR confidence level (e.g. $\alpha = 99$ per cent implies $t^\star \approx 2.33$). In the case of normally distributed investment returns and if $t^\star > SR_P$ is true for every portfolio then the reward-to-VaR ratio and the Sharpe ratio will yield the same risk-adjusted performance ranking.

Assume for example that the reward-to-VaR ratios of portfolios A and B are 0.40 per cent and 0.22 per cent, respectively, and that this ratio is equal to 0.34 per cent for the index that is used as proxy for the market portfolio. The investor in A would have earned on average an additional 0.40 per cent per year, bearing an additional percentage point of VaR by moving a fraction of wealth from risk-free security to A. In this example, A outperformed the market portfolio and portfolio B, and B underperformed both A and the market portfolio. The return on these portfolios was assumed to be normal – had it not been the case, formula (5.15) could not have been applied.

3.5 Active performance: information ratio

The information ratio is the most common tool to measure the success or failure of the active investment management of a portfolio versus its benchmark. The information ratio (sometimes also called appraisal ratio) is defined as the quotient of the active return (alpha) on a portfolio to its

active risk (see among others Goodwin (1998) for a substantial description of the information ratio). Sharpe (1994) presented the information ratio as a kind of generalized Sharpe ratio by replacing the risk-free return by the benchmark return.

The active return can be described as the component of the portfolio return which cannot be explained by the benchmark return, and the active risk represents the volatility of the active returns. In its ex post version, which is the suitable calculation method for the performance evaluation process, the information ratio is computed as follows:

$$IR_P = \frac{R_P - R_B}{TE_{ex-post,P}} \tag{5.17}$$

The actively taken risk is represented by the ex post tracking error $TE_{ex\text{-}post,P}$ of the portfolio P versus the benchmark B, which is defined by

$$\begin{aligned} TE_{ex-post,P} &= \sigma(R_P - R_B) \\ &= \sqrt{\frac{1}{N-1} \cdot \sum_{\forall \Delta t \in \Delta T} \left((R_{P,\Delta t} - R_{B,\Delta t}) - (\overline{R}_{P,\Delta T} - \overline{R}_{B,\Delta T}) \right)^2} \\ &= \sqrt{\frac{1}{N-1} \cdot \sum_{\forall \Delta t \in \Delta T} \left((R_{P,\Delta t} - \overline{R}_{P,\Delta T}) - (R_{B,\Delta t} - \overline{R}_{B,\Delta T}) \right)^2} \end{aligned} \tag{5.18}$$

where N is the number of observed single periods; $R_{P,\Delta t}$ is the return on portfolio P in single period Δt; $R_{B,\Delta t}$ is the return on benchmark B in single period Δt; $\overline{R}_{P,\Delta T}$ is the mean return on portfolio P in entire period ΔT; and $\overline{R}_{B,\Delta T}$ is the mean return on benchmark B in entire period ΔT.

The information ratio is a good measure of the effectiveness of active management, i.e. 'management efficiency'. The interpretation of the information ratio is generally simple: the larger its value, the higher the return that an active layer manages to achieve for a given level of risk. Therefore, the information ratio makes it possible to classify different portfolios according to their performance by scaling them to a unique base risk.[6] It should be noted that there are diverging opinions of how to interpret and rank negative results.[7]

[6] As a rule of thumb, in the context of investment fund management, information ratios above one are perceived to be excellent (see e.g. Kahn 1998).

[7] For example, assuming that two portfolios are both generating a loss equal to −20 and that the tracking error of portfolio A is 2 while that of portfolio B is 5, the comparison of portfolio B, with an information ratio of −4 and

Despite that, in combination with the Treynor ratio, the information ratio transmits a sound picture of the quality of a semi-passive investment management, as it is often practiced in public institutions like central banks.

4. Performance measurement at the ECB

As introduced in Chapter 2, the ECB's investment process for foreign reserves has a three-layer structure comprising: (a) strategic benchmarks (endorsed by the ECB's Governing Council; one for each currency); (b) tactical benchmarks (specified by an investment committee); and (c) the actual portfolio level (managed by national central banks). The strategic benchmarks serve a role as yardsticks for active management by providing a 'market allocation' against which the performance of active layers is measured (see particularly Section 3 of Chapter 2 for a general overview of the portfolio management process at the ECB).

In this three-layer management structure, the main objective of the tactical benchmark and the portfolio managers is to generate, subject to respecting their allocated risk budget, portfolio out-performance by searching for and implementing alpha strategies. It is not the purpose of the strategic benchmark allocation to generate out-performance relative to the market but rather to serve as an internal market portfolio for the active layers in the investment process and thereby as a neutral reference point.

Every month, the performance of the tactical benchmark versus the strategic one is reported and analysed in a working paper that is presented to the ECB investment committee. The level of achieved performance can impact on the medium-term investment decisions of the tactical benchmark's portfolio manager who is striving to exploit medium-term market movements not foreseen by the strategic benchmark (e.g. taking profit or stop loss by closing some profitable position, opening of new positions). At the third-level layer, the portfolio managers have as mandate to outperform the tactical benchmark by using superior short-term analysis skills and to exploit information that is not taken into account at the tactical level. On a monthly basis, the performance of the NCBs' portfolios are compiled and

portfolio A, which has an information ratio equal to −10, is not straightforward. On the one hand despite its higher risk portfolio B was able to restrict the negative return to the same level as portfolio A; so portfolio B should have acted better. On the other side in the context of a risk-averse investor (as central banks naturally are) for the same return levels the one portfolio should be preferred which has taken the lower risk; this would be portfolio A. Therefore negative information ratios should not be considered.

ranked in a table in order to present the best performers. This ranking is also provided for the entire calendar year in the risk management annual reports. Due to the longer horizon, statistical significance is higher compared to monthly figures.

The ECB fully complies with the GIPS requirements, as outlined in the first section of this chapter, in order to measure the performance of its foreign reserves and own funds portfolios: TWRR method; trade-date approach; and both actual basis and all-cash basis measurement for leveraged instruments. The market value of each instrument is quantified on a daily basis and return and performance results are computed with the same frequency. With regard to the risk-adjusted performance measures presented in Section 3, the following has been implemented by the ECB:

- **Sharpe ratio and Treynor ratio**. As mentioned previously, the ECB calculates the Sharpe ratio of its leading benchmarks, but considers this ratio more as an indicative ex post measure rather than an efficient performance indicator. The Sharpe ratio is applied to diverse ECB's portfolios (notably the benchmark and actual NCBs' portfolios), as well as to indices having similar modified durations to these portfolios. However, we consider the information ratio to be a more complete measure in order to compare portfolios. Indeed the Sharpe ratio of a portfolio and of its related benchmark provides similar information as a unique corresponding information ratio. The Treynor ratio is not implemented at the ECB for the sake of limiting the numbers of published measures, since it provides only a marginal added-value in comparison to the Sharpe ratio.
- **Reward-to-VaR ratio**. Since the ECB applies relative VaR limits to its investment portfolios, this performance indicator seems particularly appropriate to capture the VaR–return profile of the bank's portfolios (therefore the ECB is considering incorporating it into its performance reporting). Following equations (5.15) and (5.16), this measure assumes a normal distribution – hence it can be applied only to instruments for which returns are assumed to be normally distributed (which should be approximately the case for typical central bank portfolios without contingent claims). To circumvent this restriction for non-linear portfolios (e.g. when including optionalities), the calculation of the reward-to-VaR ratio could alternatively be based on the *t*-distribution, as discussed in Alexander and Baptista (2003).
- **Information ratio**. The ECB calculates and reports to the decision-making bodies on an annual basis the ex post information ratios of the

active management layers, as these are easily determined and considered as good measures of the effectiveness of active management (NCBs are also ranked according to this ratio). However, the ECB applies this ratio only to the portfolios generating a positive performance since it considers that the use of the information ratio for negative performance can easily lead to counter-intuitive conclusions. Compared with industry standards, the information ratios of the ECB's foreign reserves portfolios managed by NCBs tend to be above a degree of one. In the private asset management industry, as a rule of thumb, information ratios above one are treated as being superior. However, it should be also taken into account that the afore-mentioned industry standards have been established for significantly larger levels of relative risk.

6 Performance attribution

Roman Marton and Hervé Bourquin

1. Introduction[1]

Performance attribution analysis is a specific discipline in the investment process, with the prime objective to quantify the performance contributions which stem from the active portfolio management decisions and to assign them to exposures towards the various risk factors relative to the benchmark. Typical central bank foreign reserves portfolios are composed of more or less plain vanilla securities and also the levels of market and credit risk taken are naturally low. Despite these characteristics, performance attribution analysis in central banks is much harder than it might be expected. Generally, it can be very difficult to accurately separate the impacts of different fixed-income strategies, because interactions between each of them could exist in several ways. Especially for passively managed investment portfolios with their rather small risk factor-specific performance contributions it has proven in practice to be a difficult task of finding a balance between two of the main features of performance attribution: intuitive clarity and analytical precision.

Over the past several years, based on the collective work of experts involved in both practitioner and academic research, much progress has been made on the key ingredients of modern performance attribution analysis – yet most of the publications concentrated on models tailored to equity portfolios (see the seminal articles of Brinson and Fachler 1985, Brinson *et al.* 1986 and Brinson *et al.* 1991).[2] Unfortunately, equity-based techniques are not of practical relevance for investment portfolios of central banks and other public investors which predominantly consist of fixed-income instruments, because they are not related to the specific risk factors which fixed-income portfolios are exposed to.

[1] The authors would like to thank Stig Hesselberg for his contribution to this chapter.

[2] The underlying concepts attribute the performance at sector and portfolio level to the investment categories asset allocation, instrument selection and interaction.

As on the one hand only sparse literature is concerned with fixed-income attribution analysis in general (see e.g. Buchholz *et al.* 2004; Colin 2005) and on the other hand hardly any published material is known which particularly focuses on performance attribution for central banks (as first approaches an article by De Almeida da Silva Junior (2004) and a publication by Danmarks Nationalbank (2004, appendix D)[3] could be mentioned), this chapter takes up the challenge of concentrating on risk factor analysis of bond portfolios embedding the typical peculiarities of passively managed foreign reserves portfolios, and giving the reader an overview of how attribution modelling works in central bank practice. In addition to providing state-of-the-art concepts of fixed-income attribution techniques, also their roots, which are to be found in the early stages of modern portfolio theory, are discussed in this chapter.

To perform attribution modelling, the systematic portfolio risk, i.e. the share of risk that is not eliminated by diversification, is broken down with the help of multi-factor models, which allow the different sources of risk to be analysed and the portfolio to be oriented towards the most relevant risk factors. Multi-factor return decomposition models (whereof the earliest were already designed decades ago, e.g. Ross 1976; Fama and French 1992; 1993; 1995; 1996; Carhart 1997) can be considered to be the foundation of fixed-income attribution schemes. Therefore, Section 2 of this chapter serves as an introductory part by dealing with return decomposition concepts which are applicable to modern fixed-income attribution analysis.

As the first step in performance attribution modelling for interest rate-sensitive portfolios, all the relevant return-driving risk factors have to be detected which are addressed within the investment decision process of the financial institution in question. Section 3 illustrates the mathematical derivation of the return determinants of interest rate-dependent instruments under the aspect of a risk-averse management style; the subsequent Section 4 introduces a choice of concepts available for performance attribution modelling, incorporating the specific elements which are required to build fixed-income applications suitable for central banks and other public investors. Before presenting some conclusions, Section 5 is dedicated to the fixed-income performance attribution framework currently applied to the European Central Bank's investment portfolios serving on the one hand as a specific implementation example of analysis presented in the previous

[3] In Danmarks Nationalbank (2004) the Danish central bank proposes a hypothetical fixed-income performance attribution model applicable to central bank investment management.

sections and on the other hand as a paradigm of how to improve traditional schemes.

2. Multi-factor return decomposition models

As explained in the previous chapter, the Capital Asset Pricing Model (CAPM) acts as a conceptual instrument to derive economically intuitive and practically useful measures of the risk-adjusted success or failure of investment management (e.g. Sharpe ratio, Treynor ratio and information ratio). The CAPM assumes that portfolio returns that can be adequately summarized by expected return and volatility, and investor utility functions are increasing in expected return and decreasing in volatility. For many central banks both of these simplifying model assumptions are not unrealistic, because usually only a minority of instruments (if any at all) with non-linear payoff structures (like bond options or callable bonds) are held in the foreign currency reserves portfolios, and also because central banks are risk-averse investors (by passively or at least semi-passively managing their reserves), which would correspond to concave utility functions.

With the objective to capture financial market reality as accurately as possible, portfolio theorists have developed further models which have less restrictive assumptions than the CAPM. Multi-factor models embody a significant important advantage over the CAPM, since they decompose the asset returns and also portfolio returns according to their impacts from several determinants, and not solely according to a market index. This facilitates a more precise evaluation of the investment return with respect to the taken risk. Although these models were originally designed to estimate ex ante returns, the methodology can also be applied to ex post returns in order to divide the realized return into its specific risk factor contributions. In the following three subsections, multi-factor models are first presented from a theoretical perspective and subsequently from an empirical one.

2.1 Arbitrage Pricing Theory as a basis

The Arbitrage Pricing Theory (APT) which was developed by Ross (1976) is based on fewer restrictive assumptions than the CAPM. While the CAPM postulates market equilibrium, the APT just assumes arbitrage-free markets. The APT model also tries to determine the investment return via specific

factors, but instead of just a single determinant it uses a number of K risk factors providing for a more general approach. The APT assumes that a linear relationship between the realized returns and the K risk factors exists:

$$R_i = E(R_i) + \sum_{k=1}^{K} (b_{i,k} \cdot F_k) + \varepsilon_i \tag{6.1}$$

where R_i is the realized return on asset i; $E(R_i)$ is the expected return on asset i; $b_{i,k}$ is the sensitivity of asset i towards risk factor k (factor loading); F_k is the magnitude of the k-th risk factor with $E(F_k) = 0$; and ε_i is the residual (i.e. idiosyncratic) return on asset i with $E(\varepsilon_i) = 0$.

Taking into account arbitrage considerations the following relationship results:

$$E(R_i) - R_F = \sum_{k=1}^{K} (b_{i,k} \cdot \lambda_k) \tag{6.2}$$

where λ_k can be seen as the risk premium of the k-th risk factor in equilibrium and R_F is the deterministic return on the risk-free asset. Equation (6.2) can be transformed to

$$E(R_i) - R_F = \sum_{k=1}^{K} \left(b_{i,k} \cdot (\delta_k - R_F)\right) \tag{6.3}$$

where δ_k is the return on a portfolio with a sensitivity of one towards the k-th risk factor and a sensitivity of zero against the other risk factors. The beta factor $b_{i,k}$ can be estimated by

$$b_{i,k} = \frac{\mathrm{cov}(R_i, \delta_k)}{\mathrm{var}(\delta_k)} \tag{6.4}$$

where $\mathrm{cov}(R_i, \delta_k)$ is the covariance between R_i and δ_k.

As mentioned above, an advantage of the APT model versus the CAPM is that it offers the possibility to incorporate different risk factors to explain the investment return. The return on the market portfolio (in practice: an adequate benchmark portfolio) has no special role any more – it is just one risk factor among many (examples for APT-based performance measures can be found in Connor and Korajczyk 1986 and Lehmann and Modest 1987).

2.2 Parameterizing the model: choice of risk factors

After having defined the structure of the multi-factor model, one should determine its parameters (i.e. the portfolio-relevant risk factors). Multi-factor model theory does not specify the number or nature of the risk factors – the use of these models therefore requires a prior phase for seeking and identifying factors. Many studies have been carried out on this subject – they give guidance on the choices, but the combinations of factors that allow the returns on a group of assets to be explained are not necessarily unique. Generally, there are two main techniques to identify risk factors: the exogenous (or explicit) method and the endogenous (or implicit) method. Both categories are potential candidates for central banks in order to set up or choose a fixed-income risk factor model for the purpose of performance attribution analysis. Before being able to successfully opt for a specific multi-factor attribution model it is crucial to know the elementary mechanisms of those techniques. Two representative members of the explicit method, macroeconomic and fundamental models, as well as the principal components analysis as a representative of the implicit method are sketched in the following paragraphs.

In macroeconomic models, the risk factors which are considered to impact on the asset returns are all observable macroeconomic variables. The model specification procedure for a given market for a specific time period can be executed in two regression steps. The first step uses the APT equation (6.1) to determine the factor loadings (i.e. the factor-specific sensitivities) $b_{i,k}$ for every asset i, and the second step estimates the risk premia λ_k for every risk factor k using the regression equation as below:

$$E(R_i) = \lambda_0 + \sum_{k=1}^{K} (b_{i,k} \cdot \lambda_k) \tag{6.5}$$

where $\lambda_0 = R_F$ is the deterministic return on the risk-free asset. In the end of the model specification process the following relationship, that uses the estimation results of (6.5), for a given period exists:

$$R_i = \lambda_0 + \sum_{k=1}^{K} (b_{i,k} \cdot \lambda_k) + \sum_{k=1}^{K} (b_{i,k} \cdot F_k) + \varepsilon_i \tag{6.6}$$

Here, all the required parameters are known: the number of risk factors K, the factor loadings $b_{i,k}$ and the risk premia λ_k.

Alternatively, when using fundamental factor models the factor loadings are defined explicitly and the extents of the risk factors are estimated via regression analysis. The term 'fundamental' in this context stems from the fact that these models have originally been developed for the analysis of equity returns.

Differently, the second major category to identify the relevant risk factors a portfolio is exposed to – the 'implicit method' – makes use of a statistical technique called factor analysis. It comprises the set of statistical methods which can be applied to summarize the information of a group of variables with a reduced number of variables by minimizing the loss of information due to the simplification. When designing risk factor models this technique can be used to identify the factors required by the model (i.e. it can be used to determine the explaining risk factors and the corresponding factor loadings), but it does not give any information about the nature of the risk factors – these have to be interpreted individually and be given an economic meaning. Two variations of factor analysis prevail in practice: the maximum likelihood method and principal components analysis (PCA). In particular, PCA has been found to work well on yield curve changes, since in practice all yield curve changes can be closely approximated using linear combinations of the first three eigenvectors from a PCA.

2.3 Fitting to practice: empirical multi-factor models

Empirical models have less restrictive assumptions than APT-related approaches and do not use arbitrage theory. They do not assume that there is a causal relationship between the asset returns and risk factors in every period. But they postulate that the average investment returns (or risk premia) can be directly decomposed with the help of the risk factors. So in contrast to the APT only one regression step is required. For every asset i the following model relationship is given:

$$E(R_i) - R_F = \sum_{k=1}^{K} (b_{i,k} \cdot F_k) + \varepsilon_i \tag{6.7}$$

Specifically, passively oriented managers like central banks can use multi-risk factor models to help keep the portfolio closely aligned with the benchmark along all risk dimensions. This information is then incorporated into the performance review process, where the returns achieved by a particular strategy are weighed against the risk taken. The procedure of modelling asset

returns as applied within the framework of empirical models is found in many modern performance attribution systems dealing with multi-factor models – especially in the field of fixed-income attribution analysis.

To underline the importance of multi-factor models for risk management, in addition to performance attribution they can act as the building blocks at other stages of the investment management process, such as risk budgeting and portfolio and/or benchmark optimization (see e.g. Dynkin and Hyman 2002; 2004; 2006). A multivariate risk model could also be thought of being applied to the ideological 'sister' of performance attribution: risk attribution. Using exactly the same choice of risk factors it is possible to quantify the portions of the absolute risk (e.g. volatility or VaR) and the relative risk (e.g. tracking error or relative VaR) of a portfolio that each risk factor and sector would contribute (in an ex ante sense) and consequently risk concentrations could be easily identified (for the attribution of the forward-looking variance and tracking error, respectively, see e.g. Mina 2002; Krishnamurthi 2004; Grégoire 2006). Pairing both the absolute and relative risk contributions with the absolute and relative (active) return contributions, enables the implementation of a risk-adjusted performance attribution (see among others Kophamel 2003 and Obeid 2004).

Theoretically, a multivariate fixed-income risk factor model that was chosen to be adequate for the portfolio management process of a public investor like a central bank (in terms of the selection of the risk factors) could serve as the main investment control and evaluation module. Of course, for a central bank this theoretical possibility would not necessarily find practical application, because the strategies are (among other factors) subject to policy constraints. But even then, risk factor models can contribute significantly to the investment management process in central banks, by providing the 'quantitative control centre' of that process.

3. Fixed-income portfolios: risk factor derivation

In order to effectively employ fixed-income portfolio strategies that can control interest rate risk and enhance returns, the portfolio managers must understand the forces that drive bond markets. Focusing on central banks this means that, to be able to effectively manage and analyse the foreign currency reserves portfolios, it is of crucial importance that the portfolio managers and analysts (i.e. front office and risk management) are familiar with the specific risk factors to which the central bank portfolios are

exposed and that they understand how these factors influence the asset returns of these portfolios.

The model price (i.e. the present value) of an interest rate-sensitive instrument i, e.g. a bond, at time t, with deterministic cash flows (i.e. without embedded options or prepayment facilities), is dependent on its yield to maturity $y_{i,t}$ and on the analysis time t and is defined in discrete time $P_{i,t,disc}$ and continuous time $P_{i,t,cont}$, respectively, as follows:

$$\underbrace{P_{i,t,disc} = \sum_{\forall T-t} \frac{CF_{i,T-t,t}}{(1+y_{i,t,disc})^{T-t}}}_{discrete\ time} \approx \underbrace{\sum_{\forall T-t} CF_{i,T-t,t} \cdot e^{-(T-t)\cdot y_{i,t,cont}} = P_{i,t,cont}}_{continuous\ time} \tag{6.8}$$

where for asset i: $CF_{i,T-t,t}$ is the future cash flow at time t with time to payment $T-t$; $y_{i,t,disc}$ is the discrete-time version of the yield to maturity, whereas $y_{i,t,cont}$ is its continuously compounded equivalent.

The determinants of the local-currency buy-and-hold return $dP(t,y)/P$ (i.e. without the impact of exchange rate appreciations or depreciations and trading activities) of an interest rate-dependent instrument (and hence also portfolio) without uncertain future cash flows are analytically derived by total differentiation of the price as a function of the parameters time t and yield to maturity y, and by normalizing by the price level P (for the derivation see e.g. Kang and Chen 2002, 42). Restricted to the differential terms up to the second order, the analysis delivers:[4]

$$\begin{aligned}\frac{dP(t,y)}{P} \approx & \frac{1}{P}\frac{\partial P}{\partial t}\cdot dt + \frac{1}{P}\frac{\partial P}{\partial y}\cdot dy + \frac{1}{2}\left[\frac{1}{P}\frac{\partial^2 P}{\partial t^2}\cdot (dt)^2\right. \\ & \left. + \frac{1}{P}\frac{\partial^2 P}{\partial t \partial y}\cdot dtdy + \frac{1}{P}\frac{\partial^2 P}{\partial y^2}\cdot (dy)^2 + \frac{1}{P}\frac{\partial^2 P}{\partial y \partial t}\cdot dydt\right]\end{aligned} \tag{6.9}$$

This means that the return on a fixed-income instrument is sensitive to the linear change in time dt, the linear change of its yield dy, the quadratic change in time $(dt)^2$, the quadratic change of its yield $(dy)^2$ and also cross-products between the change in time and yield $dtdy$ and $dydt$, respectively, where higher order dependencies are ignored.

The most comprehensive way to determine the return contributions induced by every risk factor is by so-called 'pricing from first principles'. This means that the model price of the instrument is determined via the present value formula immediately after every *ceteris paribus* change of the

[4] For the differential analysis the subscripts of the parameters were omitted.

considered risk factors. By applying total return formulae with respect to the initial price of the instrument, the factor-specific contributions to the instrument return can then be quantified. The main difficulties in terms of practical application are the data requirements of the approach: first, all instrument pricing algorithms must be available for the analysis, second, the whole analysis must be processed in an option-adjusted spread (OAS) framework to be able to separately measure the impacts of the diverse risk factors (see Burns and Chu 2005 for using an OAS framework for performance attribution analysis) and third (in connection with the second point), a spot rate model would need to be implemented (e.g. following Svensson 1994) to be able to accurately derive the spot rates required for the factor-specific pricing.

Alternatively, return decomposition processing could be done by using an approximate solution.[5] This is the more pragmatic way, because it is relatively easy and quick to implement. Here it is assumed that the price level-normalized partial derivatives $\frac{1}{P}\frac{\partial^2 P}{\partial t^2}\cdot(dt)^2$, $\frac{1}{P}\frac{\partial^2 P}{\partial t \partial y}\cdot dtdy$ and $\frac{1}{P}\frac{\partial^2 P}{\partial y \partial t}\cdot dydt$ as of formula (6.9) are equivalent to zero and hence could be neglected for the purpose of performance attribution analysis. Therefore the following intuitive relationship between the instrument return and its driving risk factors remains:

$$\frac{dP(t,y)}{P} \approx \underbrace{\frac{1}{P}\frac{\partial P}{\partial t}\cdot dt}_{\text{time decay effect}} \underbrace{+\frac{1}{P}\frac{\partial P}{\partial y}\cdot dy + \frac{1}{2P}\frac{\partial^2 P}{\partial y^2}\cdot(dy)^2}_{\text{yield change effect}} \tag{6.10}$$

The identified return determinants due to the passage of time and caused by the change of the yield to maturity are separately examined in the subsequent sections, whereof the yield change effect is further decomposed into its influencing components; additionally, the accurate sensitivities against the several risk factors are quantified.

3.1 Risk factor: passage of time

The first risk factor impact described by expression (6.10) represents the return contribution solely due to the decay of time *dt*, i.e. shorter times to

[5] Although approximate (also called perturbational) pricing is not as comprehensive as pricing from first principles, it should not represent a serious problem, in view of other assumptions that are made when quantifying yield curve movements. An advantage of this method is that the computations of the return and performance effects can be processed very fast without the need of any detailed security pricing formulae.

cash flow maturities and therefore changing discount factors. It is to be taken into consideration that the yield change in period dt does not have any impact on the carry effect and therefore unchanged yield curves are postulated as a prerequisite for its calculation. The precise carry return $R_{i,carry}$ (also sometimes called time return or calendar return) can at instrument level be determined as follows:

$$R_{i,carry} = \frac{P_{i,t+dt} - P_{i,t}}{P_{i,t}} \tag{6.11}$$

whereof

$$P_{i,t} = \sum_{\forall T-t} \frac{CF_{i,T-t,t}}{(1 + y_{i,t})^{T-t}} \tag{6.12}$$

and

$$P_{i,t+dt} = \sum_{\forall T-t-dt} \frac{CF_{i,T-t-dt,t+dt}}{(1 + y_{i,t})^{T-t-dt}} \tag{6.13}$$

where for instrument i: $P_{i,t+dt}$ is the model price at time $t+dt$; $CF_{i,T-t-dt,t+dt}$ is a future cash flow with time to maturity $T–t–dt$ at time $t+dt$; $y_{i,t}$ is the discrete-time yield to maturity as of basis date t.

In an approximate (perturbational) model (which could methodically also directly be applied to any sector level or to the total portfolio level) the carry return on asset i is given by[6]

$$R_{i,carry} = y_{i,t} \cdot dt \tag{6.14}$$

The approximate method does not enable one to disentangle the ordinary income return (i.e. the return impact stemming from accrued interest and coupon payments) from the roll-down return which combined would yield the overall return attributable to the passage of time. This precise decomposition of the carry return would be feasible by pricing via the first principles method and applying total return formulae.[7]

[6] See e.g. Christensen and Sorensen 1994; Chance and Jordan 1996; Cubilié 2005, appendix C.

[7] Ideally, an intelligent combination of the imprecise approximate solution and the resources- and time-consuming approach via first principles should be found and implemented in particular for the derivation of the carry effect. The ECB performance attribution methodology was designed in a way to overcome the disadvantages of both methods (see Section 5).

To analytically derive the approximation representation of the price change as a function of one single source, the widely used Taylor series expansion technique is usually applied which delivers polynomial terms in ascending order of power and so in descending explanatory order. The absolute price change of an interest rate-sensitive instrument using the second-order Taylor expansion rule with respect to time decay *dt* is approximated by (for the Taylor series rule see among others Martellini *et al.* 2004, chapter 5; Fabozzi *et al.* 2006, chapter 11)[8]

$$dP(t) = P''(t)dt + \frac{1}{2} \cdot P''(t)(dt)^2 + o((dt)^2) \tag{6.15}$$

where $dP(t)$ is the price change solely caused by the change in time t; $P'(t)$ and $P'(t)$ are the first and second derivatives of P with respect to the change in time dt; and $o((dt)^2)$ is a term negligible compared to second order terms.

For the relative price change formula (6.15) becomes

$$\frac{dP(t)}{P} = \frac{P'(t)}{P} dt + \frac{1}{2} \cdot \frac{P''}{P}(t)(dt)^2 + o((dt)^2) \tag{6.16}$$

3.2 Risk factor: change of yield to maturity

To complement, the second return determinant of equation (6.10), the change of the yield to maturity, is analyzed. Also for reasons of consistency the Taylor expansion concept is used here – with the aim at deriving the polynomials which are the explaining parameters of the price change due to the yield change. The absolute price movement with respect to a yield change is given by

$$dP(y) = P'(y)dy + \frac{1}{2} \cdot P''(y)(dy)^2 + o((dy)^2) \tag{6.17}$$

where the analogous notation as for equation (6.15) is valid.

The relative price change is then approximated by

$$\frac{dP(y)}{P} = \frac{P'(y)}{P} dy + \frac{1}{2} \cdot \frac{P''}{P}(y)(dy)^2 + o((dy)^2) \tag{6.18}$$

[8] For the Taylor expansion analysis the subscripts of the parameters were dropped.

Plugging in the well-known interest rate sensitivity measures modified duration *ModDur* and convexity *Conv*, the yield change effect is decomposed into the components linear yield change effect and quadratic yield change effect (i.e. convexity effect). The relative price change representation is now given by

$$\frac{dP(y)}{P} \approx \underbrace{-ModDur \cdot dy}_{\substack{linear\ yield \\ change\ effect}} \underbrace{+\frac{1}{2} \cdot Conv \cdot (dy)^2}_{convexity\ effect} \tag{6.19}$$

The variation of the yield to maturity of an interest rate-sensitive instrument is mainly caused by the movement of the underlying basis yield curve, which is usually the country-specific government yield curve.[9] In case of pure government issues, the yield change is almost entirely explained by the basis curve motions.[10] But for credit risk-bearing instruments (e.g. agency bonds or BIS instruments) there is also a residual yield change implied by the change of the spread between the security's yield to maturity and the government spot curve. For practical reasons just the linear yield change effect is broken down into the parts related to the government yield and the spread. In principle, the quadratic term could also be divided into these components, but as the convexity contribution itself in a typical central bank portfolio environment is of minor dimension, the decomposition would not have any significant value added for the desired attribution analysis. The relationship in formula (6.19) can be then extended to[11]

$$\frac{dP(y)}{P} \approx \underbrace{-ModDur \cdot dr}_{\substack{government\ yield \\ change\ effect}} \underbrace{-ModDur \cdot ds}_{spread\ change\ effect} + \frac{1}{2} \cdot Conv \cdot (dy)^2 \tag{6.20}$$

[9] In a portfolio context, it is more precise to speak of the 'portfolio base currency-specific basis yield curve' instead of the 'country-specific basis yield curve', because single-currency portfolios could also contain assets issued in different countries, e.g. euro portfolios.

[10] In practice it will not be described at 100 per cent, because the government instrument in question might not be a member of the universe generating the basis yield curve; and even if the issue were part of it, eventual different pricing sources would imply different yield changes.

[11] Note that if instruments with embedded options (e.g. callable bonds) or prepayment facilities (e.g. asset-backed securities) are held within the portfolio (which can be part of the eligible instruments of central banks), the modified duration would have to be replaced by the option-adjusted duration (also called effective duration) to accurately quantify interest rate sensitivity; it is determined within an option-adjusted spread (OAS) framework.

where *dr* is the change of the designated basis government spot curve; *ds* is the narrowing or widening of the spread between the yield to maturity of an instrument and the government spot curve.

3.3 Risk factor: movement of basis government yield curve

In discrete time, for the model price (present value) $P_{i,t}$ of a credit risk-free interest rate-sensitive instrument i (e.g. a government bond) at time t the following functional equivalence must be true:

$$\begin{aligned} P_{i,t} &= f(t, y_{i,t}) = \sum_{\forall T-t} \frac{CF_{i,T-t,t}}{(1+y_{i,t})^{T-t}} \\ &\equiv f(t, r_{T-t,t}) = \sum_{\forall T-t} \frac{CF_{i,T-t,t}}{(1+r_{T-t,t})^{T-t}} \end{aligned} \tag{6.21}$$

where $r_{T-t,t}$ is the spot rate valid for time to maturity $T-t$ of a government zero-coupon curve as of valuation time t. So the present value of the credit risk-free instrument must be the same when discounting the future cash flows with a constant yield to maturity as when discounting each future cash flow with its maturity-congruent zero spot rate. It should be noted, however, that spot rates are not observable in the capital market and hence must be estimated by an appropriate model.

To be able to quantify the risk factor of the change of the basis government yield curve, various methods were developed to model the dynamics of term structures and to derive the resulting factor-specific sensitivities. In term structure models (i.e. interest rate models) the model factors are specifically defined to help explain the returns of credit risk-free bonds by variations of the moments of the term structure. As the factors explain the risk of interest rate changes, it is crucial that in every model a characteristic yield-curve movement is associated with every factor. Term structure models could be divided into four categories: equilibrium and no arbitrage models, principal components models, spot rate models and functional models.

Equilibrium and no-arbitrage models are generally based on the findings in option valuation as established by Black and Scholes (1973) and Merton (1973), respectively. But they widely evolved independently from the literature on option pricing, because specific peculiarities had to be taken into account for the interest rate-sensitive domain. Discussing these models in

detail would be out of scope of this chapter, so just a few representatives should be stated: equilibrium concepts – Vasicek (1977), Brennan and Schwartz (1979; 1982) and Cox *et al.* (1985); no-arbitrage models – Ho and Lee (1986), Black *et al.* (1990) and Hull and White (1990; 1993).

Also we do not deal with the principal components models in detail. One shortcoming of this group of models is the huge number of required parameters: e.g. for three principal components and twelve times to maturity thirty-six parameters are needed. These are the parameters which describe the characteristic yield-curve movements for each of the three risk factors at each of the twelve maturities. To mention a few, term structure models based on principal components analysis can be found in Litterman and Scheinkman (1991), Bliss (1997) and Esseghaier *et al.* (2004).

Spot rate models define the risk factors to be the changes of yields of hypothetical zero-coupon bonds for specific times to maturity, i.e. the changes of pre-defined spot rates. The number of the spot rates is a variable within the modelling process – so the portfolio analyst has a huge degree of freedom to specify the model the way that corresponds best with the individual investment management attributes (as a popular reference, a spot rate model is incorporated into the RiskMetrics methodology – see RiskMetrics Group 2006). The sensitivities to the changes of the defined spot rates are known as 'key rate durations' (see Reitano 1991; Ho 1992); instead of assuming a price sensitivity against the parallel change of the yield curve (as the modified duration does), key rate durations treat the price P as a function of N chosen spot rates – designated as the key rates $r_1, \ldots, r_N$:

$$P = P(r_1, ..., r_N) \tag{6.22}$$

Key rate durations KRD_i are partial durations which measure the sensitivity of price changes of first order to the isolated changes of i various segments on the government spot curve:

$$KRD_i = -\frac{1}{P}\frac{dP}{dr_i}\Bigg| \forall i \in [1; N] \tag{6.23}$$

Therefore every yield-curve movement can be represented as a vector of changes of defined key rates $(dr_1, \ldots, dr_N)$. The relative price change is approximated by

$$\frac{dP}{P} \approx -\sum_{i=1}^{N} (KRD_i \cdot dr_i) \tag{6.24}$$

In practice, key rate durations are numerically calculated as follows:

$$KRD_i = -\frac{1}{P}\frac{P_{i,up} - P_{i,down}}{2dr_i} \tag{6.25}$$

where $P_{i,up}$ and $P_{i,down}$ are the calculated model prices after shocking up and down the diverse key rates. The key rate convexity $KRC_{i,j}$ for the simultaneous variation of the i-th and j-th key rate is given by (see Ho *et al.* 1996)

$$KRC_{i,j} = \frac{1}{P}\frac{d^2P}{dr_i \cdot dr_j}\bigg| \forall i, j \tag{6.26}$$

Formula (6.24) must therefore be extended to

$$\frac{dP}{P} \approx -\sum_{i=1}^{N}(KRD_i \cdot dr_i) + \frac{1}{2}\sum_{i,j}(KRC_{i,j} \cdot dr_i \cdot dr_j) \tag{6.27}$$

Summing up the key rate shocks should accumulate to a parallel curve shock – this intuitively means that the exposure to a parallel change of the yield curve comprises the exposures to various units of the curve. It is not guaranteed that the sum of the key rate durations is equal to the modified duration; but for instruments without cash flow uncertainties, which are the usual case in risk-averse portfolios of public investors like central banks, the difference is naturally small. For more complex products this difference can be substantially bigger due to the non-linear cash flow structure.

A disadvantage of spot rate modelling approaches is that the characteristic yield-curve shifts are not defined to be continuous. This means that certain interpolations of the interest rate changes are necessary to enable applying the model to zero-coupon bonds with maturities that do not correspond to the model-defined maturities – the bigger the number of risk factors the more precise the model will be. A representative central bank which successfully implemented a performance attribution framework based on the key rate concept is the European Central Bank. The ECB model consists of eighteen specified key rate maturities up to thirty years – this large number along with a sophisticated cash flow mapping algorithm reduces to a minimum any inaccuracies due to interpolations. The methodology of the ECB is briefly described in Section 5 of this chapter.

Functional models assume that the changes of interest rates, in particular the government spot rates, are defined continuously. They belong to the

class of parsimonious models – this is due to the fact that these techniques model the spot curve just by its first (mostly) three principal (i.e. orthogonal) components which together explain most of the variance of the historical values of government yield curves.[12] The risk factors are represented by the variations of the principal components during the analysis period: parallel shift (level change), twist (slope change) and butterfly (curvature change). Functional models can be divided into the following two categories: polynomial models and exponential models.

In polynomial models the government spot rate $r_{T-t,t}$ for time to maturity $T - t$ as of time t is described by polynomials in ascending order of power, where the general form is given by

$$r_{T-t,t} = \xi_t + \psi_t \cdot (T-t) + \zeta_t \cdot (T-t)^2 + \cdots \tag{6.28}$$

where ξ_t, ψ_t und ζ_t are time-variable coefficients associated with the yield-curve components level, slope and curvature.

The parallel shift PS_{dt} in period dt is

$$PS_{dt} = a_{t+dt} - a_t \tag{6.29}$$

where a_t stands for the average level of the spot curve at time t.

The twist TW_{dt} in period dt is

$$TW_{dt} = (b_{t+dt} - b_t) + (c_{t+dt} - c_t) \cdot (T-t) \tag{6.30}$$

where $b_t + c_t \cdot (T-t)$ represents the best linear approximation of the curve at time t.

The butterfly BF_{dt} in period dt is

$$BF_{dt} = (d_{t+dt} - d_t) + (e_{t+dt} - e_t) \cdot (T-t) + (f_{t+dt} - f_t) \cdot (T-t)^2 \tag{6.31}$$

where $d_t + e_t \cdot (T-t) + f_t \cdot (T-t)^2$ describes the best quadratic approximation of the curve at time t.

Exponential models on the other hand do not use polynomials but exponentials to reconstruct the yield curve. As a benefit the yield curves can be captured more accurately and so the resulting residual is smaller. The approach by Nelson and Siegel (1987), used also in other chapters of this

[12] As an example, for the US Treasury yield curve and the German government yield curve the explained original variance by the first three principal components is about 95 per cent (first component: ~75 per cent, second component: ~15 per cent and third component: ~5 per cent).

book, is an exponential model which specifies a functional form of the spot curve. The original motivation for this way of modelling was to cover the entire range of observable shapes of the yield curves: a monotonous form, humps on different positions of the curve and S-formations. The Nelson–Siegel model has four parameters which are to be estimated: β_0, β_1, β_2 and τ_1. These coefficients identify three unique attributes: an asymptotic value, the general shape of the curve and a humped or U-shape which combined generate the Nelson–Siegel spot curve for a specific date.[13]

The spot rate $r_{T-t,t}$ is determined by

$$r_{T-t,t} = \beta_{0,t} + \beta_{1,t} \cdot \left[\frac{1 - e^{-(T-t)/\tau_{1,t}}}{(T-t)/\tau_{1,t}}\right] + \\ + \beta_{2,t} \cdot \left[\frac{1 - e^{-(T-t)/\tau_{1,t}}}{(T-t)/\tau_{1,t}} - e^{-(T-t)/\tau_{1,t}}\right] \tag{6.32}$$

The optimal parameter values are those for which the resulting model prices of the government securities (i.e. government bonds and eventually also bills) match best the observed market prices at the same point of time.[14] Regarding the model by Nelson and Siegel, the parameters $\beta_{0,t}$, $\beta_{1,t}$, $\beta_{2,t}$ can be interpreted as time-variable level, slope and curvature factors. Therefore the variation of the model curve can be divided into the three principal components parallel shift, twist and butterfly – every movement corresponds to the respective parameter $\beta_{0,t}$, $\beta_{1,t}$ or $\beta_{2,t}$. The parallel shift PS_{dt} in period dt is given by

$$PS_{dt} = \beta_{0,t+dt} - \beta_{0,t} \tag{6.33}$$

The twist TW_{dt} in period dt is covered by

$$TW_{dt} = \beta_{1,t+dt} \cdot \left[\frac{1 - e^{-(T-t)/\tau_{1,t+dt}}}{(T-t)/\tau_{1,t+dt}}\right]_{t+dt} - \beta_{1,t} \cdot \left[\frac{1 - e^{-(T-t)/\tau_{1,t}}}{(T-t)/\tau_{1,t}}\right]_{t} \tag{6.34}$$

[13] An extension to the Nelson and Siegel (1987) method is the model by Svensson (1994). The difference between both approaches is the functional form of the spot curve – the Svensson technique defines a second exponential expression which specifies a further hump on the curve.

[14] Actually there are two ways to define the objective function of the optimization problem: either by minimizing the price errors or by minimizing the yield errors. As government bond prices are traded in the market it makes sense to specify a loss function in terms of this variable which is directly observed in the market.

The butterfly BF_{dt} in period dt is modelled by

$$BF_{dt} = \beta_{2,t+dt} \cdot \left[\frac{1 - e^{-(T-t)/\tau_{1,t+dt}}}{(T-t)/\tau_{1,t+dt}} - e^{-(T-t)/\tau_{1,t+dt}} \right]_{t+dt} - \beta_{2,t} \cdot \left[\frac{1 - e^{-(T-t)/\tau_{1,t}}}{(T-t)/\tau_{1,t}} - e^{-(T-t)/\tau_{1,t}} \right]_{t} \tag{6.35}$$

Complementing the described partial motions of the yield curve, also the sensitivities towards them can be derived from an exponential model (see e.g. Willner 1996). Following the approach proposed by Kuberek,[15] which is a modification of the Nelson-Siegel technique, the price of a government security i in continuous time can be represented in the following functional form:

$$\begin{aligned} P_{i,t} &= f(t, r, \beta_0, \beta_1, \beta_2, \tau_1) \\ &= \sum_{\forall T-t} \left[CF_{i,T-t,t} \cdot e^{-(T-t)\cdot\left(r_{T-t,t} + \beta_{0,t} + \beta_{1,t}\cdot e^{-(T-t)/\tau_{1,t}} + \beta_{2,t}\cdot(t/\tau_{1,t})\cdot e^{1-(T-t)/\tau_{1,t}}\right)} \right] \end{aligned} \tag{6.36}$$

The model-inherent factor durations of every instrument (and hence also portfolio) can then be quantified analytically. The sensitivity to a parallel shift $Dur_{i,PS,t}$ is determined by

$$Dur_{i,PS,t} = \frac{1}{P_{i,t}} \frac{\partial P_{i,t}}{\partial \beta_{0,t}} = -\frac{1}{P_{i,t}} \sum_{\forall T-t} \left[-(T-t) \cdot CF_{i,T-t,t} \cdot e^{-(T-t)\cdot r_{T-t,t}} \right] \tag{6.37}$$

The sensitivity to a twist $Dur_{i,TW,t}$ is calculated by

$$Dur_{i,TW,t} = \frac{1}{P_{i,t}} \frac{\partial P_{i,t}}{\partial \beta_{1,t}} = -\frac{1}{P_{i,t}} \sum_{\forall T-t} \left[-(T-t) \cdot e^{-(T-t)/\tau_{1,t}} \cdot CF_{i,T-t,t} \cdot e^{-(T-t)\cdot r_{T-t,t}} \right] \tag{6.38}$$

The sensitivity to a butterfly $Dur_{i,BF,t}$ is given by

$$\begin{aligned} Dur_{i,BF,t} &= \frac{1}{P_{i,t}} \frac{\partial P_{i,t}}{\partial \beta_{2,t}} \\ &= -\frac{1}{P_{i,t}} \sum_{\forall T-t} \left[-(T-t) \cdot \left((T-t)/\tau_{1,t} \right) \cdot e^{1-(T-t)/\tau_{1,t}} \cdot CF_{i,T-t,t} \cdot e^{-(T-t)\cdot r_{T-t,t}} \right] \end{aligned} \tag{6.39}$$

[15] Kuberek, R. C. 1990, 'Common factors in bond portfolio returns', Wilshire Associates Inc. Internal Memo.

One major advantage of exponential functional models is the fact that they only require very few parameters (times to payment of the cash flows and the model beta factors) to be able to determine the corresponding spot rates. Many central banks use exponential functions to construct the government spot curve either by using the model by Nelson and Siegel (1987) or by Svensson (1994) – for an overview see the study developed by the BIS (2005).

For the purpose of fixed-income performance attribution analysis, exponential techniques are used most frequently when applying functional models, because they produce better approximations of the yield curve compared to the polynomial alternatives with comparable degree of complexity. Thus, for example, polynomial modelling using a three-term polynomial would only produce a quadratic approximation of the yield curve, and this would lead to distorted results (mainly for short and long maturities). As a reference, elaborations on the polynomial decomposition of the yield curve can be found in Colin (2005, chapter 6) and Esseghaier *et al.* (2004).

3.4 Risk factor: narrowing/widening of sector and euro country spreads

Alongside the movement of the government yield curve, the change of an instrument's yield to maturity is affected by the variation of the spread against the basis government yield curve – see formula (6.20).[16] When analysing at portfolio level, at least two fundamental types of spreads can be distinguished: sector and euro country spreads. Specifically in the case of evaluating central bank portfolios it is advisable to separate these categories and not to subsume them under the same expression, e.g. 'credit spread', because the intentions behind the different types of spread positions can vary. In euro portfolios, the German government yield curve could be chosen as the reference yield curve; the differences between the non-German government yield curves and the German government yield curve would be designated as euro country spreads. In central banks the euro country spread exposures (versus the benchmark) might not be taken as part of portfolio management decisions, e.g. by decisions of an investment committee, and hence would not be treated as credit spread positions. On the contrary, investments in non-government issues, like US agency bonds or instruments issued by the Bank for International Settlements (BIS), which

[16] Technically spoken, the spread could be interpreted as an option-adjusted spread (OAS), i.e. a constant spread to the term structure based on an OAS model.

imply sector spread positioning (against the benchmark), are mostly due to concrete strategic investment policy directives or tactical asset allocation decisions.

To precisely quantify the sensitivity to a sector or euro country spread change, the spread duration and not the modified duration should be used as a measure. It specifies the amount by which the price of a risky (in terms of deviating from the basis yield curve) interest rate-sensitive instrument i changes in per cent due to a *ceteris paribus* parallel shift $ds_{i,t}$ of 100 basis points of its spread. The numerical computation of the spread duration $Dur_{i,spr,t}$ is similar to the calculation of the option-adjusted duration – with the difference that it is the spread that is shifted instead of the spot rate:

$$Dur_{i,spr,t} = -\frac{1}{P_{i,t}} \frac{P_{i,SpreadUp,t} - P_{i,SpreadDown,t}}{2 \cdot ds_{i,t}} \tag{6.40}$$

where $P_{i,SpreadUp,t}$ and $P_{i,SpreadDown,t}$ are the present values which result after the upward and downward shocks of the spread, respectively. Consequently, equation (6.21) for the calculation of the price (present value) of an interest-sensitive instrument must be extended with respect to the influence of the spread:

$$P_{i,t} = f(t, r_{T-t,t}, s_{i,t}) = \sum_{\forall T-t} \frac{CF_{i,T-t,t}}{(1 + r_{T-t,t} + s_{i,t})^{T-t}} \tag{6.41}$$

where at time t: $r_{T-t,t}$ is the government spot rate associated with the time to cash flow payment $T - t$; $s_{i,t}$ is the spread calculated for instrument i against the government spot curve.

The total spread of an instrument's yield versus the portfolio base currency-specific basis yield curve is normally to a great extent described by its components sector and country spread – the residual term can be interpreted as the issue- or issuer-specific spread (whose change effects are mostly explicitly or implicitly attributed to the category 'selection effect' within a performance attribution model).

4. Performance attribution models

In the previous section of this chapter the basis elements required to set up a performance attribution model were derived: the general structure of a multi-factor model and the return-driving risk factors for fixed-income

portfolios typically managed in central banks. We now provide a method to attribute the value added arising from the active investment decisions in a way, so that the senior management, the portfolio management (front office) and the risk management (middle office) are aware of the sources of this active return, which should consequently improve the transparency of the investment management process. A suitable model must fulfil specific requirements:

1. The appropriate performance attribution model which is to be chosen must be in accordance with the active investment decision processes, primarily with respect to the factor-specific relative positioning versus the benchmark. This is most probably the crucial point when carrying out the model specification or verification. It is common knowledge that the local-currency buy-and-hold return (i.e. without the impact of trading activities) of an interest rate-sensitive instrument is mainly driven by the impacts of the decay of time and the change of its yield. But how to best disentangle those factors to be in conformity with the investment strategies and to be able to measure the distinct success of each of them in excess to the benchmark return?
2. The variable which is aimed at being explained as accurately as possible by the model components is the active return (i.e. the out- or underperformance) versus the benchmark based on the concept of the TWRR[17] – so the incorporation of solely market risk factors of course does not sufficiently cover the range of the performance determinants. As parts of the TWRR are caused by holding instruments in the portfolio that perform better or worse than the average market changes (based on the incorporated yield curves) would induce, also an instrument selection effect must be part of the analysis. Additionally, dealing with better or worse transaction prices than quoted on the market at portfolio valuation time has an impact on the TWRR and therefore the trading skills of the portfolio managers themselves also act as an explanatory variable of the attribution model.
3. The performance attribution reports are to be tailored for the objective classes of recipients, i.e. senior management, portfolio management, risk management, etc. The classes determine the level of detail reported; for example, whether reporting at individual security level is necessary or desired. The design of the model is dependent on the needs of the clients – the resulting decision is of significant influence on the model

[17] See Chapter 5 for the determination of the time-weighted rate of return.

building: defining and implementing an attribution model just for the total portfolio level is usually easier than following a bottom-up approach from security level upwards.

4. As described in the subsequent sections there are different ways to process attribution analysis within a single period and multiple periods. But which technique is the most suitable in the individual case? The model must in any case guarantee mathematical precision without causing technically (i.e. methodically) induced residuals. Central banks are rather passive investors versus the benchmarks and so naturally the active return which is to be explained by the attribution model is rather small. Using an imprecise model could therefore easily lead to a dominance of the residual effect which is of course a result to be avoided. Additionally, the results of the attribution analysis must be intuitive to the recipients of the reports – no 'exotic' (i.e. non-standard) calculation concepts are to be used.

Basically, two ways of breaking down the active return into its determining contributions are prevalent: arithmetically and geometrically – in the first case the decomposition is processed additively and in the second case it is done multiplicatively. In each of these cases the model must be able to quantify the performance contributions for single-currency and multi-currency portfolios (for the latter the currency effects must be supplemented). A further considerable component is the time factor as the single-period attribution effects must be correctly linked over time. As a reference, a good overview of the different types of attribution models is given by Bacon (2004).

In principle, the formal conversion of return decomposition models into performance attribution models is done as follows. The decomposition model for the return R_P of portfolio P is defined by (for the notation see the previous Section 2 on multi-factor return decomposition models)

$$R_P = \sum_{k=1}^{K} (b_{P,k} \cdot F_k) + \varepsilon_P \tag{6.42}$$

The representation of the performance attribution model, which explains the active portfolio return AR_P, (assuming an additive model) is given by

$$AR_P = R_P - R_B = \sum_{k=1}^{K} \left((b_{P,k} - b_{B,k}) \cdot F_k \right) + \varepsilon_P \tag{6.43}$$

where $b_{P,k}$ is the sensitivity of portfolio P versus the k-th risk factor, $b_{B,k}$ is the sensitivity of benchmark B versus the k-th risk factor and F_k is the magnitude of the k-th risk factor.

At this point the fundamental differences between empirical return decomposition models (as described in Section 2.3) and performance attribution models should be emphasized: in equation (6.42) the buy-and-hold return is the dependent variable and so the risk factors are solely market risk factors as well as a residual or idiosyncratic component representing the instrument selection return, whereas in equation (6.43) the dependent variable is the performance determined via the method of the time-weighted rate of return and so the market risk factors and the security selection effect are extended by a determinant which represents the intraday trading contribution.

Performance attribution models can be applied to various levels within the portfolio and the benchmark, beginning from security level, across all possible sector levels, up to total portfolio level. The transition of a sector-level model to a portfolio model (i.e. the conversion of sector-level performance contributions into portfolio-level performance contributions) is done by market value-weighting the determinants of the active return AR_P:

$$AR_P = R_P - R_B = \sum_{i=1}^{N} \sum_{k=1}^{K} \left(\left(b_{P,i,k} \cdot w_{P,i} - b_{B,i,k} \cdot w_{B,i} \right) \cdot F_k \right) + \varepsilon_i \qquad (6.44)$$

where $w_{P,i}$ and $w_{B,i}$ are the market value weights of the i-th sector within portfolio P and benchmark B, respectively; $b_{P,i,k}$ and $b_{B,i,k}$ are the sensitivities of the i-th sector of portfolio P and benchmark B, respectively, versus the k-th risk factor.

The flexible structure of the model allows the influences on aggregate (e.g. total portfolio) performance to be reported along the dimensions of risk factor categories and also sector classes. The contribution PC_k related to the k-th factor across all N sectors within the portfolio and benchmark to the active return is determined by

$$PC_k = \sum_{i=1}^{N} \left(b_{P,i,k} \cdot w_{P,i} - b_{B,i,k} \cdot w_{B,i} \right) \cdot F_k \qquad (6.45)$$

The contribution PC_i related to the i-th sector across all K risk factors of the attribution model to the performance is then given by

$$PC_i = \sum_{k=1}^{K} (b_{P,i,k} \cdot w_{P,i} - b_{B,i,k} \cdot w_{B,i}) \cdot F_k \tag{6.46}$$

4.1 Fundamental types of performance attribution models

In an arithmetical (additive) performance attribution model, in every single period the following theorem must be satisfied: the sum of the individual performance contributions at a given level must equal the active return AR^{add} at this level:

$$AR^{add} = \sum_{i=1}^{N} \sum_{k=1}^{K} PC_{i,k} = R_P - R_B \tag{6.47}$$

where $PC_{i,k}$ is the performance contribution related to the k-th risk factor and the i-th sector; R_P is the return on portfolio P; R_B is the return on benchmark B; N is the number of sectors within the portfolio and the benchmark; K is the number of risk factors within the model.

For a single-currency portfolio whose local-currency return is not converted into another currency, all K return drivers are represented by local risk factors. But in case of a portfolio comprising more than one currency, the local returns of the assets are to be transformed into the portfolio base currency in order to obtain a reasonable portfolio return measure. As central banks and other public investors are global financial players that invest the foreign reserves across diverse currencies, a single-currency attribution model is not sufficient to explain the (active) returns on aggregate portfolios in base currency.

This implies for the desired attribution model that currency effects affecting the portfolio return and performance additionally would have to join the local determinants:

$$R_{P,Base} = RC_{P,i,currency} + \sum_{i=1}^{N} \sum_{k=1}^{K} RC_{P,i,k,local} \tag{6.48}$$

where $R_{P,Base}$ is the return on portfolio P in base currency; $RC_{P,i,currency}$ is the contribution to the portfolio base currency return stemming from the variation of the exchange rate of the base currency versus the local currency of the i-th sector of portfolio P; $RC_{i,k,local}$ is the contribution of the i-th sector of portfolio P to the notional local-currency portfolio return, with respect to the local risk factor k.

The local return of a currency-homogeneous sector (where 'sector' can therefore also stand for an asset) and the appreciation or depreciation of the exchange rate of the base currency relative to the local currency of the sector on a given day are linked multiplicatively to get the return in base currency (see formula [5.6]), while, on the contrary, arithmetical attribution models decompose the (base currency) return additively (the same is therefore valid for the active return). Caused by these methodically heterogeneous procedures for return calculation and return decomposition, an interaction effect (i.e. an intra-temporal cross product) arises which should be visualized separately in the attribution reports, because it is model-induced and not intended by any active investment decisions. At portfolio level the interaction effect I_P is given by

$$I_P = \sum_{i=1}^{N} \left[(w_{P,i} - w_{B,i}) \cdot (R_{P,i,local} - R_{B,i,local}) \cdot R_{i,xch-rate} \right] \tag{6.49}$$

where $w_{P,i}$ and $w_{B,i}$ are the market value weights of the i-th sector within portfolio P and benchmark B, respectively; $R_{P,i,local}$ and $R_{B,i,local}$ are the local returns of the i-th sector within portfolio P and benchmark B, respectively; $R_{i,xch\text{-}rate}$ is the movement of the exchange rate of the base currency versus the local currency of the i-th sector.

A first simple, intuitive approach to determine the currency contribution $CY_{P,i}$ of the i-th sector to the performance of a multi-currency portfolio P relative to a benchmark B could be defined as follows:

$$CY_{P,i} = w_{P,i} \cdot (R_{P,i,base} - R_{P,i,local}) - w_{B,i} \cdot (R_{B,i,base} - R_{B,i,local}) \tag{6.50}$$

As this method does not explicitly incorporate the effect of currency hedging it is only of restricted applicability for typical central bank portfolios. In the following paragraphs two alternative theoretical attribution techniques which explicitly include the impact of hedging are sketched and subsequently a pragmatic solution is introduced.

In the Ankrim and Hensel (1994) approach the currency return is defined to consist of two components – the unpredictable 'currency surprise' and the anticipatable interest rate differential (the forward premium) between the corresponding countries, respectively. By adopting the models by Brinson and Fachler (1985) and Brinson *et al.* (1986) the performance contributions resulting from asset allocation and instrument selection decisions as well as from the interaction between those categories are derived and the contributions attributable to the currency surprise and the forward premia are added.

Alternatively, the method by Karnosky and Singer (1994) incorporates continuous-time returns[18] and treats forward premia as so called 'return premia' above the local risk-free rates. Again, the local asset allocation, instrument selection and interaction concepts are used and the currency effect (with a separate contribution originating from currency forwards) is added (see the articles by Laker 2003; 2005 as exemplary evaluations of the Karnosky–Singer model). Although both of the mentioned multi-currency models already exist for more than a decade, the portfolio managers who use it in daily practice represent a minority. This is probably due to the fact that these two approaches are too complex and academic to be applied in a practical environment.

To overcome the methodical obstacles and interpretational disadvantages which are prevalent within both of the above-described approaches, a pragmatic way to disentangle the currency hedging effect from the overall currency effect is presented. This scheme is also suitable for the attribution analysis of foreign reserves portfolios of central banks and other public wealth managers. The currency impact $CY_{P,i}$ of the i-th sector on the portfolio performance can be broken down by

$$\begin{aligned} CY_{P,i} &= \left(w_{P,i} - w_{B,i}\right) \cdot R_{i,xch-rate} \\ &= \underbrace{\left(w_{P,i,\,invested} - w_{B,i,\,invested}\right) \cdot R_{i,xch-rate}}_{invested\ currency\ \ exposure\ effect} \\ &\quad + \underbrace{\left(w_{P,i,\,hedged} - w_{B,i,\,hedged}\right) \cdot R_{i,xch-rate}}_{hedged\ currency\ exposure\ effect} \end{aligned} \tag{6.51}$$

where $w_{P,i,invested}$ and $w_{B,i,invested}$ are the invested weights, and $w_{P,i,hedged}$ and $w_{B,i,hedged}$ are the hedged weights of the i-th sector within portfolio P and benchmark B, respectively.

To be able to determine the hedged weights, each currency-dependent derivative instrument within the portfolio, e.g. a currency forward, must be split into its two currency sides – long and short. The currency-specific contribution to the hedged weights stemming from each relevant instrument is its long market value divided by the total portfolio market value and its short market value divided by the total portfolio market value, respectively. Subsequently, for every i-th sector within the portfolio, the sum of the currency-specific hedged weights contributions and the sum of

[18] Continuous-time returns were applied to enable the simple addition and subtraction of returns.

the currency-specific invested weights contributions (e.g. from bonds) must be compared with their benchmark equivalents to obtain the hedged currency exposures and invested currency exposures of the *i*-th sector needed for the attribution analysis.

Contrary to the arithmetical attribution model, where the active return is quantified as the difference between the portfolio and the benchmark return, in a geometric attribution model it is based on the factorized quotient of both returns. Burnie *et al.* (1998) and Menchero (2000b) describe such models for single-currency portfolios while Buhl *et al.* (2000) present a geometric attribution model for multi-currency portfolios. An explicit comparison of arithmetic and geometric models based on case studies can be found in Wong (2003).

In the multi-period case it is a fundamental law that the result of linking the single-period performance contributions over the entire period must be equal to the total-period active return. In additive models, however, the model-inherent problem exists that the sum of the active returns (and correspondingly the sum of the risk factor-specific performance effects) of every single period does not equal the active return of the total period.[19] Different methods of arithmetic multi-period attribution analysis attempt to solve the inter-temporal residual-generating problem in different ways (see e.g. Carino 1999; Kirievsky and Kirievsky 2000; Mirabelli 2000; Davies and Laker 2001; Campisi 2002; Menchero 2000a; 2004). An algorithm which should be explicitly mentioned is the recursive periodization method that was first published by Frongello (2002a) and then also confirmed by Bonafede *et al.* (2002). This approach completely satisfies the requirements on a linking algorithm as defined by Carino (1999, 6) and Frongello (2002a, 13) – the mathematical proof for the residual-free compounding of the single-period contributions was presented by Frongello (2002b). As a central bank reference, the suggested performance attribution model of Danmarks Nationalbank also uses this linking technique (see Danmarks Nationalbank 2004, appendix D).[20]

[19] Also replacing the summation of the single-period effects by the multiplication of the factorized effects (in analogy to the geometric compounding of discrete returns over time) would not lead to correct total-period results.

[20] For the sake of completeness we also want to point at an alternative approach of how to accurately link performance contributions over time. When carrying out the attribution analysis based on absolute (i.e. nominal) currency units instead of relative (i.e. percentage) figures, the performance effects for the total period can simply be achieved by summing up the single period contributions. This is exactly the way how the compounding over time is done within the ECB attribution framework which is described in Section 5.

4.2 Fixed-income performance attribution models

For central bank currency reserves portfolios, i.e. interest rate-sensitive portfolios, the categories of asset allocation, security selection and interaction taken from the classical equity-oriented attribution models are not adequate to mirror the investment process (an equivalent statement in general terms can be found in Spaulding 2003, 74). In the context of a fixed-income portfolio there are at least five independent ways to deviate from the benchmark in terms of risk factor exposures: duration position (e.g. by trading futures contracts), term structure / yield position (e.g. by preferring specific maturity buckets), country weighting (e.g. by overweighting high-yield government bond markets), sector weighting / credit position (e.g. by investing in market segments like agency bonds or Pfandbriefe) and instrument selection. An accurate performance attribution model for a central bank should be able to measure the distinct contributions of each type of active investment decisions that is relevant for its individual portfolio management process. To study, Colin (2005) provides an introduction into the field of fixed-income attribution analysis and also gives an overview of the relevant concepts in a compact form; see also Buchholz *et al.* (2004).

In this section we will concentrate on *additive* attribution models as they seem more adequate than their geometric alternatives for the investment analysis practice, in particular related to central banks and other public investors. This is, among others, due to the fact that the resulting performance contributions are more intuitive to understand from a methodical point of view. The basic structure of a fixed-income performance attribution model – independent of the aggregation level and the analysis period – can be represented by

$$AR_{base} = PC_{market-risk} + PC_{selection,intraday,residual} \tag{6.52}$$

where AR_{base} is the active time-weighted rate of return in base currency; $PC_{market-risk}$ is the contribution to active return related to market risk factors; $PC_{selection,intraday,residual}$ is the portion of the performance which is unexplained by the applied market risk factor model and which usually is due to instrument selection, intraday trading activities and the real model residual term.

The model can then take the form of

$$AR_{base} = PC_{carry} + PC_{govtYldChg} + PC_{spreadChg} + PC_{convexity} + PC_{currency-invested} + PC_{currency-hedged} + PC_{selection,intraday,residual} \quad (6.53)$$

where the market risk factor contributions to the performance are divided into local and currency effects and can be classified as follows: carry effect PC_{carry}; government yield change effect $PC_{govtYldChg}$; spread change effect $PC_{spreadChg}$; convexity effect $PC_{convexity}$; invested currency exposure effect $PC_{currency\text{-}invested}$; hedged currency exposure effect $PC_{currency\text{-}hedged}$.

The following passage introduces an *example* of an explicit performance attribution proposal which (among other alternatives) can be thought of being appropriate for the investment process of fixed-income central bank currency reserves portfolios. The evaluation of the government yield change effect is based on the concept of parsimonious functional models (see Section 3.3) which derive a defined number of the principal components of the entire government yield curve motion, representing the government return-driving parameters. By explicitly modelling the unique movements parallel shift, twist and butterfly, the isolated contributions of typical central bank term structure positioning strategies versus the benchmark, like flatteners, steepeners or butterfly trades, can be clearly quantified.[21] In publications on fixed-income attribution analysis, the basis government yield change effect is regularly broken down into those three partial motions (see e.g. Ramaswamy 2001; Cubilié 2005; Murira and Sierra 2006). As a reference publication by a central bank, the performance attribution proposal outlined by Danmarks Nationalbank (2004, appendix D) decomposes the government curve movement effect into the impacts originating from the parallel shift and from the variation of the curve shape (additionally it disentangles the sector-specific spread change contribution from the instrument-specific spread change effect).

Applying a perturbational technique (see Section 3) to our illustrative example, the risk factor-related representation of the performance attribution model at portfolio level is defined as follows:[22]

[21] In literature sometimes the parallel shift effect is designated as 'duration effect' and the combined twist and butterfly effect is called 'yield curve reshaping effect'.

[22] The duration against the parallel shift, twist and butterfly could either be a modified duration or an option-adjusted duration. The most appropriate measure with respect to the diverse instrument types should be used; so in case of portfolios with e.g. callable bonds the option-adjusted duration would be a more accurate measure than the modified duration.

$$AR_{P,base} = \sum_{i=1}^{N} \begin{cases} (y_{P,i} \cdot w_{P,i} - y_{B,i} \cdot w_{B,i}) \cdot dt \\ +(Dur_{P,i,PS} \cdot w_{P,i} - Dur_{B,i,PS} \cdot w_{B,i}) \cdot (-PS) \\ +(Dur_{P,i,TW} \cdot w_{P,i} - Dur_{B,i,TW} \cdot w_{B,i}) \cdot (-TW) \\ +(Dur_{P,i,BF} \cdot w_{P,i} - Dur_{B,i,BF} \cdot w_{B,i}) \cdot (-BF) \\ +(Dur_{P,i,sector} \cdot w_{P,i} - Dur_{B,i,sector} \cdot w_{B,i}) \cdot (-ds_{sector}) \\ +(Dur_{P,i,country,euro} \cdot w_{P,i} - Dur_{B,i,country,euro} \cdot w_{B,i}) \cdot (-ds_{country,euro}) \\ +1/2 \cdot (Conv_{P,i} \cdot w_{P,i} - Conv_{B,i} \cdot w_{B,i}) \cdot (dy)^2 \\ +(w_{P,i,invested} - w_{B,i,invested}) \cdot R_{i,xch-rate} \\ +(w_{P,i,hedged} - w_{B,i,hedged}) \cdot R_{i,xch-rate} \\ +\varepsilon_{selection,intraday,residual} \end{cases} \tag{6.54}$$

where for the i-th of N sectors within portfolio P with weighting $w_{P,i}$: $y_{P,i}$ is the yield to maturity; $Dur_{P,i,PS}$ is the duration against a 100 basis point basis curve parallel shift PS; $Dur_{P,i,TW}$ is the duration towards a 100 basis point basis curve twist TW; $Dur_{P,i,BF}$ is the duration with respect to a 100 basis point basis curve butterfly BF; $Dur_{P,i,sector}$ is the duration related to a 100 basis point change of the spread between the yield of credit instruments and the basis curve ds_{sector}; $Dur_{P,i,country,euro}$ is the duration versus a 100 basis point tightening or widening of the spread between the yield of euro-denominated government instruments and the basis curve $ds_{country,euro}$; $Conv_{P,i}$ is the convexity of the price/yield relationship; $w_{P,i,invested}$ is the weight of the absolute invested currency exposure and $w_{P,i,hedged}$ is the weight of the absolute hedged currency exposure, respectively, towards the appreciation or depreciation of the exchange rate of the portfolio base currency versus the local currency of the considered sector $R_{i,xch-rate}$; $\varepsilon_{selection,intraday,residual}$ is the remaining fraction of the active return. The analogous notation is valid for sector i within benchmark B.

Equation (6.54) can be rewritten as

$$\begin{aligned} AR_{P,base} = {} & PC_{P,carry} \\ & + PC_{P,PS} + PC_{P,TW} + PC_{P,BF} + PC_{P,sector} + PC_{P,country,euro} \\ & + PC_{P,convexity} + PC_{P,currency-invested} + PC_{P,currency-hedged} \\ & + PC_{P,selection,intraday,residual} \end{aligned} \tag{6.55}$$

where compared with equation (6.53): $PC_{P,govtYldChg}$ is split into the components impacted by the parallel shifts $PC_{i,PS}$, twists $PC_{i,TW}$ and butterflies

$PC_{i,BF}$; $PC_{P,spreadChg}$ is divided into the partitions related to sector spread changes $PC_{i,sector}$ and euro country spread changes $PC_{i,country,euro}$.

After having determined a market risk factor model to be used for performance attribution, a procedure must be chosen to estimate the model coefficients. Our example is based on fundamental factor models as they were described in Section 2.2; this means that the sensitivities towards the risk factors (i.e. the factor loadings) are determined explicitly, e.g. via the formulas (6.37), (6.38) and (6.39), and the risk factor magnitudes themselves are estimated implicitly via cross-sectional regression analysis (following the concept of empirical multi-factor models, the regression could be directly done in one step and it does not necessarily have to be divided into two parts, like e.g. the Arbitrage Pricing Theory prescribes).

Though, to be able to clearly separate government yield change effects from spread change effects, the regression should be divided into a government risk part and a credit risk part. So first the regression is done based on the individual returns of a government securities universe to determine the parallel shift, twist and butterfly movements of the government curve and additionally the shifts of the euro country spreads in a *discrete* single period Δt (e.g. one day). The regression equation applicable to every instrument i of the universe within period Δt would then be as follows (with durations being as of the basis date):[23]

$$\begin{aligned} & R_{i,base} - y_i \cdot \Delta t - R_{i,convexity} - R_{i,currency} \\ & = Dur_{i,PS} \cdot (-PS) + Dur_{i,TW} \cdot (-TW) + Dur_{i,BF} \cdot (-BF) \\ & \quad + Dur_{i,country,euro} \cdot (-\Delta s_{country,euro}) + \varepsilon \end{aligned} \tag{6.56}$$

The dependent regression variable is the difference between the buy-and-hold return[24] of instrument i in base currency $R_{i,base}$ and those return components that are deterministic or directly observable on the market, i.e. caused by the passage of time and due to the variation of the exchange rate of the portfolio base currency versus the local currency of the instrument in period Δt.[25] Based on the factor sensitivities of the universe instruments

[23] Note: the subscripts P and B are omitted because of the universe instruments' independencies of any portfolio or benchmark allocations.

[24] The time-weighted rate of return cannot be used as the dependent variable, because the incorporated influences of trading activities naturally cannot be explained by a market risk factor regression model.

[25] Additionally the convexity return contribution is subtracted as no regression beta values need to be determined for this risk factor category.

(representing the independent variables of the regression model), the risk factor changes can then be estimated via standard ordinary least squares analysis (OLS).

Subsequently, the basis date spread durations of a representative universe of credit risk-bearing securities are sector-specifically regressed on the residual returns (after subtracting the government curve-induced return portions) ε_{credit} in period Δt to derive the best-fitting sector spread changes, e.g. via OLS:

$$\varepsilon_{credit} = Dur_{i,spr} \cdot (-\Delta s_{sector}) + \varepsilon_{specific} \tag{6.57}$$

The residual term $\varepsilon_{specific}$ could be interpreted as the specific contribution from instrument selection which represents market movement-related elements not attributed so far, like the effect originating from issue- or issuer-specific spread changes.

For central banks with eligible and traded instruments with embedded options (e.g. callable bonds) and/or prepayment facilities (e.g. asset-backed securities) the causes of the active return on the portfolio versus the benchmark would probably not satisfactorily be captured by the exemplary performance attribution model as described so far, because two return-driving determinants would be missing: volatility changes of the basis instruments and prepayment rate fluctuations. By also incorporating these risk factors into the analysis, the attribution model grows to[26]

$$\begin{aligned} AR_{P,base} = {} & PC_{P,carry} + PC_{P,PS} + PC_{P,TW} + PC_{P,BF} + PC_{P,sector} \\ & + PC_{P,country,euro} + PC_{P,convexity} + PC_{P,currency-invested} \\ & + PC_{P,currency-hedged} + PC_{P,volatility} + PC_{P,prepayment} \\ & + PC_{P,selection,intraday,residual} \end{aligned} \tag{6.58}$$

where for portfolio P: $PC_{P,volatility}$ is the performance contribution with respect to the volatility variations of the underlying securities of the derivative instruments; $PC_{P,prepayment}$ is the effect of fluctuations of the prepayment rates on the active return.

[26] Equation (6.58) is also applicable to levels beneath the total portfolio level, i.e. from a security level upwards; for the aggregation of attribution effects across portfolio sector levels see formula (6.45) as well as equation (6.63) as used in a specific model context.

Last but not least, a performance-generating effect which has not been explicitly discussed so far is the intraday trading effect.[27] In principle, there are two methodically alternative ways how to determine it: implicitly by applying a holdings-based attribution system and explicitly by applying a transaction-based attribution system. The dependent variable which is to be explained within the holdings-based attribution framework naturally is the buy-and-hold return, whereas the variable to be decomposed within the transaction-based attribution analysis is the time-weighted rate of return. So the first approach is solely based on the instrument and portfolio market value changes, respectively, and the second method also incorporates transactions data into the analysis, so that it allows the direct calculation of the explicit effects induced by the transactions prices (in comparison with the valuation prices of the same day). In order to complement the holdings-based attribution method by the intraday trading effect, the latter could be indirectly determined by relating the buy-and-hold return provided by the performance attribution system to the time-weighted rate of return delivered by the performance measurement system. There is no consensus among experts as to which method should be preferred and would be more appropriate for practical use and there are pros and cons for each of the two approaches. Explicitly including the intraday trading effect $PC_{P,intraday}$ into the model, equation (6.58) becomes

$$\begin{aligned} AR_{P,base} = {} & PC_{P,carry} + PC_{P,PS} + PC_{P,TW} + PC_{P,BF} + PC_{P,sector} \\ & + PC_{P,country,euro} + PC_{P,convexity} + PC_{P,currency-invested} \\ & + PC_{P,currency-hedged} + PC_{P,volatility} + PC_{P,prepayment} \\ & + PC_{P,intraday} + PC_{P,selection,residual} \end{aligned} \tag{6.59}$$

The component 'residual' of the composite effect item $PC_{P,selection,residual}$[28] represents any inaccuracies, e.g. due to different pricing sources and/or freezing times for the securities' prices in the performance measurement system and for the yield curves in the performance attribution system. For multi-currency portfolios the residual effect is also caused by the following

[27] The impact of the intraday trading effect will naturally be of greater significance and importance for the return and performance attribution analysis of active investment portfolios than of passive benchmark portfolios.

[28] In attribution modelling it is impossible to disentangle the performance contribution stemming from security selection from the model noise. The only way to quantify the magnitude of the real residual effect and hence to assess the explanatory quality of the model would be to define some clear-cut positions (e.g. separate outright duration and curve positions) for testing purposes, to run the attribution analysis for the positions and to verify whether the model attributes the active return accordingly and the remaining performance portion is equivalent to zero.

two contradictory concepts: on the one hand the local return and the currency return are combined multiplicatively and on the other hand the base currency return (and performance) is decomposed additively in arithmetic models – the derivation of this intra-temporal cross product is shown in formula (6.49).

The above-described way to carry out performance attribution analysis is only one example among others. To give the reader an impression of a completely diverging approach which was published in a renowned journal, the Lord (1997) model is outlined. It is most probably the first publication of an explicit performance attribution technique for interest rate-sensitive portfolios (to differentiate, the approach proposed in Fong *et al.* (1983) is probably the first published model for the return decomposition of interest rate-dependent portfolios). It incorporates the concept of the so-called duration-matched Treasury bond (DMT) which represents the duration-specific level of the corresponding government yield curve. In the attribution model, to every portfolio bond a synthetic DMT (originating from the government yield curve) is assigned – per definition the duration of the DMT at the beginning of the analysis period is identical to the duration of the bond it was selected to match.

Contrary to the exemplary scheme described before, the Lord model is based on pricing from first principles and decomposes the local-currency buy-and-hold return on an interest rate-sensitive instrument $R_{i,\Delta t}$ in period $\Delta t = [t\text{–}1;t]$ generally into the components income return and price return – according to the total return formula:

$$R_{i,\Delta t} = \frac{P_{i,t} + AI_{i,t} + CP_{i,t} - P_{i,t-1}}{P_{i,t-1}} = \underbrace{\frac{AI_{i,t} + CP_{i,t}}{P_{i,t-1}}}_{\textit{income return}} + \underbrace{\frac{P_{i,t} - P_{i,t-1}}{P_{i,t-1}}}_{\textit{price return}} \qquad (6.60)$$

where for security i as of day t: $P_{i,t}$ is the market price per end of the day; $AI_{i,t}$ are the accrued interest; $CP_{i,t}$ are the coupon payments.

The income return comprises the accrued interest and coupon payments during the analysis period (i.e. the deterministic ordinary income). By further dividing the price return into respective components, the model is expressed in terms of risk factor-induced sub-returns in the following way (by omitting the subscript dt):[29]

[29] A similar approach to the Lord (1997) model can be found in Campisi (2000) – but herein the return is broken down into fewer components.

$$R_i = R_{i,income} + R_{i,carry} + R_{i,govtYldChg} + R_{i,spreadChg} + R_{i,residual} \tag{6.61}$$

where $R_{i,income}$ is the income return; $R_{i,carry}$ is the carry return;[30] the government yield change return $R_{i,govtYldChg}$ is quantified based on the yield change of the corresponding DMT; the spread change return $R_{i,spreadChg}$ is the return portion that was generated by the narrowing or widening of the spread between the yield of the instrument and the corresponding DMT; $R_{i,residual}$ is the remaining return fraction.

Breaking the yield and spread change returns further down leads to

$$\begin{aligned} R_i = {} & R_{i,income} + R_{i,carry} + R_{i,PS} + R_{i,YC} + R_{i,sector} \\ & + R_{i,issuer} + R_{i,residual} \end{aligned} \tag{6.62}$$

where the parallel shift return $R_{i,PS}$ measures the partition of the government yield change return that was induced by the change of the yield of the five-year government bond, and the yield-curve return $R_{i,YC}$ is the remaining partition; the sector spread change return $R_{i,sector}$ is the component of the spread change return due to the variation of the option-adjusted spread; and the issue- or issuer-specific spread change return $R_{i,issuer}$ is the remaining component.

Due to several oversimplifying assumptions (e.g. the parallel shift return is based on a single vertex at the government yield curve), the Lord model obviously would generate substantial methodically-induced distorted return (and consequently also performance) decompositions and hence residual effects. Therefore it will most probably not find its way to investment management practice in central banks and other public wealth management institutions where only a limited leeway for position taking versus the benchmark and correspondingly relatively small active returns are prevalent. It was incorporated into the chapter to demonstrate an alternative way how to combine diverse elementary attribution concepts within one model, i.e. in this case the DMT approach and pricing from first principles, thus representing a bottom-up scheme.[31] The factor-specific sub-returns on security level can then be aggregated to any level A via their market value weights (up to the total portfolio level). Taking the differences between the

[30] Here, the carry return consists of the contributions from rolling down the yield curve as well as the accretion (or decline) of the instrument's price toward par.

[31] One significant input variable for the decision to opt for or against a security level-based attribution modelling technique will naturally be the objective class of recipients of the attribution reports.

factor-related portfolio and benchmark return contributions provides the corresponding risk factor contributions to the active return at level A:

$$PC_{A,k} = \sum_{i=1}^{N} (R_{P,A,i,k} \cdot w_{P,A,i}) - \sum_{i=1}^{N} (R_{B,A,i,k} \cdot w_{B,A,i}) \tag{6.63}$$

where $PC_{A,k}$ is the performance contribution of the k-th risk factor to the active return at level A; $R_{P,A,i,k}$ and $R_{B,A,i,k}$ are the sub-returns related to the k-th risk factor and the i-th sector within level A of portfolio P and benchmark B, respectively; $w_{P,A,i}$ and $w_{B,A,i}$ are the weights of the i-th sector within level A of portfolio P and benchmark B, respectively; N is the number of sectors within level A.

A version of to the duration-matched Treasury bond (DMT) concept as described by Lord (1997) was applied at the European Central Bank some years ago in a first approach to develop a fixed-income performance attribution technique. The following functionality was characterizing: a spread product (e.g. an agency bond, a swap or BIS product), which de facto simultaneously embodies a duration and spread position, was matched to a risk-free alternative with similar characteristics regarding maturity and duration. This risk-free alternative (which is defined as a government security in some central banks like the ECB[32] or the swap curve in others like the Bank of Canada) was called the reference bond. The duration effect of the position was then calculated assuming the portfolio manager had bought the reference bond and the spread effect was deduced using the spread developments of the e.g. agency bond and its reference bond.[33] In comparison, the current ECB solution is explained in the subsequent section.

5. The ECB approach to performance attribution[34]

The management of the ECB's foreign reserves is a decentralized process where Eurosystem National Central Banks (NCBs) act as agencies of the

[32] In that case the reference bond of a government security is the government security itself.

[33] Alternatively to bottom-up approaches, like the Lord (1997) model, also top-down fixed-income attribution analysis techniques can be found in the literature. As an example, the methodology proposed by Van Breukelen (2000) is based on top-down investment decision processes which rely on weighted duration bets. The Van Breukelen method delivers the following attribution results: performance contribution of the duration decisions at total portfolio level as well as the effects attributed to asset allocation, instrument selection and interaction at sector level.

[34] We would like to thank Stig Hesselberg for his contribution to the ECB performance attribution framework.

ECB. The ECB has therefore worked closely with NCBs in developing its approach to performance attribution. The multivariate solution (following the idea of multi-factor return decompositions models as explained in Section 2) also preferred by the NCBs takes the following risk factors into account: carry (i.e. the passage of time), duration (i.e. the parallel shift of the basis government yield curve), yield curve (i.e. the change of slope and curvature of the basis government yield curve) and the change of credit spreads (in terms of sector spreads)[35]. Furthermore, the coverage of the impacts from intraday trading, securities lending and instrument selection was also considered as being important.

As visualized in formula (6.10) the local-currency return on an interest rate-sensitive instrument (and thus also portfolio) is generally composed of the time decay effect and the yield change effect. Following equation (6.19) the linear and the quadratic component of the yield change effect could be disentangled from each other, and by further decomposing the linear yield change effect into a basis government yield change effect and spread change effect, the price/yield change relationship could be represented as done in formula (6.20). Footing on these theoretical foundations, a conceptually *unique* framework was developed for fixed-income performance attribution at the ECB and is sketched in this section.

The market risk factor model developed by the Risk Management Division of the European Central Bank and applied to the performance attribution analysis of the management of the currency reserve and own funds portfolios of the ECB is based on a key rates approach (see Section 3.3) to derive the carry effect, government yield change effect, sector and euro country spread change effects and convexity effect. The key rates represent the points on a term structure where yield or spread changes are considered important for the evolution of the market value of the portfolio under consideration. Typically, higher granularity would be chosen at the short end of the term structure; for the ECB performance attribution model the following key rate maturities are defined: 0, 0.25, 0.50, 0.75, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 15, 20, 25 and 30 years. As all of the coding was done with the programming language 'Matlab', which is a very fast matrix-operating and code-interpreting tool, despite of the relative large number of chosen key rate maturities with respect to diverse yield curves, the procedure of running the attribution analysis (even for longer periods) only takes a very short time.

[35] Complementary to the sector spread effect the euro country spread effect is also covered by the ECB model.

The idea of the methodology is based on decomposing the whole portfolio into cash flows belonging to different fixed-income asset classes (related to different yield or spread curves). Each of these cash flows can be viewed as a zero-coupon bond, characterized by the time to maturity and the future (nominal) payment. All the payments are then reallocated to the nearest key rates in such a way, that the net present value and the modified duration of the payments are preserved. For example, a payment due in 1.5 years would be redistributed to the one-year key rate and the two-year key rate with roughly half of the nominal amount at each. This decomposition transforms the vast array of portfolio cash flows into a limited number of synthetic zero-coupon bonds (corresponding to the number of key rates) that are much more tractable and still approximate closely those of the actual portfolio.

The main strength of this approach is that it gives a more precise approximation of actual position size and impact than allocation by time bucket does. This is achieved partly because no discrete jumps occur in the allocation of the exposures of a bond when it moves from one key rate maturity (and hence also key rate maturity bucket) to another. The high precision is especially important for portfolio management with low tracking error, as it is mostly true for central banks and other public investors. Further, the method establishes a clear correspondence between performance attribution and the positions as they are taken by the portfolio managers. In a set-up with limited leeway for position taking, portfolio managers are likely to choose their over- and underweight across credit classes and across the maturity spectrum carefully. This is captured by the key rate exposures and directly attributed to the performance. It should also be noted that the method is equally useful for performance attribution and return attribution.

Recall that in continuous time the price of a zero-coupon bond P is calculated as

$$P = CF_{T-t} \cdot e^{-y_{T-t} \cdot (T-t)} \tag{6.64}$$

where CF_{T-t} is a cash flow with time to payment $T–t$ and y_{T-t} is the continuously-compounded bond yield. Further we note that the first derivative of the bond price, with respect to the yield to maturity, is

$$\frac{\partial P}{\partial y} = -(T-t) \cdot CF_{T-t} \cdot e^{-y_{T-t} \cdot (T-t)} = -(T-t) \cdot CF_{T-t} \cdot D_{T-t} \tag{6.65}$$

where the discount factor is defined as $D_{T-t} = e^{-y_{T-t} \cdot (T-t)}$.

Correspondingly, the price of a synthetic zero-coupon bond P at a key rate maturity $X - t$ is written as

$$P = CF_{X-t} \cdot e^{-y_{X-t} \cdot (X-t)} = CF_{X-t} \cdot D_{X-t} \tag{6.66}$$

When an actual cash flow CF_{T-t} is distributed to the two neighbouring key rates as the cash flows CF_{S-t} and CF_{U-t}, the total market value and the interest rate sensitivity should be preserved. These restrictions can be written as[36]

$$CF_{S-t} \cdot D_{S-t} + CF_{U-t} \cdot D_{U-t} = CF_{T-t} \cdot D_{T-t} \tag{6.67}$$

$$\begin{aligned}(S-t) \cdot CF_{S-t} \cdot D_{S-t} + (U-t) \cdot CF_{U-t} \cdot D_{U-t} \\ = (T-t) \cdot CF_{T-t} \cdot D_{T-t}\end{aligned} \tag{6.68}$$

Rearranging provides the following result:

$$CF_{S-t} = \frac{CF_{T-t} \cdot D_{T-t} \cdot [(T-t) - (U-t)]}{D_{S-t} \cdot [(S-t) - (U-t)]} \tag{6.69}$$

and

$$CF_{U-t} = \frac{CF_{T-t} \cdot D_{T-t} \cdot [(T-t) - (S-t)]}{D_{U-t} \cdot [(U-t) - (S-t)]} \tag{6.70}$$

These formulas allow us to distribute all actual cash flows to the nearest key rates (and thereby express all the exposure in terms of a few zero-coupon bonds) and at the same time preserve the market value and the modified duration.[37]

To facilitate a simple interpretation of the results, one curve (i.e. the reference government curve) is quoted in absolute levels in order to form the basis, while all other yield curves are quoted in terms of the spread to this curve. The theoretical market value impact on a portfolio can be calculated using observed levels and changes for the relevant curves; likewise the impact of the difference between the portfolio and its benchmark can

[36] For a comparison of cash flow mapping algorithms see Henrard (2000).

[37] Note that the convexity (second-order derivative of the bond price with respect to the yield) is not preserved. It can be shown that in most cases convexity will not be the same. This effect, however, is very limited and will not generally be an issue of concern.

also be calculated. Referencing to Sections 3.1 and 3.2, the impact of any interest exposures on the zero-coupon bond price, irrespective of being in absolute (i.e. nominal) terms or relative (i.e. percentage) terms, can be quantified using Taylor expansions.[38] For further purposes the price change effects are normalized to one unit of cash equivalent exposure. The formulae are applied to the lowest discrete increment in time available, i.e. $dt = \Delta t$ is one day (1/365 years) and $dy_{X-t} = \Delta y_{X-t}$ is a daily change in yields.

The price effects of first and second order $\Delta P_{keyRateChg,X-t}$ per unit of cash equivalent exposure due to a key rate change are approximated by[39]

$$\Delta P_{keyRateChg,X-t} \approx -(X-t) \cdot CF_{X-t} \cdot e^{-y_{X-t} \cdot (X-t)} \cdot \Delta y_{X-t} + \frac{1}{2} \cdot (X-t)^2 \cdot CF_{X-t} \cdot e^{-y_{X-t} \cdot (X-t)} \cdot (\Delta y_{X-t})^2 \Big|_{CF_{X-t}=1} \quad (6.71)$$

By applying the concept of distinguishing between a reference government curve and the corresponding spread curves, the first-order approximation in formula (6.71) quantifies both the basis government yield change effect $\Delta P_{govtYldChg,X-t}$ when related to the reference government curve and the spread change effect $\Delta P_{spreadChg,X-t}$ when related to the spread curves. Consequently, applying this method avoids to explicitly compute the securities' spread durations and to estimate the spread changes via regression analysis as presented in the example of Section 4.2.

The basis government yield change effect is further broken down into the components related to the parallel shift and the curve reshaping. Alternatively to the return-based approaches outlined in Section 3.3, the segregation was done at the exposure side by dividing the relative positions versus the benchmark into outright duration and curve positions following a specific algorithm which was defined by the ECB Risk Management Division. As there is no unique way to do the decomposition, the objective was to use an algorithm that corresponds best with the intentions and the style of the portfolio managers when taking positions. The ECB approach satisfies the following three conditions. First, the sum of the key rate positions assigned to the overall 'curve position' is zero. In other words, the derived 'curve

[38] The Taylor expansion technique guarantees the determination of distinct price effects by avoiding any overlapping effects and therefore represents a potential alternative to regression analysis in the context of performance attribution analysis.

[39] To adequately capture the key rate change impacts which arise from the cash flows assigned to key rate zero (but for which the original times to maturity are naturally greater than zero), appropriate portfolio- and benchmark-specific values must be chosen for the expression $X-t$ in formula (6.71).

position' is neutral (to the first-order approximation) to a parallel shift in yields. Second, the 'outright duration position' does not contain two key rate positions of different sign. And finally, no curve and outright duration positions are assigned to key rates where no position is actually taken. The applied modelling technique elegantly circumvents the requirement to especially establish a functional model, like the frequently used schemes by Nelson and Siegel (1987) and Svensson (1994), to derive the partial government curve movements and/or durations as described in Section 3.3.

Again following the Taylor expansion rule, the first- and second-order price effects $\Delta P_{carry,X-t}$ per unit of cash equivalent exposure due to the carry are approximated by

$$\Delta P_{carry,X-t} \approx -y_{X-t} \cdot CF_{X-t} \cdot e^{-y_{X-t} \cdot (X-t)} \cdot \Delta t + \frac{1}{2} \cdot y_{X-t}^2 \cdot CF_{X-t} \cdot e^{-y_{X-t} \cdot (X-t)} \cdot (\Delta t)^2 \Big|_{CF_{X-t}=1} \tag{6.72}$$

To avoid potential inaccuracies of a perturbational model approach, as mentioned in Section 3.1, three further components are included into the effect calculation. The first is the so-called 'roll-down', related to the fact that the yield of a security may change only due to the passage of time if the yield curve is not flat. This is not taken into account in the static approach for the carry above. If the yield curve is upward sloping and unchanged then the effective yield of a zero-coupon bond will decrease (and the price will increase) simply due to the passage of time. The steepness of the yield curve and the corresponding yield change of one day (1/365 years) are calculated and inserted into the formula for the key rate change effect (6.71) to determine the price effect $\Delta P_{roll-down,X-t}$ due to the roll-down per unit of cash equivalent exposure:

$$\Delta P_{roll-down,X-t} \approx (X-t) \cdot CF_{X-t} \cdot e^{-y_{X-t} \cdot (X-t)} \cdot \Delta y_{roll-down,X-t} - \frac{1}{2} \cdot (X-t)^2 \cdot CF_{X-t} \cdot e^{-y_{X-t} \cdot (X-t)} \cdot (\Delta y_{roll-down,X-t})^2 \Big|_{CF_{X-t}=1} \tag{6.73}$$

where the yield or spread change caused by the roll down $\Delta y_{roll-down,X-t}$, i.e. the change in yield or spread due to the slope between the key rate maturity $X-t$ and the key rate maturity to the left $(X-t)_{-1}$ on an unchanged curve, is given by[40]

[40] The impact of the roll-down on the yield is expressed with opposite sign in equation (6.74) to fit with formula (6.73).

$$\Delta y_{roll-down,X-t} = \frac{y_{X-t} - y_{(X-t)_{-1}}}{(X-t) - (X-t)_{-1}} \cdot \Delta t \tag{6.74}$$

The second complementary price effect is derived as a part of the Taylor expansion and is the cross effect (i.e. the interaction) between the change in time and the change in yield. The price effect $\Delta P_{cross,X-t}$ per unit of cash equivalent exposure from this is approximated as

$$\Delta P_{cross,X-t} \approx \left(1 - (X-t)_{X-t}\right) \cdot CF_{X-t} \cdot e^{-y_{X-t}\cdot(X-t)} \cdot \Delta t \cdot \Delta y_{X-t}\big|_{CF_{X-t}=1} \tag{6.75}$$

The third supplementary effect is due to the alternative cost which is associated with one of the sources for extra income when a credit bond is in the portfolio. Because all cash flows are expressed in terms of cash equivalents, the effect of buying a bond that is cheaper then a government bond (besides the carry and roll-down) needs to be captured separately. The money saved by buying the cheap bond is invested at the risk-free rate and is hence also increasing the total return from this position. Even though this effect is also small it is added for completeness. If i_{X-t} is the yield of the credit bond then the net present value of the money V_0 invested at the risk-free rate will be

$$V_0 = CF_{X-t} \cdot e^{-y_{X-t}\cdot(X-t)} - CF_{X-t} \cdot e^{-i_{X-t}\cdot(X-t)} \tag{6.76}$$

In turn, this is inserted into the formula for the carry effect (6.72) using for y_{X-t} the risk-free rate to achieve the corresponding price effect $\Delta P_{alt\text{-}cost,X-t}$ per unit of cash equivalent exposure:

$$\begin{aligned}\Delta P_{alt-cost,X-t} = {} & CF_{X-t} \cdot e^{-y_{c=govt,X-t}\cdot(X-t)} \\ & - CF_{X-t} \cdot e^{-(y_{c\neq govt,X-t}+y_{c=govt,X-t})\cdot(X-t)}\big|_{CF_{X-t}=1}\end{aligned} \tag{6.77}$$

where $c{=}govt$ indicates the reference government yield curve and $c{\neq}govt$ designates the set of relevant spread curves.

For the final reporting the effect of carry, the cross (interaction) effect, the roll down effect and the effect due to alternative cost are all added together to become the aggregate $\Delta P_{carryEtc,X-t}$:

$$\begin{aligned}\Delta P_{carryEtc,X-t} = {} & \Delta P_{carry,X-t} + \Delta P_{roll-down,X-t} + \Delta P_{cross,X-t} \\ & + \Delta P_{alt-cost,X-t}\end{aligned} \tag{6.78}$$

The terminology (denotation) for this composite effect is still 'carry' in the final report – intuitively this is the most obvious choice and since the effect of carry is by far the largest of the four effects, this improves the readability of the final report without any significant loss of information.

At the exposure side, by differentiating between a basis government curve and dependent spread curves, a government bond affects only the exposures related to the reference government curve,[41] while a credit bond affects both the reference government curve exposures and the exposures towards the spread curve associated with the bond's asset class. This naturally divides the performance into a part related to government exposures and the parts related to exposures to different classes of credit instruments (consequently the application of an artificial separation approach like the duration-matched Treasury bond method as described in Lord 1997, and sketched in Section 4.2, is not of relevance). At total portfolio level,[42] this can be formalized by the following two equations:

$$Exp_{P,govt,X-t,t} = \sum_{\forall c} \sum_{\forall i} CF_{P,i,c,X-t,t} \tag{6.79}$$

$$Exp_{P,c,X-t,t} = \sum_{\forall i} CF_{P,i,c,X-t,t} \; \forall c \neq govt \tag{6.80}$$

where for analysis day t: $Exp_{P,govt,X-t,t}$ is the absolute cash equivalent exposure of portfolio P with respect to the reference government curve at key rate maturity X–t; $Exp_{P,c,X-t,t}$ related to $c \neq govt$ is the absolute cash equivalent exposure of portfolio P related to a curve c different from the reference government curve at key rate maturity X–t; $CF_{P,i,c,X-t,t}$ is a cash flow of instrument i held in portfolio P which is assigned to key rate maturity X–t of curve c.

In order to quantify the market model-related performance contributions, the absolute cash equivalent exposures of the portfolio at all key rates across all included curves on all analysis days under consideration are compared with those of the re-scaled[43] benchmark (based on the two sets of

[41] In this context the ECB own-funds portfolio represents an exceptional case as it contains euro-denominated assets. The German government yield curve was chosen as the basis government yield curve and therefore positions in non-German government issues will contribute to the euro country spread exposure.

[42] The ECB performance attribution framework was designed to report the effects at total portfolio level.

[43] Due to the fact that the exposures are quoted as cash equivalents, the risk factor exposures of a benchmark have to be adjusted by the market value ratio of the considered portfolio and the benchmark on every day of the analysis period.

synthetic zero bonds) and the differences reflect the position of the portfolio manager on a given day:

$$RelExp_{P-B,c,X-t,t} = Exp_{P,c,X-t,t} - Exp_{B,c,X-t,t} \cdot \frac{MV_{P,t}}{MV_{B,t}} \tag{6.81}$$

where for analysis day t: $RelExp_{P-B,c,X-t,t}$ is the relative cash equivalent exposure of portfolio P versus benchmark B at key rate maturity X–t of curve c; $Exp_{B,c,X-t,t}$ is the corresponding absolute exposure of benchmark B; $MV_{PF,t}$ and $MV_{BM,t}$ are the market values of portfolio P and benchmark B, respectively.

The over- and underweight versus the benchmark in each of the synthetic zero-coupon bonds are then combined with the diverse above-mentioned price effects per cash equivalent exposure unit:

$$PC_{P,k,c,X-t,t} = RelExp_{P-B,c,X-t,t-1} \cdot \Delta P_{k,c,X-t,t} \tag{6.82}$$

where for analysis day t: $PC_{P,k,c,X-t,t}$ is the contribution to the performance of portfolio P with respect to risk factor k, curve c and key rate maturity $X - t$; $\Delta P_{k,c,X-t,t}$ is the price change effect related to risk factor k at key rate maturity $X - t$ of curve c.

As the various performance contributions on every single analysis day are quantified in terms of effects on cash equivalent exposures and hence expressed in currency units, the model inherently possesses the outstanding feature that the performance contributions over a longer analysis period are simply accurately determined by adding up the single-period results. The transformation to percentage (or basis point) numbers is finally done by relating the total period results to the basis portfolio market value.[44] This leads to the identical result as geometrically correctly linking the single-period percentage effects over time, and therefore the application of an explicit correction mechanism for the multi-period case of additive models as explained in Section 4.1 (e.g. the algorithm published in Frongello 2002a) is not of relevance.

The portion of the performance which is not influenced by any of the specified market risk factors was – in the ECB case – identified to be mostly caused by a combination of the following categories: intraday trading

[44] For the case of external portfolio cash flows (i.e. injections and/or withdrawals) during the analysis period, a more complex two-step re-scaling algorithm is applied to the attribution framework. First, the cumulative attribution effects are converted into basis point values by relating them to the simple Dietz basis market values and then the adjustment is done with respect to the performance based on the time-weighted rates of return taken from the performance measurement system.

activities, securities lending and instrument selection. As the ECB performance attribution model builds on a transaction-based approach (for a methodical comparison with the holdings-based alternative see the end of Section 4.2), the intraday trading effect is based on comparing the transaction prices with the corresponding bid/ask end-of-day valuation prices which were frozen in the portfolio evaluation system, and it shows the impact of trading better or worse than at freezing time.[45] To intention-specifically distinguish, the intraday trading effect is decomposed into the partition due to transactions on benchmark rebalancing days (on which also the relevant portfolios have to be adequately adjusted) and the part related to the other (i.e. trader-initiated) transactions. After having also included the impact resulting from securities lending activities into the analysis, the attribution effect item 'selection, residual' technically represents the fraction of the performance which was not explicitly explained by the ingredients of the methodology considered so far. From a fundamental point of view, this composite effect should for the most part be originated by superior or inferior instrument selection relative to the average market movements and also relative to the specific benchmark. Consequently, this means for the local-currency ECB model that the magnitude of this contributory category is mainly generated by rich/cheap trading and issue- or issuer-specific spread change effects (i.e. the part of the securities' yield change effects not described by the government yield change and credit spread change effects based on the defined curves and the applied key rate concept), and it is usually just to a minor extent caused by the model noise, i.e. the real residual or model error term which could be due to pricing inaccuracies.

The local-currency additive ECB fixed-income performance attribution model is structured as follows:

$$\begin{aligned} R_P - R_B = {} & PC_{P,carryEtc,govt} + PC_{P,carryEtc,sector} \\ & + PC_{P,duration,govt} + PC_{P,YC,govt} + PC_{P,convexity} \\ & + PC_{P,sector} + PC_{P,country,euro} \\ & + PC_{P,intraday,rebalancing} + PC_{P,intraday,rest} \\ & + PC_{P,securities-lending} \\ & + PC_{P,selection,residual} \end{aligned} \tag{6.83}$$

[45] This procedure perfectly coincides with the concept of the time-weighted rate of return for whose determination the intraday transaction-induced cash flows are related to the end-of-day market values.

where R_P is the portfolio return and R_B is the benchmark return; the groups of impacts on the performance are as follows: carry effect and the complementary effects of roll down, interaction and alternative cost with respect to relative exposures towards the basis government curve $PC_{P,carryEtc,govt}$ and separately related to the sector spread curves $PC_{P,carryEtc,sector}$; the effect of outright duration positions and the parallel shift of the basis government curve $PC_{P,duration,govt}$; the effect of curve positions and the reshaping of the basis government yield curve $PC_{P,YC,govt}$; the effect of the quadratic yield changes $PC_{P,convexity}$; the effect of spread positions and the narrowing and widening of sector spreads $PC_{P,sector}$ and euro country spreads $PC_{P,country,euro}$.[46] The remaining contributory group is composed of: intraday trading on benchmark rebalancing days $PC_{P,intraday,rebalancing}$ and on other days $PC_{P,intraday,rest}$; gains from securities lending $PC_{P,securities-lending}$; and a composite influence from security selection and the real residual $PC_{P,selection,residual}$.

6. Conclusions

Fixed-income performance attribution analysis should be an integral part of the investment process of central banks and other public investors that enables them to accurately identify the sources of out- or underperformance. By their ability to explicitly demonstrate the consequences of distinct managerial decisions, attribution reports contribute a significant portion to the transparency of the investment process. What makes performance attribution in general and fixed-income attribution in particular a rather challenging discipline is the fact that no standardized theoretical approaches (i.e. no 'recipes') exist and that the individual model identification is not the result of any mathematical procedure. An additional difficulty is faced in the case of most central bank portfolios for which active position taking is limited and consequently the size of out- and underperformance to be attributed to risk factors is small. Having chosen an appropriate risk factor model that fits with the individual investment decision process of a central bank by analytically deriving the portfolio-relevant return-driving components, the model could furthermore be applied to other quantitative aspects of the investment process like risk attribution, risk budgeting or portfolio and also benchmark optimization.

[46] Note that the euro country spread effect is solely relevant for the ECB own-funds portfolios and not for the foreign reserves portfolios.

An expert team of the Risk Management Division of the European Central Bank (together with selected quantitative analysts from European national central banks) recently designed and programmed its second approach for an explicit fixed-income performance attribution framework, specifically tailored to the investment process of the foreign reserves portfolios as well as the own funds portfolios of the ECB. The implemented methodology was specified in a high-level conceptual way such as to avoid shortcomings of techniques widely used in practice, and the output of the realized ECB system is already used in regular reporting.

Part II

Policy operations

7 Risk management and market impact of central bank credit operations

Ulrich Bindseil and Francesco Papadia

1. Introduction[1]

This chapter provides an overview of the risk management issues arising in central bank repurchase operations conducted to implement monetary policy. The topic will be further deepened in the next three chapters. In some sense, Chapters 7 to 10 are more focused on central banks than the rest of the book, since the risk management design of monetary policy operations is obviously also guided by policy needs. Still, many of the considerations made in this chapter are also relevant for any institution entering an agreement on the collateralization of its exposures with counterparties. In such a collateral agreement, a set of eligible assets needs to be defined, as well as risk mitigation measures, including valuation principles, haircuts and limits. The counterparties to the agreement are then constrained by it with regard to the type and amount of collateral they submit to cover exposures. However, typically, the agreements also allow some flexibility and discretion to counterparties in choosing different types of collateral, since their respective availability cannot be anticipated. As a consequence, the party receiving collateral cannot anticipate exactly what risks it will take. Even if one were to impose very tight constraints regarding the type of collateral to be used, the degree of control would not be perfect, since one cannot anticipate to what extent exposures will be created. At the end, there is a trade-off between the precision of a given collateral agreement and the flexibility allowed to the counterparties in choosing collateral,

[1] The authors are indebted to Younes Bensalah, Ivan Fréchard, Andres Manzanares, Tommi Moilainen, Ken Nyholm and in particular Vesa Poikonen for their input to the chapter. Useful comments were also received from Denis Blenck, Isabel von Köppen, Marco Lagana, Paul Mercier, Martin Perina, Francesco Mongelli, Ludger Schuknecht and Guido Wolswijk. Any remaining mistakes as well as the opinions expressed are, of course, the sole responsibility of the authors.

which brings about uncertainty on the residual risks which is taken when entering into a transaction.

Central banks implement monetary policy by steering short-term market interest rates around a target level. They do this essentially by controlling the supply of liquidity, i.e. of the deposits held by banks with the central bank, mostly by means of open market operations. Specifically, major central banks carry out open market operations, in which liquidity is provided on a temporary basis. In the case of the Eurosystem, an overall amount of close to EUR 500 billion was provided at end June 2007, of which more than EUR 300 billion was in the form of operations with a one-week maturity and the rest in the form of three-months operations.

In theory, these temporary operations could take the form of *unsecured* short-term loans to banks, offered via a tender procedure. It is, however, one of the oldest and least-disputed principles that a central bank should, under no circumstance, provide unsecured credit to banks.[2] This principle is enshrined, in the case of the Eurosystem, in article 18.1 of the Statute of the European System of Central Banks and of the European Central Bank (hereafter referred to as the ESCB/ECB Statute), which prescribes that any Eurosystem credit operation needs to be 'based on adequate collateral'.

There are various reasons behind the principle that central banks should not provide lending without collateral,[3] namely:

- Their function, and area of expertise, is the implementation of monetary policy aimed at price stability, not the management of credit risk.
- While access to central bank credit should be based on the principles of transparency and equal treatment, unsecured lending is a risky art, requiring discretion, which is neither compatible with these principles nor with the central bank accountability.
- Central banks need to act quickly in monetary policy operations and, exceptionally, also in operations aiming at maintaining financial stability. Unsecured lending would require careful and time-consuming analysis and limit setting.
- They need to deal with a high number of banks, which can include banks with a rather low credit rating.[4]

[2] For the reasons mentioned, also banks have a clear preference for collateralized inter-bank operations, and impose strict limits on any unsecured lending.

[3] For a general modelling of the role of collateral in financial markets see Bester (1987).

[4] Some central banks, including the US Federal Reserve System, conduct their open market operations only with a limited number of counterparties. However, all central banks, including the Fed, offer a borrowing facility under which they lend at a preset rate to a very wide range of banks and accept a wide set of collateral.

- They should avoid establishing credit lines reflecting the creditworthiness of different banks. A central bank can hardly stop transacting with a counterparty because its limit would have been exhausted. Such an action could be interpreted as a proof of deterioration of that counterparty's credit quality, resulting in its inability to get liquidity from the market, with potential negative financial stability consequences.
- To reflect the different degrees of counterparty risk in unsecured lending, banks charge different interest rates. By contrast, central banks have to apply uniform policy rates and thus cannot compensate the different degree of risk.

In analysing central bank collateral frameworks, this chapter, and in particular Section 3, will take a broader perspective than the rest of the book, as it will not only take a risk management perspective but also an economic perspective. The principle that all temporary refinancing operations need to be secured with collateral implies that these operations have two legs: one in central bank deposits (liquidity) and the other in collateral. The liquidity leg obviously has a decisive impact on the market for deposits: indeed, the implementation of monetary policy, consisting in achieving and maintaining a given level of interest rate, is based on this impact. It is less recognized, instead, that the collateral leg also has an influence on the market for the underlying asset. This effect is less important, but it is surprising how little it has been researched, also considering that central banks face some important choices in the specification of their collateral framework. In addition to the description of collateral frameworks in some technical documentation (see ECB 2006b for the case of the Eurosystem), there is, to our knowledge, only one comprehensive and analytical study on central bank collateral, namely the one the Federal Reserve System published in 2002 (Federal Reserve System 2002). Section 3 of this chapter aims to help fill this gap, also following the critical analyses of Fels ('Markets can punish Europe's fiscal sinners', *Financial Times* April 1, 2005), and Buiter and Sibert (2005).

The setting-up of a central bank's collateral framework may be summarized in **five phases**, which are also reflected in the organization of the sections in this chapter as well as in the other chapters on monetary policy operations in this book:

1. First, a **list of all asset types** that could be eligible as collateral in central bank credit operations has to be established. The assets in the list will have different risk characteristics, which implies that different risk mitigation measures are needed to deal with them.

2. The specific aim of **risk mitigation measures** is to bring the risks that are associated with the different types of assets to the same level, namely the level that the central bank is ready to accept.[5] Risk mitigation measures are costly and, since they have to be differentiated across asset types, their costs will also differ. The same applies to *handling* costs: some types of collateral will be more costly to handle than others. Thus, the fact that risk mitigation measures can reduce residual risks for a given asset to the desired, very low level is, of course, not sufficient to conclude that such asset should be made eligible. This also requires the risk mitigation measures and the general handling of such a type of collateral to be *cost-effective*, as addressed in the next two steps.
3. The potential collateral types should be **ranked in increasing order of cost**.
4. The central bank has to choose a **cut-off line** in the ranked assets on the basis of a comprehensive cost–benefit analysis, matching the demand for collateral with its increasing marginal cost.
5. Finally, the central bank has to **monitor** how the counterparties use the opportunities provided by the framework, in particular which collateral they use and how much concentration risk results from their choices. The actual use by counterparties, while being very difficult to anticipate, determines the residual credit risks borne by the central bank. If actual risks deviate much from expectations, there may be a need to revise the framework accordingly.

The first two and the last step are discussed in Section 2 (step 5 is also dealt with in Chapter 10). Steps 3 and 4 are dealt with in Section 3. Section 3 also discusses the effect of eligibility decisions on spreads between fixed-income securities. Section 4 concludes.

2. The collateral framework and efficient risk mitigation

This section illustrates how the collateral framework can protect the central bank, up to the desired level, against credit risk. Any central bank, like any commercial bank, has to specify its collateral and risk mitigation framework. Central banks have somewhat more room to impose their preferred specifications, while commercial banks have to follow market conventions to a larger extent. Section 2.1 discusses the desirable characteristics of

[5] See also Cossin *et al.* (2003).

eligible collateral, Section 2.2 looks at risk mitigation techniques, the specification of which may be different from asset type to asset type, and finally Section 2.3 stresses that the actual functioning of the collateral framework has to be checked against expectations.

2.1 Desirable characteristics of eligible collateral

There are a number of properties that assets should have to be suitable as collateral. Some, but not all, relate to the risks associated with the asset.

2.1.1 Legal certainty

There should be legal certainty about the transfer of the collateral to the central bank and the central bank's ability to liquidate the assets in case of a counterparty default. Any legal doubts in this regard should be removed before an asset is accepted as eligible.

2.1.2 Credit quality and easy availability of credit assessment

To minimize potential losses, the probability of a *joint* default of the counterparty and of the collateral issuer should be extremely limited. For this, both a limited correlation of default between the collateral issuer and the counterparty and a very small probability of default of the collateral issuer are important. To limit the correlation of default between the counterparty and the collateral issuer, central banks (and banks in the interbank market) normally forbid 'close links' between the counterparty and the collateral issuer. The ECB assumes the existence of 'close links' when the counterparty (issuer) owns at least 20 per cent of the capital of the issuer (counterparty), or when a third party owns the majority of the capital of both the issuer and the counterparty (see ECB 2006b). Ensuring a limited probability of default requires a credit assessment. For most marketable assets, a credit assessment is publicly available from rating agencies. For other assets (e.g. loans from banks to corporations), the central bank may have to undertake its own credit assessment, or require the counterparty to obtain such an assessment from a third party or provide its own assessment, when this is judged to be of adequate quality. Central banks typically set a minimum credit quality threshold. In the case of the ECB, this has been set to an A – rating by at least one of the three international rating agencies for rated issuers, and a corresponding 10 basis point probability of default for other debtors. The setting of a minimum rating is also standard in the interbank use of collateral, and in particular in triparty repurchase arrangements,

in which systematic eligibility criteria need to be defined. The need to define a rating threshold is particularly acute in the case of the ECB, which accepts bonds from a plurality of governments and also a wide variety of private paper. Obviously, a trade-off exists between the credit quality threshold and the amount of collateral available.

2.1.3 Easy pricing and liquidity

Preferably, assets should be *easy to price* and *liquid* so that, in case of counterparty default, they can be sold off quickly at prevailing prices, especially in case of troubled conditions in financial markets.

2.1.4 Handling costs

Handling costs should be limited: while some collateral, such as standard bonds, can be easily transferred through an efficient securities settlement system, other types of collateral may require manual handling or the setting-up of specific IT applications.[6]

2.1.5 Available amounts and prospective use

The amounts available of an asset type and its actual (or prospective) use as collateral are important to determine whether it is worth investing the resources required for its inclusion in the list of collateral (in terms of acquiring the needed expertise, financial and legal analysis, data collection, setting up / adapting IT systems, maintenance, etc.).

The asset class which ranks highest on the basis of these criteria is normally central government debt: this has a credit rating and generally a relatively high one, is highly liquid, easily handled and abundant. Also marketable, private and rated debt instruments, in particular if they have a standard structure and are abundantly available, are rather attractive. In the euro area, Pfandbriefe and other bullet bonds of banks, as well as local government debt and corporate bonds, have these characteristics. Asset-backed securities (ABSs) or collateralized debt obligations (CDOs) also normally have ratings, but tend to have special characteristics and are often less liquid. Non-marketable assets, such as bills of exchange or bank loans, rarely have credit ratings and may have higher handling costs. Finally, commodities or real estate could also be considered as eligible collateral, as they were in the past. However, the handling costs of such assets tend to be

[6] E.g. according to the Federal Reserve System (2002, 3–80): 'Securities (now most commonly in book-entry form) are very cost effective to manage as collateral; loans are more costly to manage because they are non-marketable.'

very high (see e.g. Reichsbank 1910) and there is, to our knowledge, no industrial country's central bank that currently accepts them.

The Eurosystem eligibility criteria are described in detail in ECB 2006b. The actual use of collateral in Eurosystem credit operations is described for instance in ECB (2007a, 8) – see also Chapter 10.

2.2 Risk mitigation techniques – the Eurosystem approach

Different potential collateral types imply, before the application of risk mitigation measures, differing degrees of risk for the central bank. For instance, a refinancing operation is, everything else equal, riskier if the counterparty submits as collateral an illiquid corporate bond rather than a government security. Similarly, in case of counterparty default, it is more likely that the central bank would realize a loss when liquidating an ABS, relative to a government security. Section 3.1 presents very briefly the principles applied by the Eurosystem in setting risk mitigation measures (more details on this are provided in Chapters 8 and 9), while Section 3.2 briefly deals with inter-bank standards for collateral eligibility and risk mitigation techniques.

The central bank cannot (and should not) protect itself 100 per cent from risks, since some extremely unlikely events may always lead to a loss (e.g. the sudden simultaneous defaults of both the counterparty and the issuer of the asset provided as collateral). Therefore, some optimal risk tolerance of the central bank needs to be defined and adequate mitigation measures should reduce risk to the corresponding level. Since the risk associated with collateralized operations depends, before the application of credit risk mitigation measures, on the type of collateral used, the risk mitigation measures will need to be differentiated according to the collateral type to ensure compliance with the defined risk tolerance of the central bank.

The following risk mitigation measures are typically used in collateralized lending operations.

- **Valuation and margin calls**: collateral needs to be *valued* accurately to ensure that the amount of liquidity provided to the counterparty does not exceed the collateral value. As asset prices fluctuate over time, collateral needs to be revalued regularly, and new collateral needs to be called in whenever a certain trigger level is reached. In a world without monitoring and handling costs, collateral valuation could be done on a real-time basis, and the trigger level for margin calls would at the limit be zero. In practice, costs create a trade-off. The Eurosystem, in line with

market practice, values collateral daily and has set an symmetric trigger level of 0.5 per cent, i.e. when the collateral value, after haircuts (see below), falls below 99.5 per cent of the cash leg, a margin call is triggered.

- **Haircuts**: in case of counterparty default, the collateral needs to be sold. This takes some time and, for less liquid markets, a sale in the shortest possible time may have a negative impact on prices. To ensure that there are no losses at liquidation, a certain percentage of the collateral value needs to be deducted when accepting the collateral. This percentage depends on the price volatility of the relevant asset and on the prospective liquidation time. The higher the haircuts, the better the protection, but the higher also the collateral needed for a given amount of liquidity. This trade-off needs to be addressed by setting a certain confidence level against losses. The Eurosystem, for instance, sets haircuts to cover 99 per cent of price changes within the assumed orderly liquidation time of the respective asset class. Chapter 8 provides the Eurosystem haircuts for marketable tier one assets. Haircuts increase with maturity, because so does the volatility of asset prices. In addition, haircuts increase as liquidity decreases.
- **Limits**: to avoid concentration, limits may be imposed, which can take the following form: (i) Limits for exposures to individual counterparties (e.g. limits to the volume of refinancing provided to a single counterparty). (ii) Limits to the use of specific collateral by single counterparties: e.g. percentage or absolute limits per issuer or per asset type can be imposed. For instance, counterparties could be requested to provide not more than 20 per cent in the form of unsecured bank bonds. (iii) Limits to the total submitted collateral from one issuer, aggregated over all counterparties. This is the most demanding limit specification in terms of implementation, as it requires that the aggregate use of collateral from any issuer is aggregated and, when testing collateral submission, counterparties are warned if the relevant issuer is already at its limit. This specification is also problematic as it makes it impossible for counterparties to know in advance whether a given security will be usable as collateral.

As the usage of limits always creates some implementation and monitoring costs and constrains counterparties, it is preferable, when possible, to try to set the other parameters of the framework to avoid the need for limits. This is what the Eurosystem has done so far, including the application of different haircuts to different assets. The differentiation of haircuts should also contribute to reduce concentration risk, avoiding that counterparties have

Table 7.1 Shares of different types of collateral received by 113 institutions responding to the 2006 ISDA margin survey

Type of collateral	Per cent of total	Per cent of total non-cash
Cash	72.9%	–
Bonds – total	16.4%	66.4%
• *Government securities*	*11.8%*	*47.8%*
• *Government agency securities*	*4.2%*	*17.0%*
• *Supranational bonds*	*0.4%*	*1.6%*
• *Covered bonds*	*0.0%*	*0.0%*
Letters of credit	2.2%	8.9%
Equities	4.2%	17.0%
Metals	0.2%	0.8%
Others	1.7%	6.9%

Source: ISDA. 2006. 'ISDA Margin Survey 2006', Memorandum, Table 3.1.

incentives to provide disproportionately one particular type of collateral. This could happen, in particular, if the central bank would set too lax risk control measures, thereby making the use of a given collateral type too attractive (in particular if compared to the conditions in which that asset is used in private sector transactions).

2.3 Collateral eligibility and risk control measures in inter-bank transactions

As noted by the Basel Committee on the Global Financial System (CGFS 2001), collateral has become one of the most important and widespread risk mitigation techniques in wholesale financial markets. Collateral is in particular used in: (i) secured lending; (ii) to secure derivatives positions, and (iii) for payment and settlement purposes (e.g. to create liquidity in a RTGS system). Regular updates about the use of collateral in inter-bank markets are provided by ISDA (International Swap and Derivatives Association) documents, like the *2006 ISDA Margin Survey*. According to this survey (Table 7.1.), the total estimated collateral received and delivered in 2006 would have had a value of USD 1,329 trillion. The collateral received by the 113 firms responding to the survey (which are estimated to cover around 70 per cent of the market) would have had a composition as indicated in Table 7.1.

Obviously, cash collateral is not suitable for inter-bank (or central bank) secured lending operations, in which the purpose is just to get cash. The

high share of cash collateral therefore indicates that secured lending is not the predominant reason for collateralization. Amongst bonds, Government securities and, to a lesser extent, Government agencies dominate. Also the use of equities is not negligible. The 113 respondents also reported in total 109,733 collateral agreements being in place (see ISDA Marginal Survey, 9), of which 21,889 were bilateral, i.e. created collateralization obligations for both parties, the rest being unilateral (often reflecting the higher credit quality of one of the counterparties). The most commonly used collateralization agreements are ISDA Credit Support Annexes, which can be customized according to the needs of the counterparties. Furthermore, the report notes that, in 2006, 63 per cent of all exposures created by OTC derivatives were collateralized.

The ISDA's *Guidelines for Collateral Practitioners*[7] describe in detail principles and best practices for collateralization, which are not fundamentally different from those applied by the Eurosystem (see above). Table 7.2 summarizes a few advices from this document, and checks whether or not, or in which sense, the Eurosystem practices are consistent with these advices.

It should also be noted that haircuts in the inter-bank markets may change over time, in particular they are increased in case of financial market tensions which are felt to affect the riskiness of certain asset types. For instance Citigroup estimated that, due to the tensions in the sub-prime US markets, haircuts applied to CDOs of ABSs have more than doubled in the period from January to June 2007. In particular, haircuts on AAA rated CDOs of ABSs would have increased from 2–4 per cent to 8–10 per cent, on A rated ones from 8–15 per cent to 30 per cent and on BBB rated ones even from 10–20 per cent to 50 per cent. (Citigroup Global Markets Ltd., Matt King, 'Short back and sides', July 3, 2007). The same analysis notes that 'the level of haircuts varies from broker to broker: too high, and the hedge funds will take their business elsewhere; too low, and the broker could face a nasty loss if the fund is wound up'. Changing risks, combined with this competitive pressure, thus lead to changes of haircuts across times; such changes, however, will be more limited for the standard types of collateral used in the inter-bank market, in particular for Government bonds. In contrast, central banks will be careful in raising haircuts in case of financial tensions, as they should not add to potentially contagious dynamics, possibly leading to financial instability.

[7] International Swaps and Derivatives Association. 2007. 'Guidelines for Collateral practitioners', Memorandum.

Table 7.2 Comparison of the key recommendations of ISDA Guideline for Collateral Practitioners with the Eurosystem collateralization framework

Recommendation according to ISDA Guidelines for Collateral Practitioners	Eurosystem approach
Importance of netting and cross-product collateralization for efficiency (pp. 16–9).	Netting is normally not relevant as all exposures are one-sided. Cross-product pooling is ensured in a majority of countries (one collateral pool for all types of Eurosystem credit operations with one counterparty).
Collateral should preferably be liquid, and risk control measures should depend on liquidity. Liquidity can be assumed to depend on the credit rating, currency, issue size, and pricing frequency (pp. 19–25).	Eurosystem accepts collateral of different liquidity, but has defined haircuts which differentiate between four liquidity categories.
Instruments with low price volatility are preferred. Higher volatility should be reflected in higher haircuts and lower concentration limits (p. 20).	Low volatility is not an eligibility criterion and also not relevant for any limit. However, volatilities impact on haircuts.
A minimum credit quality should be stipulated for bonds, such as measured e.g. by rating agencies (p. 20).	For securities at least one A - rating by one recognized rating agency (for credit claims an equivalent 10 basis point probability of default).
Collateral with longer duration should have higher haircuts due to higher price volatility (p. 20).	Maturities are mapped into price volatilities and therefore on haircuts (see above).
Avoid negative correlation of collateral value with exposure value (in OTC derivatives) (p. 21).	Not relevant (exposure is given by cash leg).
Avoid positive correlation between collateral value and credit quality of the issuer (p. 21).	Not specifically addressed – with the exception of the prohibition of close links (of a control type). Potential weaknesses: large amounts of unsecured bank bonds submitted (sector correlation), Pfandbriefe and ABSs originated by the counterparty itself.
Haircuts should be designed to cover losses of value due to the worst expected price move (e.g. at a 99 per cent confidence level) over the holding period, as well as costs likely to be incurred in liquidating the assets, such as commissions and taxes (pp. 21–5).	99 per cent confidence level over holding period, but nothing for commissions or taxes.
The holding period should span the maximum time lapse possible between the last valuation and possibility of a margin call, and actually being able to liquidate collateral holding in the event of default. Traditionally, the assumed holding period was one month, but practice seems to have been moving to 10 business days (p. 24).	For Government bonds, Eurosystem assumes one week (five business days) holding period, for the other three liquidity categories 2, 3 and 4 weeks, respectively.

Table 7.2 (cont.)

Recommendation according to ISDA Guidelines for Collateral Practitioners	Eurosystem approach
Low rated debt, such as that rated below investment grade, might warrant an additional haircut (p. 23).	Eurosystem does not accept BBB rated (i.e. still investment grade) collateral, so no need for additional credit haircut – see also Chapter 8
Concentration of collateral should be avoided; maximum single issuer concentration limits are best expressed as a percentage of the market capitalization of the issuer. There should be haircut implications if diversification is compromised (p. 26).	Not applied by Eurosystem.
Collateral and exposures should be marked-to-market daily (p. 38).	Yes.

2.4 Monitoring the use of the collateral framework and related risk taking

Even if thorough analytical work underlies a given collateral framework, the actual use of collateral and the resulting concentration of risks cannot be fully anticipated. This is particularly important because, in practice, an appropriate point in the flexibility/precision trade-off must be chosen when building a framework. Indeed, to remain flexible, as well as simple, transparent and efficient, a collateral framework has to accept a certain degree of approximation. But the degree of approximation which is thought acceptable ex ante may appear excessive in practice, for instance because a specific collateral type is used in a much higher proportion than anticipated.

The point can be better made with an example: the Eurosystem has defined, as mentioned above, four liquidity categories and has classified assets in these categories on the basis of institutional criteria, as shown in Chapter 8. Obviously liquidity also differs within these categories, as Table 7.3, which takes bid–ask spreads as an indicator of liquidity, shows.

For instance, while government bonds are normally very liquid, euro-denominated government bonds of e.g. Slovenia and of new EU Member States are less so – mainly due to their small size. The Eurosystem's classification of all government bonds under the most liquid category is thus a simplification. The justification for this simplification is that it does not imply substantial additional risks: even if there would be larger than expected use of such bonds, this could not create really large risks, as their

Table 7.3 Bid–ask spreads as an indicator of liquidity for selected assets (2005 data)

Liquidity category	Issuers (ratings in parentheses)	Liquidity indicator: bid–ask spread (in cent)[a]
1	Germany, France, the Netherlands and Spain (AAA)	0.5–1
1	Austria, Finland, Ireland (AAA)	1
1	Italy and Belgium (AA)	0.5–1
1	Portugal (AA/A)	1
1	Greece (A)	1
1	Slovenia (AA)	20
1	Non-euro area new EU Member States (mostly A rated)	15–20
2	German Länder (AAA-A)	3–5
2	Agencies/supranationals and Jumbo Pfandbriefe (mostly AAA)	3–5
3	Non-Jumbo Pfandbriefe (mostly AAA)	3–5

[a] Bid–offer spreads observed in normal times on five-year euro-denominated bonds in Trade Web (when available) in basis points of prices (so-called cents or ticks). Indicative averages for relatively small tickets (less than EUR 10 million). Bid–offer spreads very much depend on the size of the issue and how old it is. The difference in bid–offer spreads between the various issuers tends to increase rapidly with the traded size.

maximum use is still extremely small compared to total outstanding operations. The table also reports, for information, the ratings of the different bonds, revealing that the effect of ratings on bid–offer spreads is rather small. This is another justification (see also Chapter 8) for not introducing credit risk related haircuts into the collateral framework, as the value added of doing this would be more than compensated by the costs in terms of added complication. At the end, the approximation deriving from classifying the assets in four liquidity categories is acceptable provided that: (i) the average liquidity of each class is correctly estimated; (ii) the heterogeneity within each asset class is not too high; and (iii) the prevailing heterogeneity does not lead to severe distortions and concentration risk.

In general, the central bank, as any institution offering a collateral framework, should monitor the actual use of collateral, not only on aggregate, but also on a counterparty-by-counterparty basis, to determine whether an adjustment of the framework may be needed. It should also aim at calculating aggregate risk measures, such as a portfolio VaR figure reflecting both credit and liquidity risk. A methodology for doing so is

presented and applied to the Eurosystem in detail in Chapter 10 of this book. In practice defining an efficient collateral and risk mitigation framework has to be seen as a continuous interactive process.

3. A cost–benefit analysis of a central bank collateral framework

A central bank should aim at economic efficiency and base its decisions on a comprehensive cost–benefit analysis. In the case of the Eurosystem, this principle is enshrined in article 2 of the ESCB/ECB Statute, which states that 'the ESCB shall act in accordance with the principle of an open market economy with free competition, favouring an efficient allocation of resources'. The cost–benefit analysis should start from the condition, established in Section 2, that risk mitigation measures make the residual risk of each collateral type equal and consistent with the risk tolerance of the central bank. Based on this premise, the basic idea of an economic cost–benefit analysis is that all collateral types can be ranked in terms of the cost of their use. This will in turn depend on the five characteristics listed in Section 2.1. Somewhere on the cost schedule between the least and the most costly collateral types, the increasing marginal cost of adding one more collateral type will be equal to its declining marginal value. Of course, estimating the 'cost' and 'benefit' curves is challenging, and will probably rarely be done explicitly in practice. Still, such an approach establishes a logical framework to examine the eligibility decisions. The next sub-section provides an example of such a framework in the context of a simple model.

3.1 A simple model

The following model simplifies drastically in one dimension, namely by assuming homogeneity of banks, both in terms of needs for central bank refinancing and in terms of holdings of the different asset types. Even with this simplification, the estimation of the model appears difficult. Still, it illustrates certain aspects that might escape attention if eligibility decisions were not dealt with in a comprehensive model. For instance, if a central bank underestimated the handling costs of a specific asset type, and thus overestimated its use by counterparties in central bank operations, then it may take a socially sub-optimal decision when making it eligible.

$A = \{1 \ldots n\}$	Set of all asset types that may potentially be eligible as collateral.
$E \subset A$	Set of eligible assets, as decided by the central bank. Ineligible assets are $(A \backslash E)$ (i.e. set A excluding set E).
W_j	Available amount of asset j in the banking system which can be potentially used as collateral. This is, where relevant, after application of the relevant risk mitigation measures needed to achieve the desired low residual risk; obviously $j \in E$.
V_j	Amount of collateral j that is actually submitted to the central bank (again, after haircuts).
D	Aggregate refinancing needs of banking system vis-à-vis the central bank ('liquidity needs'). Exogenously given in our model.
K_j	Fixed cost for central bank to include asset j for one year in the list of eligible assets.
$k_j V_j$	Total variable cost for central bank of handling asset j. The costs include the costs of risk mitigation measures.
$c_j V_j$	Total variable cost for banks of handling asset j. Again, this includes all handling and assessment costs. If haircuts are high, obviously costs are increased proportionally. Moreover, this includes opportunity costs: for some collateral, there may be use in the inter-bank repurchase market, and the associated value is lost if the collateral is used for central bank refinancing.

When deciding which collateral to make eligible, the central bank has first to take note of the banking system's refinancing needs vis-à-vis the central bank (D) and it should in any case ensure that

$$\sum_{j \in E} W_j \geq D \tag{7.1}$$

Inequality (7.1) is a precondition for a smooth monetary policy implementation. A failure of monetary policy implementation due to collateral scarcity would generate very high social costs. For the sake of simplicity, we assume that D is exogenous and fixed; in a more general model, it could be a stochastic variable and the constraint above would be transformed into a confidence level constraint. In addition, collateral provides utility as a buffer against inter-bank intraday and end-of-day liquidity shocks. We assume that one has to 'use' the collateral to protect against liquidity shocks, i.e. one

has to bear the related fixed and variable costs (one can imagine that the collateral has to be pre-deposited with the central bank). For the sake of simplicity, we also assume that, as long as sufficient collateral is available, liquidity-consuming shocks do not create costs. If however the bank runs out of collateral, costs arise.

We look at one representative bank, which is taken to represent the entire banking system, thus avoiding aggregation issues. Let $r = -D + \sum_{j \in E} V_j$ be the collateral reserves of the representative bank to address liquidity shocks. Let ε be the liquidity shock with expected value zero and variance σ^2 and let F be a continuous cumulative density function and f be the corresponding symmetric density function. The costs of a liquidity shortage are p per euro. Assume that the bank orders collateral according to variable costs in an optimal way, such that $C(r)$ is the continuous, monotonously increasing and convex cost function for pre-depositing collateral with the central bank for liquidity purposes. The risk-neutral representative bank will choose $r \in [0, \sum_{i \in E} W_i]$ that minimizes expected costs G of collateral holdings and liquidity shocks:

$$E(G(r)) = E(C(r) + p\max(-r + \varepsilon, 0)) = \left(C(r) + p\int_r^{\infty} f_x(x - r)dx \right) \tag{7.2}$$

The first-order condition of this problem is (see e.g. Freixas and Rochet 1997, 228)

$$\partial C/\partial r - pF(-r) = 0 \tag{7.3}$$

The cost function $\partial C/\partial r$ increases in steps as r grows, since the collateral is ordered from the cheapest to the most expensive. The function $pF(-r)$ represents the gain from holding collateral, in terms of avoidance of costs deriving from insufficient liquidity, and is continuously decreasing in r, starting from p/2. While the first-order condition (7.3) reflects the optimum from the commercial bank's point of view, it obviously does not reflect the optimum from a social point of view, as it does not include the costs borne by the central bank. If social costs of collateral use are $C(r) + K(r)$, then the first-order condition describing the social optimum is simply

$$\partial C/\partial r + \partial K/\partial r - pF(-r) = 0 \tag{7.4}$$

Table 7.4 Example of parameters underlying a cost–benefit analysis of collateral eligibility

Category (j)	Available amount[a] (W)	Fixed costs for central bank[a] (V)	Variable unitary cost for central bank[b] (k)	Variable unitary cost for banks[b] (c)
α (e.g. government securities)	1,000,000	0	0.5	0.5
β (e.g. Pfandbriefe)	1,000,000	5	0.5	0.5
γ	500,000	5	1	1
δ	500,000	5	1	1
ε (e.g. bank loans)	500,000	20	1	2
ζ (e.g. commodities)	500,000	50	10	5

[a] in EUR billions.
[b] in basis points per year.

Consider now a simple numerical example (Table 7.4) that illustrates the decision-making problem of both the commercial and the central bank and its welfare effects. Note that we assume, in line with actual central bank practice, that no fees are imposed on the banking system for the posting of collateral. Obviously, fees, like any price, would play a key role in ensuring efficiency in the allocation of resources. In the example, we assume that liquidity shocks are normally distributed and have a standard deviation of EUR 1,000 billion and that the cost of running out of collateral in case of a liquidity shock is five basis points in annualized terms. We also assume that the banking system has either a zero, a EUR 1,500 billion or a EUR 3,000 billion structural refinancing need towards the central bank. The first-order condition for the representative bank (3) is illustrated in Figure 7.1. The intersection between the bank's marginal costs and benefits will determine the amount of collateral posted, provided the respective collateral type is eligible.

It can be seen from the chart that if $D = 0$, 1,500 or 3,000, the bank (the banking system) will post EUR 1,280, 2,340 and 3,250 billion as collateral, respectively, moving from less to more costly collateral. In particular, where $D = 3{,}000$, it will use collateral up to type ε – provided this collateral and all the cheaper ones are eligible. How does the social optimality condition on eligibility (equation (7.4)) compare with that of the commercial bank (7.3)? First, the central bank should make assets eligible as collateral to respect constraint (7.1), e.g. when $D = 1{,}500$ it needs to make eligible all category α and β assets. Beyond this, it should decide on eligibility on the basis of a social cost–benefit analysis. Considering (unlike the commercial bank that does not internalize the central bank costs) all costs and benefits,

Table 7.5 Social welfare under different sets of eligible collateral and refinancing needs of the banking system, excluding costs and benefits of the provision of collateral for refinancing needs (in EUR billions)

Eligible assets	$D=0$	$D=1{,}500$	$D=3{,}000$
A	30.2	Mon. pol. failure	Mon. pol. failure
$\alpha+\beta$	**42.8**	**45.0**	Mon. pol. failure
$\alpha+\beta+\gamma$	37.8	15.2	Mon. pol. failure
$\alpha+\beta+\gamma+\delta$	32.8	10.2	**0**
$\alpha+\beta+\gamma+\delta+\varepsilon$	12.8	−9.8	−40.0
$\alpha+\beta+\gamma+\delta+\varepsilon+\zeta$	−37.2	−59.8	−89.0

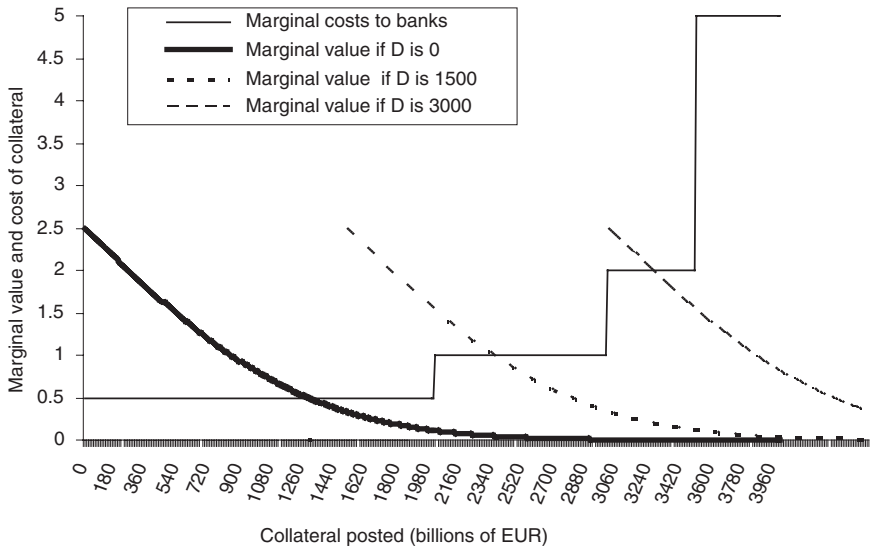

Figure 7.1. Marginal costs and benefits for banks of posting collateral with the central bank, assuming structural refinancing needs of zero, EUR 1,500 billion and EUR 3,000 billion.

Table 7.5 provides, for the three cases, the total costs and benefits for society of various eligibility decisions.

The highest figure in each column, highlighted in bold, indicates the socially optimal set of eligible collateral. It is interesting that while in the first scenario ($D = 0$) the social optimum allows the representative bank to post as much collateral as it wishes, taking into account its private benefits and costs, this is not the case in the second and third scenarios ($D = 1{,}500$ and 3,000 respectively). Here, the social optimum corresponds to a smaller

set of collateral than the one that commercial banks would prefer. The result is not surprising since the costs for the central bank enter into the social optimum but are ignored by the representative bank. Of course, the result also depends on the absence of fees, which could make social and private optima coincide.

When interpreting this model, it should be borne in mind that the model is simplistic and ignores various effects relevant in practice. Most importantly, the heterogeneity of banks in terms of collateral holdings, refinancing needs and vulnerability to liquidity shocks makes a big difference, also for the welfare analysis. As the marginal utility of collateral should be a decreasing function of the amount of collateral available, not only at the level of the aggregate banking system but also for individual banks, the heterogeneity of banks implies that the actual total social value of collateral eligibility will be higher than the aggregate represented in the model.[8] Another notable simplification is the assumption that the value of the collateral's liquidity service is constant over time. This will instead vary, and peak in the case of a financial crisis. This should be taken into account by the central bank when doing its cost–benefit analysis.

It is interesting to consider, within the example provided, the effects of eligibility choices on the spreads between different assets. Let us concentrate on the case where refinancing needs are 1,500 and the central bank has chosen the socially optimal set of eligible collateral, which is $\alpha + \beta$. The representative bank will use the full amount of available collateral (2,000) and there is a 'rent', i.e. marginal value of owning collateral of type α or β of around 1 basis point, equal to the marginal value for this amount minus the marginal cost (the gross marginal value being $pF(-r/\sigma) = 1.5$, for $p = 5$ basis points, $r = 2{,}000 - D = 500$ and $\sigma = 1{,}000$). Therefore, assuming that the ineligible asset γ would be equal in every other respect to α and β, it should trade at a yield of 1 basis point above these assets. Now assume that the central bank deviates from the social optimum and also makes γ eligible. The representative bank will increase its use of collateral to its private optimum of 2,340 and the marginal rent disappears, as private marginal cost and marginal benefit are now equalized for that amount. At the same time, the equilibrium spread between γ and α/β is now only 0.5 basis point,

[8] This is because if the utility of having collateral is for all banks a falling and convex function, then the average utility of collateral across heterogeneous banks is always higher than the utility of the average collateral holdings of banks (a bit like Jensen's inequality for concave utility functions). One could aim at numerically getting some idea of the difference this makes, depending on assumptions that would somehow reflect anecdotal evidence, but this would go beyond the scope of this chapter.

since this is the difference in the cost of using these assets as collateral. What now are the spreads of these three assets relative to asset δ? Before making γ eligible, these were −1, −1 and 0 for α, β and γ, respectively. After making γ eligible, these are −0.5, −0.5 and 0, respectively, i.e. the spread between γ and δ remains zero, and the spread between α/β and δ has narrowed down to the cost difference between the different assets. The increased 'supply of eligibility' from the central bank reduces the 'rent' given by the eligibility premium. This shows how careful one has to be when making general statements about a constant eligibility premium.

Within this numerical model, further cases may be examined. If, for $D =$ 1,500, in addition to α, β and γ, δ is also made eligible, which represents a further deviation from the social optimum due to the implied fixed costs for society, nothing changes in terms of spreads, and the amount of collateral used does not change either. The same obviously holds when asset classes ε and ξ are added. In the case $D = 3{,}000$, the social optimum is, following Table 7.5, to make assets α, β, γ and δ eligible. Very similar effects to the previous case can be observed. The rent for banks of having collateral of types α and β is now two basis points, and the rent of owning collateral of types γ and δ is, due to the higher costs, 1.5 basis points. Therefore, the spread between the two groups of assets is again 0.5 basis point. The spread between assets of type α or β and the ineligible assets of types ε and ζ is 2 basis points. After making ε eligible, the spreads between ε and all other eligible asset classes do not change (because at the margin, having ε is still without special value). However, due to the increased availability of collateral, the spreads against asset category ζ shrink by 0.5 basis point.

Finally, an alternative interpretation of the model, in which the variable costs of using the assets as collateral also include opportunity cost, is of interest and could be elaborated upon further in future research. Indeed, it could be argued that financial assets can, to a varying extent, be used as collateral in inter-bank operations, as an alternative to the use in central bank operations. Using assets as central bank collateral thus creates opportunity costs, which are high for e.g. government bonds, and low for less liquid assets, such as ABSs and bank loans, as these are normally not used as collateral in inter-bank markets. Therefore, the order in which banks would rank eligible assets according to their overall costs could be different from a ranking based only on handling and credit assessment costs, as implied above. According to this different ranking, for instance, bank loans may be 'cheaper' for banks to use than government bonds. While this underlines that the model above is a considerable simplification and should

be considered only as a first conceptual step towards a comprehensive theoretical framework, it also shows that the model can be extended to encompass different assumptions.

3.2 Empirical estimates of the effect of eligibility on yield: normal times

In the previous section, a simple model was presented to provide a framework for the decision of the central bank to make different types of asset eligible and to look at the interest rate differential between eligible and ineligible assets, dubbed the 'eligibility premium'. In this section, we seek empirical indications of the possible size of this premium. As argued above, eligibility as central bank collateral should make, everything else equal, one asset more attractive and thus increase its price and lower its yield.[9] The additional attractiveness results from the fact that the asset can provide a liquidity service, which has a positive value. The eligibility premium depends on conditions which may change over time. As was seen in the model presented above, the first time-varying condition is the overall scarcity of collateral: if the banking system has a liquidity surplus and the need for collateral for payment system operations is limited, or if there is ample government debt outstanding, then declaring an additional asset eligible will have no measurable effect on prices, as that asset would anyway not be used to a significant extent. If, in contrast, the need for central bank collateral is high, and the amounts of eligible collateral are limited, then the price effects of declaring one asset type eligible will be substantial. Similarly, the relative amount of the collateral assets newly made eligible also matters, as it also changes the overall availability of collateral and therefore its value. Thus, the price of the eligible asset A should be affected more strongly by the decision to make asset B eligible, if asset B is in abundant supply. Moreover, the eligibility premium will change in case of financial tensions, shifting the demand curve for collateral to the right. This was illustrated during the global 2007 'sub-prime' turmoil, as discussed in Section 3.3. In the following, four different approaches to quantifying the eligibility premium are presented. The values of two of these measures during times of market turmoil are then considered in Section 3.3.

[9] This effect should only be relevant if the asset will effectively be used as collateral under the chosen risk control measures and handling solutions. If, for instance, the handling solution is extremely inconvenient, or if the haircuts applied to the asset are extremely high, eligibility may not lead to practical use of the asset as collateral and would therefore be hardly relevant in terms of eligibility premium.

Table 7.6 Information on the set of bonds used for the analysis

Rating	Number of EEA bonds	Number of non-EEA bonds	Number of EEA issuers	Number of non-EEA issuers
AAA	220	18	43	5
AA	348	27	63	8
A	624	50	171	14
TOTAL	1192	95	277	27

Source: ECB Eligible Assets Database.

3.2.1 Measuring the effect on spreads of a change in eligibility

For the reasons mentioned above, an ideal opportunity to measure the effects of eligibility on spreads arises when a *small* asset category is added to a large eligible set. Such a case occurred recently in the Eurosystem when, on 1 July 2005, selected euro-denominated securities from American, Canadian, Japanese and, potentially, Swiss issuers (non-European Economic Area – non-EEA – issuers) were added to the list of eligible assets (see the ECB press releases of 21 February 2005 and 30 May 2005). This change should have lowered the yield of these instruments relative to comparable assets that were already eligible. Therefore, yields of the newly eligible assets issued by the non-EEA issuers mentioned above were compared with yields of a sample of assets of EEA issuers which had been eligible for a long time.

The set of non-EEA bonds was taken from the ECB's Eligible Assets Database on 5 October 2005. The sample of EEA bonds used for benchmarking was selected by taking all the corporate and credit bonds issued by EEA entities. Bonds issued during 2005 were removed as well as bonds having a residual maturity of less than one year since bonds near maturity tend to have a volatile option-adjusted spread. A number of bonds with extreme spread volatility were also removed. Finally, the EEA sample was adjusted to match the relative rating distribution of the non-EEA bonds. The rating classes are Bloomberg composites, i.e. averages or lowest ratings.[10] Table 7.6 shows information on the sample of bonds that were used in the analysis.

[10] The Bloomberg composite rating (COMP) is a blend of Moody's and Standard & Poor's ratings. If Moody's and Standard & Poor's ratings are split by one step, the COMP is equivalent to the lower rating. If Moody's and Standard & Poor's ratings are split by more than one step, the COMP is equivalent to the middle rating.

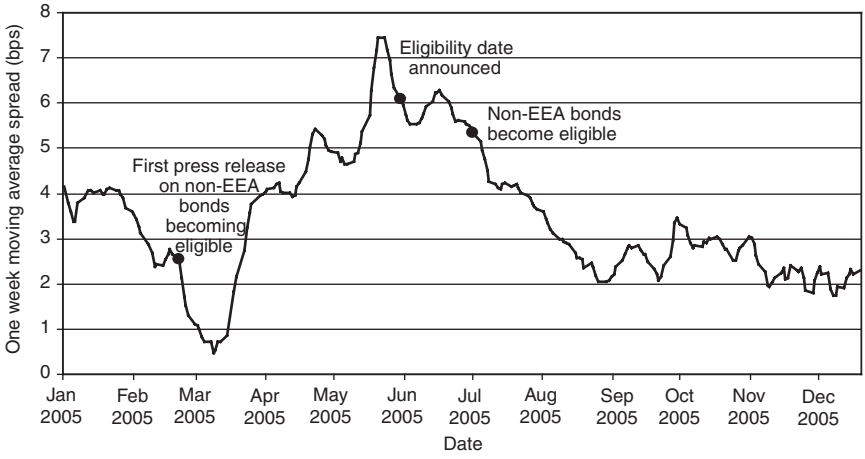

Figure 7.2. One-week moving average spread between non-EEA and EEA issuers in 2005. The spread is calculated by comparing average option-adjusted bid spreads between bonds from non-EEA and EEA issuers. The option-adjusted spread for each security is downloaded from Bloomberg. *Sources:* ECB Eligible Assets Database and Bloomberg.

Figure 7.2 shows a plot of average daily yield spreads in 2005 between non-EEA and EEA issuers. The spread is calculated by comparing average option-adjusted bid spreads between bonds from non-EEA and EEA issuers. The use of option-adjusted spreads makes bonds with different maturities and optionalities comparable. The resulting yield differential is quite volatile, ranging between 0.5 and 7.5 basis points during the year. The upcoming eligibility of bonds from non-EEA issuers was originally announced on 21 February, but the eligibility date was not yet published at that stage. The eligibility date of 1 July was announced in a second press release, on 30 May. Following each of these dates the spread seems to be decreasing, but, in fact, it had already been doing so prior to the announcements. Therefore, it is difficult to attribute the changes to the Eurosystem eligibility. Overall, the level of spreads does not seem to have changed materially from before the original eligibility announcement to the last quarter of the year in which the eligibility status was changed. To identify the possible source of the changes, one may note that eighty-seven out of ninety-five non-EEA bonds are issued by US-based companies, which suggests that the main driving forces behind the evolution of the spread are country-specific factors. Especially the major credit events during the second quarter of the year, such as problems in the US auto industry, can be assumed to have caused the widening of the spread during that period.

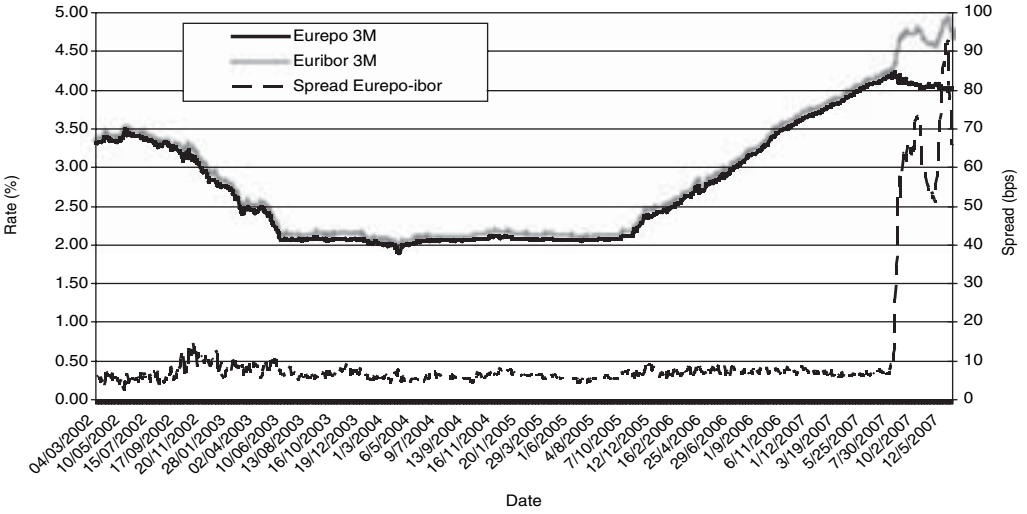

Figure 7.3. Spread between the three-month EURIBOR and three-month EUREPO rates since the introduction of the EUREPO in March 2002 – until end 2007.
Source: EUREPO (http://www.eurepo.org/eurepo/historical-charts.html).

3.2.2 Spread between collateralized and uncollateralized inter-bank repurchase operations

Another possible indicator for the eligibility premium is the spread between inter-bank uncollateralized deposits and repurchase operations, which is normally in the range of 3 to 5 basis points for the relevant maturity (see Figure 7.3). It can be argued that a bank can have access to the cheaper secured borrowing if it has eligible collateral. Thus, the spread between the two kinds of borrowing corresponds to the value of having collateral eligible for inter-bank operations. Of course, this reasoning directly holds only for the large majority of banks (in the AA and A rating range), which can indeed refinance at close to EURIBOR rates. For the few worse-rated banks, the eligibility premium will be higher. In addition, this measurement only holds in normal times: in case of liquidity stress, the spreads should widen. This is indeed what seemed to have happened in 2002, and even more in 2007, as Figure 7.3 suggests. In 2002, in particular the German banking system was considered to be under stress, including rumours of liquidity problems of individual banks, which led to a sort of flight into collateralized operations and the spread surpassed 10 basis points at the end of 2002. The second half of 2007 will be dealt with in Section 3.3. A further caveat in looking at this measure of the liquidity premium is that the set of eligible

assets for standard inter-bank repurchases ('General Collateral'[11]) is a subset of the one eligible for operations with the Eurosystem, and central bank and inter-bank repurchases have some other differences impairing a close comparison between the two.

3.2.3 Securitization aimed at having more central bank collateral

Finally, according to anecdotal evidence from the euro area, a few banks have securitized assets with the sole purpose of making them eligible as central bank collateral. Current estimates are that such securitization would have cost them around 3 basis points (per annum). The fact that this phenomenon has been observed only rarely, but that more banks have assets suitable for similar securitization, suggests that other banks are not willing to pay the 3 basis points for obtaining eligible assets. Again, this indication of the eligibility premium is subject to some caveats, as not all banks may hold sufficient assets suitable for securitization and since the cost of securitization may be higher for some banks.

3.2.4 Collateral for central bank operations and for inter-bank repurchase markets

A last remark can shed light on the specific eligibility premium that derives from the fact that some assets are only eligible for central bank operations but not for operations in the private repurchase markets. For this purpose, it is interesting to jointly consider, on the one hand, the difference between the collateral accepted in the euro area for standard inter-bank operations (so-called General Collateral, or GC) and that eligible for Eurosystem operations and, on the other hand, the relationship between the rates of interest prevailing on the Eurosystem and on the GC market operations. As regards the first point, GC essentially includes all central government bonds of the euro area, and partially Jumbo Pfandbriefe (e.g. for Eurex repurchases), while Eurosystem collateral is much wider, including many other fixed-income private instruments and some non-marketable claims.

With regard to the relationship between the interest rates prevailing on the two types of operations, the striking fact is that they are so close, both in level and in behaviour. Bindseil *et al.* (2004b) calculate, for the one-year period starting in June 2000, the spread between weighted average rates on short-term Eurosystem main refinancing operations (MROs) and rates on private repurchase operations. They note (page 14) that the former are even,

[11] 'General Collateral' according to the EUREPO definition is any euro area government debt (see www.eurepo.org).

on average, marginally lower (0.487 basis point) than the latter. This is surprising since, as stated earlier, the set of collateral eligible for inter-bank operations is smaller than the one for central bank operations and thus banks should be willing to pay a higher rate of interest on the latter operations. The result of a very close relationship between the two types of rates is confirmed by more recent observations, as illustrated by a comparison of the one week EUREPO rate with the weighted average MRO tender rate. Again, EUREPO rates tend to exceed MRO rates, but by mostly 1 or 2 basis points. This also reflects the fact that EUREPO rates are offered rates, with a typical spread in the repurchase market of around 1–3 basis points. Overall, one can interpret the results deriving from the comparison between the cost of market repurchase transactions with the cost of central bank financing as meaning that the relevance of 'collateral arbitrage', i.e. using for central bank operations the assets not eligible for inter-bank operations, is relatively limited, otherwise competitive pressure should induce banks to offer higher rates to get liquidity from the Eurosystem rather than in the GC market. However, it should also be noted that a degree of collateral arbitrage can be seen in quantities rather than in rates, as banks tend to use over-proportionally less liquid, but highly rated, private paper, such as ABSs or bank bonds, in the Eurosystem operations.

Interestingly, the relationship between the rates prevailing in the Eurosystem's three-month refinancing operations (LTROs), which have been studied by Linzert *et al.* (2007), and those determined in three-month private repurchase operations, such as reflected in EUREPO, is rather different from that prevailing for MROs. On average, the weighted average rate of LTROs was 3 basis points above the corresponding EUREPO rate in the period from March 2002 to October 2004, thus giving some evidence of collateral eligibility effects.

In summary, all the four estimates considered above consistently indicate that the eligibility premium deriving from the fact that one specific asset type is eligible as collateral for Eurosystem operations is in the order of magnitude of a few basis points. However, again, the following caveats should be highlighted: (i) In times of financial tensions, the eligibility premium is likely to be much higher – as demonstrated by the summer 2007 developments, summarized in Section 3.3. (ii) For lower-rated banks (e.g. banks with a BBB rating), the value of eligibility is likely to be significantly higher. (iii) The low eligibility premium in the euro area is also the result of the ample availability of collateral. If availability were to decrease or demand increase, the premium would increase as well.

3.3 The eligibility premium in times of a liquidity crisis: the 'sub-prime turmoil' of 2007

Section 3.2 presented a number of approaches to estimate the eligibility premium under normal circumstances, as they prevailed most of the time between 1999 and July 2007. This section turns to the period of financial turmoil whish started in August 2007, focusing on two of the measures developed in the previous section. We do not summarize the chronology of events of the summer 2007 turmoil, nor do we try to analyse here its origins, only referring to Fender and Hördahl (2007), who provide an overview of the events until 24 August 2007. We will only show that these measures of the eligibility premium suddenly took values never seen since the introduction of the euro.

Figure 7.3 shows the evolution of the three-months EURIBOR–EUREPO spread, also for the period until mid January 2008. The swelling spread for three months operations suggests that there was a general unwillingness to lend at three months. Anecdotal information also indicates that actual unsecured turnover at this maturity was very low. Thus, the better measure for the GC eligibility premium was probably the 1-week spread, which increased from an average level of 3.5 basis points between January and July 2007 to an average of 12.3 basis points in the period 1 August to 15 October 2007, indicating that the GC eligibility premium more than tripled (these averages were 6.8 and 54.9 basis points for the three-months rates, respectively).

As a second indication of a higher eligibility premium in crisis times, Figure 7.4 shows the evolution, during the summer of 2007, of the spread between the weighted average rate of MROs, and one week EUREPO and EURIBOR rates, in analogy to what was shown in the last section of the previous paragraph. As from the beginning of August, as already seen, the spread between the EURIBOR and the EUREPO increased dramatically and, on average the weighted average at the MRO operations was closer to the EURIBOR than to the EUREPO rate, indicating that counterparties bid more aggressively in the Eurosystem operations, arguably because they could use collateral that could not easily be used in private repurchase transactions, evidencing a much higher value for the eligibility premium.

Finally, Figure 7.5 shows the evolution of the weighted average rate of LTROs during 2007 and the three-month EUREPO and EURIBOR rates.

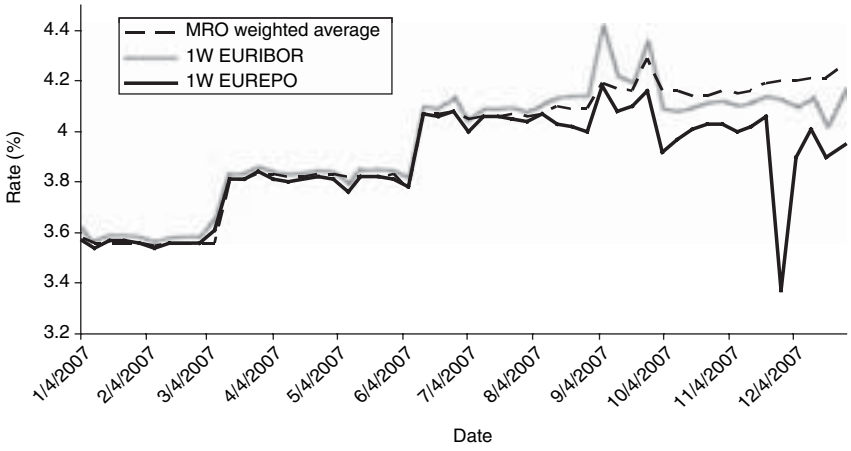

Figure 7.4. Evolution of MRO weighted average, 1 Week repo, and 1 Week unsecured interbank rates in 2007.

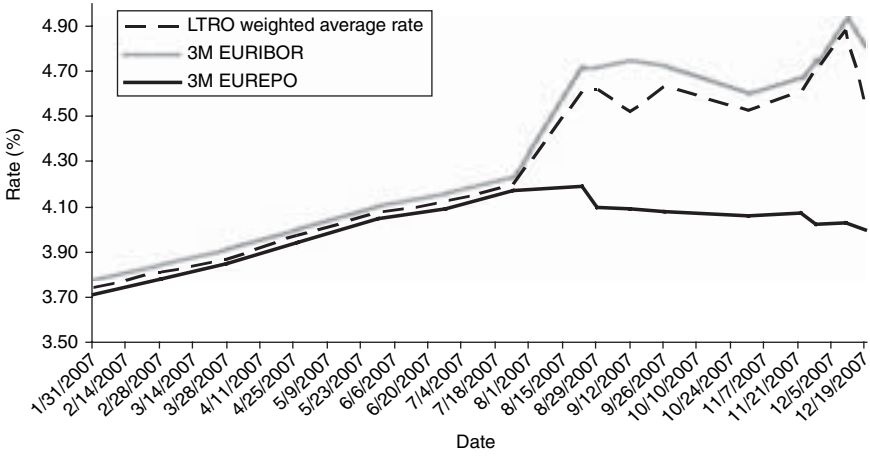

Figure 7.5. Evolution of LTRO weighted average, 3M repo, and 3M unsecured interbank rates in 2007.

Until including July, weighted average LTRO rate was very close to, albeit slightly higher (2 basis points on average) than, the EUREPO rate as seen above. The EURIBOR, on its turn, was also close to the weighted average LTRO but somewhat higher (on average 7 basis points). Since the beginning of August, as seen above, the spread between EURIBOR rate and the EUREPO has grown dramatically (on average to 64 basis points) and the weighted average LTRO has tended to follow much more closely the EURIBOR than the EUREPO, so much that its spread to the latter increased

Table 7.7 Spreads containing information on the GC and Eurosystem collateral eligibility premia – before and during the 2007 turmoil

	EURIBOR minus EUREPO		EURIBOR minus OMO[a]	
	1 week	3 months	1 week	3 months
Jan – July 2007	3	7	3	4
August – December 2007	15	67	−2	11

[a] weighted average OMO rates. *Source:* ECB.
Source: ISDA. 2006. 'ISDA Margin Survey 2006', Memorandum, Table 3.1.

to 50 basis points. This behaviour manifests, even more clearly than in the case of the one week MRO, a very aggressive bidding by commercial banks at Eurosystem operations, facilitated by the ability to use a much wider range of collateral in these operations as compared with private repurchase transactions. Indeed, it is surprising that the secured operations with the Eurosystem are conducted at rates which are closer to those of unsecured operations than to those prevailing in private secured operations.

Table 7.7 summarizes again all spread measures during the pre-turmoil and turmoil period. Overall, the second half of 2007 episode shows that indeed, eligibility premia for being acceptable collateral, either in interbank operations or for central bank operations, soar considerably in the case of financial market turmoil and implied liquidity fears.

3.4 Effects of eligibility on issuance

The preceding analysis has maintained as simplifying assumption that the amounts of securities of different types are given. However, issuance activity should react to yield effects of eligibility decisions. First, there may be a substitution effect, with debtors seeking to fund themselves in the cheapest way; thus, eligible instruments should substitute, over time, ineligible instruments. Second, agents may decide to issue, in the aggregate, more debt since the lower the financing costs, the greater the willingness to issue debt should be. While the substitution effect could, at least in theory, be significant even for an eligibility premium of only a few basis points, the second effect would require more substantial yields differentials to be relevant. Here, it suffices to note that the assumption (maintained so far) of a zero elasticity of issuance to yield changes caused by eligibility decisions biases any estimate of the eligibility premium to the upside, particularly in the long term. In the

extreme case of infinite elasticity, the only consequence of a changing eligibility premium would be on the amounts issued, not on yields.

4. Conclusions

This chapter has presented an analytical approach to the establishment of a central bank collateral framework to protect against credit losses. A collateral framework should ensure that the residual risks from credit exposures (e.g. lending) are in line with the central bank credit risk tolerance. At the same time, such a framework should remain reasonably simple. If a central bank accepts different types of collateral, it should apply differentiated risk mitigation measures, to ensure that the risk remaining *after* the application of these measures complies with its risk tolerance, whatever asset from its list of eligible collateral is used. The differentiation of risk control measures should also help avoid that counterparties provide in a very disproportionate way one particular type of collateral. This could happen, in particular, if too lax risk control measures were applied to one given asset, thus making its use as collateral too attractive, in particular if compared to standard market practice. Once the necessary risk mitigation measures have been defined for each type of asset, the central bank can rank each asset type according to its costs and benefits and then set a cut-off point which takes into account collateral demand.

The chapter stresses, however, that the collateral framework needs to strike a balance between precision and flexibility for counterparties in choosing collateral. In addition, any framework needs to maintain a degree of simplicity, which implies that it is to be seen as an approximation to a theoretically optimal design. Its actual features have to be periodically reviewed, and if necessary modified, in the light of experience, in particular in the light of the actual exposures and use of the different types of collateral and resulting concentration risks.

If the collateral framework and associated risk mitigation measures follow the above-outlined methodology, aiming at socially optimal configurations, one should not denote an effect on asset prices as *distortion*. This also implies that market neutrality is not necessarily an objective of a central bank collateral framework: effects on market equilibria are acceptable as far as they move towards optimality.

The chapter concentrates on one particular effect of the Eurosystem collateral framework on market equilibrium, namely on the 'eligibility

premium', i.e. the reduction of the yield of a given asset with respect to another asset, which is similar in all other respects but eligibility as collateral with the central bank. First, it shows how the proposed model allows to understand the origin and nature of the eligibility premium. Second, it carries out an empirical analysis to get an idea of the size of such a premium. While the size of the eligibility premium is likely to change over time, in the case of the euro area, the broad range and large amount of eligible collateral makes the eligibility premium small under normal circumstances. Some empirical measures, the limitations of which need to be stressed, consistently indicate an average level of the eligibility premium not higher than 5 basis points. However, this premium will of course be different for different assets and possibly also for different counterparties. More importantly, the eligibility premium rises with an increase in the demand for collateral, as occurring particularly in the case of a financial crisis, as illustrated by the financial market turmoil during 2007. An increase in the eligibility premium should also be observed in case the supply of available collateral were to shrink.

Independently from the conclusion reached about the complex empirical issue of the eligibility premium, there are good reasons why a central bank should accept a wider range of collateral than private market participants: First, central bank collateral serves monetary policy implementation and payment systems, the smooth functioning of which is socially valuable. While in the inter-bank market uncollateralized operations are always an alternative, central banks can, for the reasons spelled out in the introduction, only lend against collateral. A scarcity of collateral, which could particularly arise in periods of financial tensions, could have very negative consequences and needs to be avoided, even at the price of having 'too much' collateral in normal times. Second, as a consequence of the size of central bank operations, it may be efficient to set up specific handling, credit assessment or risk mitigation structures which the private sector would find more costly to set up for inter-bank operations. Finally, there is no guarantee that the market can establish efficient collateralization conventions, since the establishment of these conventions involves positive network externalities (see e.g. Katz and Shapiro 1985 for a general presentation of network externality issues). Indeed, the central bank, as a large public player, could positively influence market conventions. For instance, trade bills became the dominant financial instrument in the inter-bank market in the eighteenth, nineteenth and early twentieth century in the United Kingdom and parts of Europe (see e.g. King 1936; Reichsbank 1910)

because central banks accepted them for discounting. The last two points can be summarized by noting that the central bank is likely to have a special collateral-related 'technology' compared with private market participants, either because economies of scale or because of its ability to exploit network externalities. This in turn confirms the point that it can positively impact market equilibria, as argued above.

8 Risk mitigation measures and credit risk assessment in central bank policy operations

Fernando González and Phillipe Molitor

1. Introduction

Central banks implement monetary policy using a variety of financial instruments. These instruments include repurchase transactions, outright transactions, central bank debt certificates, foreign exchange swaps and the collection of fixed-term deposits. Out of these instruments, repurchase transactions are the most important tool used by central banks in the conduct of monetary policy. Currently the Eurosystem alone provides liquidity to the euro banking system through repurchase transactions with a total outstanding value of around half a trillion euro.

Repurchase transactions, also called 'reverse transactions' or 'repos', consist of the provision of funds against the guarantee of collateral for a limited and pre-specified period of time. The transaction can be divided into two legs, the cash and the collateral leg.

The cash leg is akin to a classical lending operation. The lender transfers an amount of cash to a borrower at the initiation of a transaction. The borrower commits to pay the cash amount lent plus a compensation (i.e. interest) back to the lender at maturity.

By the nature of lending, any lender bears credit risk, namely the risk that the borrower will fail to comply with its commitments to return the borrowed cash and/or provide the required compensation (i.e. interest) at the maturity of the transaction. Several tools are available to the lender to mitigate this risk.

First, counterparty risk can be reduced by conducting operations only with counterparties of a high credit quality, so that the probability of a default is small. In a central banking context, the set of institutions having access to monetary policy operations is generally specified with the goal of

guaranteeing equal treatment to financial institutions while also ensuring that they fulfil certain operational and prudential requirements.

Second, and reflecting the same idea, counterparty risk can also be reduced by implementing a system of limits linking the exposure to each counterparty to its credit quality, so that the potential loss is kept at low levels. For central banks, however, such a system is generally incompatible with an efficient and transparent tender procedure for allotting liquidity.

Finally, counterparty risk can be mitigated by requiring the borrower to provide adequate collateral. This approach mitigates financial risks without limiting the number of counterparties or interfering with the allotment procedure. It is a common approach chosen by major central banks when conducting repurchase operations. When combined with the appropriate risk management tools, collateralization can reduce the overall risk to negligible levels.

The collateral leg of a repurchase transaction consists, hence, of providing collateral amounting at least in value to the cash borrowed to the lender, which is returned by the borrower upon receiving back the cash lent and the compensation at maturity of the transaction.

The lender in a collateralized reverse transaction may still incur a financial loss. However, this would require more than one adverse event to occur at the same time. This could happen as follows: the borrower would first default on his obligation to the lender, resulting in the lender taking possession of the collateral. Assuming that at the time of the default the value of the collateral covered the value of the liquidity provided through the reverse transaction, financial risk could arise from the following two possible sources:

- Credit risk associated with the collateral. The issuer of the security or the debtor of the claim accepted as collateral could also default, resulting in a 'double default'. The probability of such a combination of defaults can be considered negligible if eligible assets satisfy high credit quality standards and if the lender does not accept assets issued by the borrower or entities having close financial links to the borrower.
- Market and liquidity risk. This would arise if the value of the collateral fell in the period between the counterparty's default and the realization of the collateral. In the time between the last valuation of the collateral and the realization of the collateral in the market, the collateral price could decrease to the extent that only a fraction of the claim could be recovered by the borrower. Market risk may be defined in this context as the risk of financial loss due to a fall of the market value of collateral caused by

exogenous factors. Liquidity risk may be defined as the risk of financial loss arising from difficulties in liquidating a position quickly without this having a negative impact on the price of the asset. Market and liquidity risk can also be reduced considerably by following best practices in the valuation of assets and the risk control measures applied.

The collateral leg is hence intended to mitigate the credit or default risk of the counterparty borrowing the cash and therefore plays a crucial role in this type of operations. In case of default of the counterparty, the collateral taker which in the context of this book is the central bank, can sell the collateral received and make good any loss incurred in the failed repo transaction. When the collateral received is default-risk free, as for example with government bonds, collateralization transforms credit risk (i.e. the risk of default of the counterparty) into market and liquidity risk (i.e. the risk of incurring an adverse price movement in the collateral position and the risk of impacting the price due to the liquidation of a large position over a short period of time).

Figure 8.1 provides a visual summary of a reverse transaction and the risks involved.

Two main risk factors need to be considered in the risk management of collateral underlying repo operations. First, the credit quality of the collateral needs to be sufficiently high so as to give enough reassurance that the

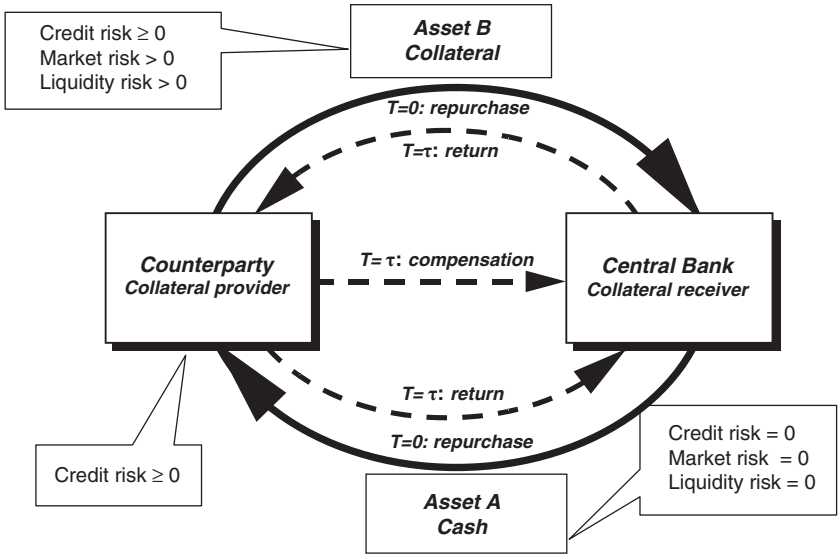

Figure 8.1 Risks involved in central bank repurchase transactions. T is a time indicator that is equal to zero at the starting date and equal to τ at the maturity date of the credit operation.

collateral would not quickly deteriorate into a state of default after the default of the counterparty. In this regard, it is also crucial that the collateral quality would be independent from that of the counterparty (i.e. no close links). To assess the credit quality of the collateral, central banks tend to rely on external ratings as issued by rating agencies or internal credit quality assessments as produced by in-house credit systems. This chapter will review the main sources of credit quality assessments used by central banks in the assessment of collateral and the main parameters that a central bank needs to define in its credit assessment framework such as the minimum credit quality threshold (e.g. a minimum rating threshold) and a performance monitoring of the credit assessment sources employed. Second, the intrinsic market risk of the collateral should be controlled. As discussed above, in case of default of the counterparty the collateral taker will sell the collateral. This sale is exposed to market risk or the risk of experiencing an adverse price movement. This chapter provides a review of different methods and practices that have been used to manage the intrinsic market risk of collateral in such repurchase or repo agreements. In general terms, such practices can rely on three main pillars: *marking to market* which helps reduce the level of risk by revaluing more or less frequently the collateral using market prices,[1] *haircuts* which help reduce the level of financial risk by reducing the collateral value by a certain percentage and *limits* which help reduce the level of collateral concentration by issuer, sector or asset class. In this chapter we consider all of these techniques in the establishment of an adequate central bank risk control framework. Given the central role of haircuts in any risk control framework, we put considerable emphasis on haircut determination.

Any risk control framework of collateral should be consistent with some basic intuitions concerning the financial asset risk that it is trying to mitigate. For example, it should support the perception that a higher haircut level should be required to cover for riskier collateral. In addition, the lower the marking-to-market frequency, the higher the haircuts need to be. Higher haircut levels should also be required in case the time to capture the assets in case of default of the counterparty or the time span needed before actual liquidation of the assets in case of default of the counterparty increases (Cossin *et al.* 2003, 9). Liquidity risk or the risk of incurring a loss in the liquidation due to illiquidity of the assets should directly impact the

[1] If the collateral value is below that of the loan and beyond a determined trigger level, the counterparty will be required to provide additional collateral. If the opposite happens the amount of collateral can be decreased.

level of haircuts. Finally, higher credit risk of the collateral received should also produce higher haircuts.

Despite the central role of collateral in current financial markets and in particular central bank monetary policy operations, little academic work exists on risk mitigation measures and risk control determination. Current industry practice is moving towards a more systematic approach in the derivation of haircuts by applying the Value-at-Risk approach to collateral risks but some reliance on ad hoc rule-based methods still persists. On the whole, despite recent advances in financial modelling of risks, the discussion among academics and practitioners on the precise framework of risk mitigation of collateral is still in its infancy (see for example Cossin *et al.* 2003; ISDA 2006). This chapter should also be seen in this light; a comprehensive and unique way of how to mitigate risk in collateralized transactions has yet to emerge. What exists now is a plethora of methods for risk control determination that are used based on context and user sophistication. The chapter reviews some of these risk mitigation determination methods of which some are used by the Eurosystem.

This chapter is organized as follows. Section 2 describes how central banks could assess first credit quality of issuers of collateral assets and the main elements of a credit assessment framework. Section 3 discusses the basic set-up of a central bank as a collateral taker in a repurchase transaction where marking-to-market policy is specified. In Section 4 we discuss various methods for haircut determination, focusing on asset classes normally used by central banks as eligible collateral (i.e. fixed-income assets), and review how to incorporate credit risk and liquidity risk in haircuts. Section 5 briefly discusses the use of limits as a risk mitigation tool for minimizing collateral concentration risks and Section 6 concludes.

2. Assessment of collateral credit quality

2.1 Scope and elements

To ensure that accepted collateral fulfils sufficient credit quality standards, central banks tend to rely on external or internal credit quality assessments. While many central banks today rely exclusively on ratings by rating agencies, there are also central bank internal credit quality assessment systems in operation. Historically, the latter was the standard. This section reviews the main credit quality assessment systems at the disposal of central

banks to assess the credit quality of collateral used in monetary policy operations. These are external credit rating agencies, in-house credit assessment systems, counterparties' internal rating systems and third-party credit scoring assessment systems.

Before any credit quality assessment is taken into account, the central bank must stipulate a minimum acceptable level of credit quality below which collateral assets would not be accepted. Typically, this minimum level or credit quality threshold is given in the form of a rating level as issued by any of the major international rating agencies. For example, the minimum threshold for credit quality could be set at a 'single A' credit rating.[2] Expressing the minimum credit quality level in the form of a letter rating is convenient because its meaning and information content is well understood by market participants. However, not all collateral assets carry a rating from one of the major rating agencies. An additional credit quality metric is needed, especially when the central bank accepts collateral issued by a wide set of entities not necessarily rated by the main rating agencies.

The probability of default (PD) over one year is such a metric. It expresses the likelihood of an issuer or debtor defaulting over a specified period of time, normally a year. Its meaning is similar to that of a rating, which takes into account the probability of default as well as other credit risk factors such as recovery in case of default. Both measures, ratings and probability of default, although not entirely equivalent, are highly correlated, especially for high levels of credit quality.

The Eurosystem Credit Assessment Framework (ECAF), which is the set of standards and procedures to define credit quality of collateral used by the Eurosystem in its monetary policy operations, uses both metrics interchangeably. In this respect, a 'translation' from ratings to probability of default levels is required (see Coppens *et al.* 2007, 12). In the case of the Eurosystem, a PD value of 0.10 per cent at a one-year horizon is considered to be equivalent to a 'single A' rating, which is the minimum level of rating accepted by the Eurosystem. These minimum levels of credit quality should be monitored and confirmed regularly by the decision-making bodies of the central bank so as to reflect the risk appetite of the institution when accepting collateral.

2.1.1 Rating agencies

The core business of public rating agencies such as Standard & Poor's, Moody's and Fitch is the analysis of credit quality of issuers of debt

[2] This means a minimum long-term rating of A- by Fitch or Standard & Poor's, or A3 by Moody's.

instruments, as regards their ability to pay back their debt to investors. These public rating agencies, or External Credit Assessment Institutions (ECAIs) as they are called in the Basel II capital requirements, usually play a key role in the credit quality assessment of any central bank. The credit assessment is summarized into different letter rating classes: Aaa to C for Moody's and AAA to C for Standard & Poor's and Fitch.

Ratings issued by rating agencies should be revised and updated at regular intervals to reflect changes in the credit quality of the rated obligor. Ratings are meant to represent a long-term view, normally trying to strike a balance between rating accuracy and rating stability. These ratings 'through the cycle' can be slow to adjust. This sometimes causes a significant mismatch between market perceptions of credit quality, which are inherently more focused on shorter time horizons, and that of rating agencies, which are more longer term. In case multiple ratings exist for a single obligor it is common prudent practice to use the second-best rating rather than the first-best available rating.

In addition to the more classical ratings, newer quantitative credit rating providers such as Moody's KMV and Kamakura have recently entered the market with ratings based on proprietary quantitative models that can be more directly interpreted as a probability of default. Contrary to the ratings of classical rating agencies, these are 'point-in-time' ratings that do not attempt to average out business cycle effects.

2.1.2 Central bank's in-house credit assessment systems

It can be valuable for central banks to develop and run internally a credit risk assessment system that caters for the different needs of a central bank in its core business of monetary policy formulation and implementation as well as (in countries where the central banking and supervisory function are allocated to a single institution) supervisory tasks. Due to their privileged institutional position, central banks might have direct access to a rich statistical data set and factual information on local obligors that permit the development of such an internal credit assessment system. As it is the case for commercial banks, central banks also tend to prefer relying on an internally developed approach for setting-up an internal credit risk assessment model rather than building on models from market providers adapted to the available dataset.

In most countries, institutional and regulatory policy considerations lead central banks to use credit assessments in order to fulfil their supervisory, regulatory or monetary policy objectives, and do not permit the disclosure

or sharing of such credit assessment information. Box 8.1 describes the historical background that triggered the set-up of in-house credit assessment systems and the subsequent development in some Eurosystem central banks. Box 8.2 introduces the in-house system implemented by the Bank of Japan.

Box 8.1. Historical background in the creation of in-house credit assessment systems in four Eurosystem central banks[3]

Deutsche Bundesbank

Prior to the launch of the European monetary union, the Deutsche Bundesbank's monetary policy instruments included a discount policy. In line with section 19 of the Bundesbank Act, the Bundesbank purchased 'fine trade bills' from credit institutions at its discount rate up to a ceiling (rediscount quota) set individually for each institution. The Bundesbank ensured that the bills submitted to it were sound by examining the solvency and financial standing of the parties to the bill. In the early seventies, the Bundesbank began to use statistical tools. In the nineties, a new credit assessment system was developed, introducing qualitative information in the standardized computer-assisted evaluation. The resulting modular credit assessment procedure builds on a discriminant analysis and a 'fuzzy' expert system.

Banque de France

The rating activities is one of the activities that originated from the intense business relations between the Banque de France and companies since its creation at the start of the nineteenth century. From the 1970s onwards, the information collection framework of Banque de France and all the resulting functions building on it were consequently developed and explain the importance of this business nowadays. The 'Companies' analysis methodology unit' and the 'Companies' Observatory unit' are both located in the directorate 'Companies' of the General Secretariat. The independence and prominence of the rating function within the Banque de France has its seeds in the multiple uses of ratings. In addition to the usage for bank refinancing purposes, credit assessments are also used for banking supervision, bank services and economic studies.

Banco de España

The Banco de España started rating private paper due to the scarcity of collateral in Spain in 1997. The scarcity of collateral in Spain was increasing as central bank Deposit Certificates were phasing out. Equities were one of the first asset classes subject to in-house assessment as local banks had equities in their portfolios, but also because of the liquidity of this type of instrument. Bank loans were added in September 2000.

[3] For information on the Deutsche Bundesbank in-house system see Deutsche Bundesbank (2006) and for information on the Banque de France see Bardos *et al.* (2004).

Box 8.1. (cont.)

Oesterreichische Nationalbank

The Oesterreichische Nationalbank (OeNB) started its credit assessment business after World War 2. The main reasons leading to the development of this business area over the years were the discount window facility, the European Recovery Program (ERP), export promotion and the development of the banking supervision activities (especially since the start of the discussions on the new capital adequacy framework around 1999). Additionally, the information serving as input to the credit assessment process is used to support economic analyses. Although historically the credit assessment function originates from the discount facility and ERP loan business, credit assessments are now mainly used for supervisory purposes.

Box 8.2. In-house credit assessments by the Bank of Japan[4]

According to its 'Guidelines on Credit Ratings of Corporations', Bank of Japan confers credit ratings to corporations, excluding financial institutions, whose head offices are located in Japan. These evaluations are made in a comprehensive manner based on quantitative analyses of the financial statements of debtors and qualitative assessments of their future profitability and the soundness of their assets.

The Bank gives the credit rating, upon request of a counterpart financial institution, taking into consideration the following factors:

(a) Quantitative factors: Mainly financial indicators of the corporation, including the net worth and stability of cash-flows.
(b) Qualitative factors: Profitability, soundness of assets, business history, position in the relevant industry, management policy, evaluation by financial institutions, information obtained through examinations of a financial institution by the Bank, and the ratings of the corporation by appropriate rating agencies, when available.

It also takes into account other information relevant for the assessment of the credit-worthiness of the corporation. The credit ratings are accorded on the basis of consolidated financial statements, when available. In principle, the credit ratings are reviewed once a year. However, the Bank can conduct irregular reviews, when judged necessary.

2.1.3 Counterparties' internal ratings based (IRB) systems

Due to the fact that credit rating agencies have traditionally concentrated on larger corporate bond issuers, the set of obligors covered by public credit rating agencies is only a fraction of all obligors that make up a counterparty's credit portfolio. Important issuer categories of obligors are small- and medium-sized enterprises. To the extent that the central bank wants to make

[4] See Bank of Japan (2004).

debt instruments issued by these types of obligors eligible, a credit assessment system needs to be in place to assess them. Following the new capital requirements as prescribed by Basel II, commercial banks can use their own IRB systems to rate their credit exposures in order to obtain the necessary risk weights for capital requirements purposes. The use of such internal models is subject to banking supervision certification following the procedures foreseen under the new capital adequacy framework (Basel II) or the EU Capital Requirements Directive (CRD)[5]. Central banks would normally reserve a right to override or adjust the rating produced by the IRB system.

Box 8.3 describes the approach followed by the Federal Reserve when accepting credit assessments issued by commercial banks to rate collateral used in the discount window facility.

Box 8.3. The Qualified Loan Review programme of the Federal Reserve[6]

The Qualified Loan Review (QLR) programme is a vehicle that allows financially sound depository institutions to pledge commercial loans as collateral to secure discount window advances, and, if applicable, Treasury Tax and Loan (TT&L) deposits. To maximize efficiency of the Reserve Bank and the pledging bank, the programme relies on the commercial bank's internal credit risk rating system to ensure that only loans of high credit quality are pledged as collateral to the discount window. Under the programme, the Reserve Bank will accept and assign collateral values to pledged loans based on the commercial bank's internal rating scale rather than an individual credit assessment by the Reserve Bank.

The discount window seeks written notification from the commercial bank's primary regulator regarding eligibility to participate or remain in the QLR programme based on team examination findings. A depository institution's qualification for the programme is contingent upon the examiner's review regarding financial strength and sophistication of the candidate's internal credit risk rating system.

To qualify for the QLR programme, an institution must submit copies of their credit administration procedures for evaluation. The internal loan review system must also prove satisfactory to the institution's primary bank supervisor in order to meet QLR qualifications. Components of an acceptable loan review system include, but are not limited to the following requirements: an independent loan review performed by qualified personnel at the institution; the internal loan review function should be independent of the lending function; the quality, effectiveness and adequacy of the loan review staff should reflect the size and

[5] The CRD comprises Directive 2006/48/EC of the European Parliament and of the Council of June 14, 2006 relating to the taking up and pursuit of the business of credit institutions (recast) (OJ L177 of June 30, 2006, page 1) and Directive 2006/49/EC of the European Parliament and of the Council of June 14, 2006 on the capital adequacy of investment firms and credit institutions (recast) (OJ L177 of June 30, 2006, page 201).

[6] See www.newyorkfed.org/banking/qualifiedloanreview.html.

Box 8.3. (cont.)

complexity of the institution; mechanisms should be in place to inform the loan review function of credit quality deterioration; the function should have the ability to follow-up with prompt corrective action when unsound conditions and practices are identified.

The frequency and scope of the internal review process should be deemed adequate by bank supervisors. Systems for the continuous surveillance of asset quality to monitor deterioration should be in place.

2.1.4 Third-party credit rating tools

Third-party credit scoring/rating tools (RTs) refer to a credit assessment source that consists of third-party applications which rate obligors using, among other relevant information, audited annual accounts (i.e. balance sheet and income statement data). Such rating tools assess the credit risk of obligors through various statistical methods. These methods aim at estimating the default probability of obligors, usually relying on accounting ratios.[7] Typically, these tools are operated by independent third-party RT providers. As it is the case with internal rating based systems of banks, they aim at filling the rating gap left by publicly recognized rating agencies.

Central banks need to make sure that the RT meets some minimum quality criteria. Typical elements of controls are the assessment of rating accuracy and methodological objectivity, coverage, availability of detailed documentation of procedures for data collection and credit assessment methodology. If the RT is run by a third-party provider outside the central bank, some minimum standards also need to be imposed. In this respect, typical control elements are the assessment on the independence of the provider, sufficient resources (i.e. economic and technical resources, know-how and an adequate number of qualified staff), credibility (i.e. a track record in the rating business) and internal governance, among other factors.

2.2 The Eurosystem Credit Assessment Framework

To ensure the Eurosystem's requirement of high credit standards for all eligible collateral, the ECB's Governing Council has established the so-called Eurosystem Credit Assessment Framework (ECAF) (see ECB 2006b, 41). The ECAF comprises the techniques and rules which establish and ensure

[7] Typical examples are working capital/total assets, EBITDA/total assets, retained earnings/total assets, etc.

Table 8.1 Summary of ECAF by credit assessment source in the context of the Single List

Rating sources	Scope by asset type	Who operates	Rating output	Supervision of credit source	Role of Eurosystem
Rating Agencies (RA)	Public sector and corporate issuers and their debt instruments. Asset-backed securities	E.g. Moody's, S&P, Fitch or any other recognized ECAI	Rating/ Probability of default (PD) of security Rating/PD of bank loan debtor	National supervisory authority/ market	Monitoring
Counterparty internal rating-based (IRB) system	Debt instruments issued by public sector and non-financial corporate issuers	IRB certified banks	Probability of default of obligor	National supervisory authority/ market	Monitoring
Central bank in-house credit-assessment systems (ICAS)	Non-financial corporate obligors	Eurosystem National Central Banks	Rating/ Probability of default of obligor	Eurosystem/ ECB	Operation of systems Monitoring, certification of eligibility
Third-party rating tool (RT)	Non-financial corporate obligors	Authorized/ eligible third party providers	Probability of default of obligor	Eurosystem/ ECB	Monitoring, supervision/ certification of eligibility

the Eurosystem's requirement of high credit standards for all eligible collateral. The ECAF makes use not only of ratings from (major) external rating agencies, but also from other credit quality assessment sources, including the in-house credit assessment systems of national central banks, the internal ratings-based systems of counterparties and third-party rating tools. Table 8.1 summarizes the key elements of the Eurosystem framework in terms of the type of credit assessment sources used, the scope of these sources as regards the asset types covered, the rating output, the operative output and the credit source supervision. Given the variety of credit assessment sources it is imperative that these systems are monitored and checked in their performance behaviour in order to maintain the principles of comparability and accuracy. Obviously, it would not be desirable that within such array of systems, one or more systems would stray away from an average performance. With this aim, the ECAF contains a performance monitoring framework.

2.2.1 Performance monitoring framework

The ECAF performance monitoring process consists of an annual ex post comparison of the observed default rate for the set of all eligible debtors (the static pool) with a credit quality better or equal than the credit quality threshold and the credit quality threshold.. It aims to ensure that the results from credit assessments are comparable across systems and sources. The monitoring process takes place one year after the date on which the static pool is defined (see Coppens *et al.* 2007).

The first element of the process is the annual compilation by the credit assessment system provider of a static pool of eligible debtors, i.e. a pool consisting of all corporate and public debtors, receiving a credit assessment from the system satisfying the following condition:

$$PD(annual\ horizon) \leq 0.10\%(benchmark\ PD)$$

All debtors fulfilling this condition at the beginning of the period constitute the static pool for this period. At the end of the foreseen twelve-month period, the realized default rate for the static pool of debtors is computed. On an annual basis, the rating system provider has to submit to the Eurosystem the number of eligible debtors contained in the static pool and the number of those debtors in the static pool that defaulted in the subsequent twelve-month period.

The realized default rate of the static pool of a credit assessment system recorded over a one-year horizon serves as input to the ECAF performance monitoring process which comprises an annual rule and a multi-period assessment. In case of a significant deviation between the observed default rate of the static pool and the credit quality threshold over an annual and/or a multi-annual period, the Eurosystem consults the rating system provider to analyse the reasons for that deviation. This procedure may result in a correction of the credit quality threshold applicable to the system in question.[8]

3. Collateral valuation: marking to market

In a monetary policy operation conducted via a repurchase transaction, there is a contract between the central bank who acts as the collateral taker

[8] The Eurosystem may decide to suspend or exclude the credit assessment system in cases where no improvement in performance is observed over a number of years. In addition, in the event of an infringement of the rules governing the ECAF, the credit assessment system will be excluded from the ECAF.

and the commercial bank (i.e. the counterparty), who borrows cash from the central bank. The central bank requires the counterparty to provide α_0 units of collateral, say a fixed-term bond (where $B(t,T)$ denotes the value of one unit of the bond at time t maturing at time T) to guarantee the cash C_0 lent at the start of the contract. The central bank detracts a certain percentage h, the haircut, from the market value of the collateral.

The time length of the repurchase transaction can be divided into K periods, where margin calls can occur K times. The central bank can introduce a trigger level for the margin call, i.e. as soon as the (haircut-adjusted) value of the collateral diverges from the underlying cash value lent by the central bank beyond this trigger level, there is a margin call to re-establish the equivalence of value between collateral and cash lent. Typically this trigger level is given in percentage terms of the underlying cash value.

At the end of each period k ($k = 1, 2, \ldots, K$) the central bank faces three possible situations:

1. The adjusted collateral value taking into account the haircut is higher than the underlying cash borrowed, i.e. $C_k < \alpha_{k-1}B(t_k,T)(1-h)$, where α_{k-1} is the amount of collateral at the beginning of period k and the collateral $B(t_k,T)$ is valued using closing market prices at the end of period k. In this situation, the counterparty could demand back some of the collateral so as to balance the relationship between cash borrowed and collateral pledged, i.e. choose α_k such that $C_k = \alpha_k B(t_k,T)(1-h)$. The repo contract continues.
2. The adjusted collateral value is below the value of the underlying cash borrowed, i.e. one has $C_k > \alpha_{k-1}B(t_k,T)(1-h)$. In this situation, a margin call happens and the counterparty will be required to deposit more collateral so as to balance the relationship, i.e. choose α_k such that $C_k = \alpha_k B(t_k,T)(1-h)$. If the counterparty does not default at the end of period k, it will post the necessary extra collateral and the contract continues.
3. In case the margin call happens and the counterparty defaults it will not be able to post the necessary extra collateral and the central bank may have a loss equal to $C_k - \alpha_{k-1}B(t_k,T)$, i.e. the difference between the cash borrowed by the counterparty and the unadjusted market value of the collateral. The contract at this stage enters into a liquidation process. If in this process the central bank realizes the collateral at a price lower than C_k, it will make a loss.

Obviously, the central bank is most interested in the third situation. Given the default of a counterparty, the central bank may be faced with a loss, especially in one of the following two situations: (a) the mark-to-market value assigned

to the collateral is far away from fair and market transacted prices for such collateral, or (b) the haircut level does not offer sufficient buffer for the expected price loss in the liquidation process. The determination of haircuts will be treated in the next section, in this section we emphasize the first aspect: without a good quality estimate for the value of the collateral, any efforts made in the correct determination of haircuts could be rendered futile. Central banks, therefore, need to pay close attention and invest sufficient resources to ensure correct valuation of the collateral received.

The valuation of marketable and liquid collateral is typically determined by current market prices. It is important for pricing sources to be independent and representative of actual transacted prices. Bid prices, if available, are generally preferred as they represent market prices at which it is expected to find buyers. If a current market price for the collateral cannot be obtained, the last trading price is sometimes used as long as this price is not too old: as a general rule, if the market price is older than five business days, or if it has not moved for at least five days, this market price is no longer deemed representative of the intrinsic fair value of the asset. Then other valuation methods need to be used. Such alternative valuation methods could for example rely on the pooling of indicative prices obtained from market dealers or on a theoretical valuation, i.e. mark-to-model valuation. Theoretical valuation is the method of choice by the Eurosystem whenever the market price is not existent or deemed to be of insufficient quality. Whatever the method chosen (market or theoretical valuation), it is accepted practice that the value of collateral should include accrued interest.

The frequency of valuation is also important. In effect, it should be apparent from the description of the three different situations that could face the central bank above, that marking to market and haircuts are close substitutes in a collateral risk control framework, albeit not perfect. In the extreme, haircuts could be lowered significantly if the frequency of marking to market were very high, with equally high frequency of collateral margin calls. This is due to the fact that the expected liquidation price loss would be small when the asset has been valued recently. On the contrary, if marking-to-market frequency is low, say once every month, the haircut level should be higher. It has to account for the higher likelihood that the price at which the collateral is marked could be far away from transacted prices when the central bank needs to liquidate. Current practice relies on daily valuation of collateral valued as of close of business. As discussed earlier, the revaluation frequency should be taken into account in the determination of haircuts treated in the next section.

4. Haircut determination methods

As a general principle, haircuts should protect against adverse market value changes before the liquidation of the collateral. The fact that different assets present different market liquidity characteristics makes it impossible to have a unique method of haircut calculation for all assets. For example, haircuts applied to non-marketable assets should reflect the risk associated with non-marketability.[9] This risk comes in the form of an opportunity cost of possibly having to hold an asset until maturity if buyers for the asset cannot be found. Marketable assets do not present this opportunity cost. Instead, the risk associated with perfectly or semi-liquid assets mainly stems from the possibility of incurring a loss if the value of the collateral decreases due to an adverse market move before liquidation. Differences in asset characteristics therefore imply a haircut determination methodology that takes those characteristics into account. It has to be estimated either statistically or dynamically what the collateral would be worth if the collateral taker ever had to sell it.

The generally most accepted methodology for calculating haircuts is based on the Value at Risk (VaR) concept. VaR can be calculated on a value amount or percentage basis and is interpreted as the value amount or percentage loss in value that will be equalled or exceeded on ξ per cent of the time.[10] It depends on the value chosen for ξ as well as on the time horizon considered. This risk measure is a basic indicator of the price volatility for any debt or equity instrument. It is the first building block of any haircut calculation. Figure 8.2 illustrates the main components of a haircut calculation.[11]

Additional VaR adjusts the basic VaR to account for sources of risk other than just pure market risk. These are specific risks that affect the value of the collateral. For example, they could comprise the extra risk due to lower rated collateral or lower liquidity characteristics. The additional types of VaR will usually require some more specific instrument type analysis.

[9] This concept of non-marketability refers to tradable assets that do not enjoy a market structure that supports their trading. A typical non-marketable asset would be a bilateral bank loan. Bilateral bank loans can be traded or exchanged on an over-the-counter basis. The cost of opportunity risk is equal to the difference between the yield to maturity on the collateral and the yield that would have been realized on the roll-over of monetary policy operations until the maturity date of the collateral.

[10] For example, VaR (5%) is the loss in value that will be equalled or exceeded only 5 per cent of the time.

[11] See also ISDA (2006) for a similar exposition of variables.

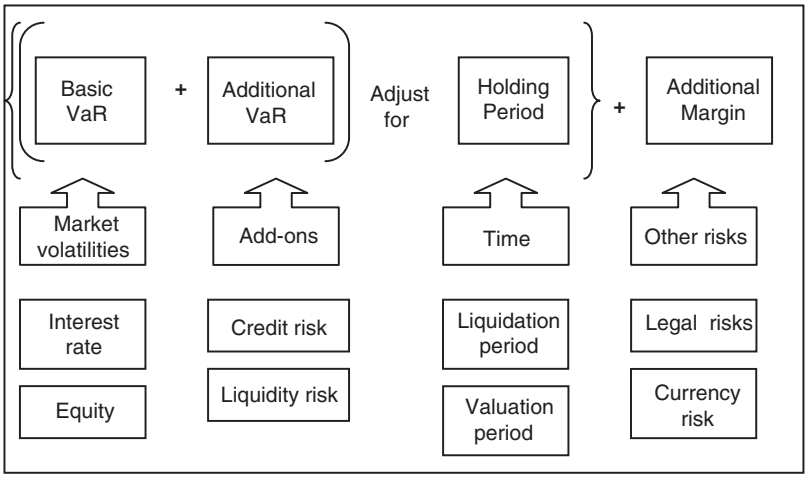

Figure 8.2 Basic determinants of haircut calculations.

The holding period should cover the maximum time period that is estimated possible between the last establishment of the correct amount of collateral and actually being able to liquidate the collateral in case of default. This is depicted in Figure 8.3. The holding period consists of the so-called 'valuation period', 'grace period' and 'actual realization time'. The length of the holding period is therefore based on assumptions regarding these three components. The valuation period relates to the valuation frequency of the collateral. In a daily valuation framework, if the default event time is t, it is assumed that the valuation occurred at $t-1$ (i.e. prices refer to the closing price obtained on the day before the default event time) and is common to all collateral types.

The grace period time is the time allocated to find out whether the counterparty has really defaulted or merely has operational problems to meet its financial obligations.[12] The grace period may also encompass the time necessary for decision makers to take the decision to capture the collateral and the time necessary for legal services to analyse the legal implications of such a capture. When the grace period has elapsed and it is clear that the counterparty has defaulted and the collateral is captured, the collateral is normally sold in the market immediately.

[12] The repo agreement specifies the type of default events that could trigger the capturing of the collateral. Among those events are the failure to comply with a daily margin call or the more formal bankruptcy proceeding that a counterparty may initiate to protect its assets. However, the triggering event may be due to operational problems in the collateral management system of the counterparty and not because of a real default which provides some degree of uncertainty in the 'capture' of the collateral guaranteeing the repo operation. Following the master repurchase agreement, the central bank issues a 'default notice' to the counterparty in case of a default event, in which three business days are given to the counterparty to rectify the event of default.

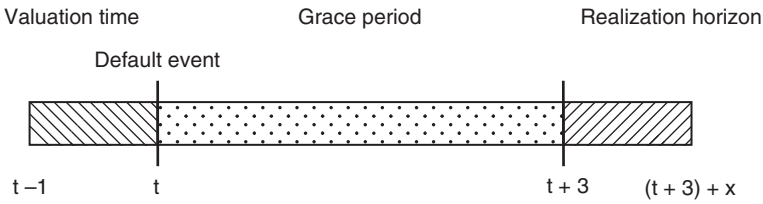

Figure 8.3 Holding period.

The realization horizon refers to the time necessary to orderly liquidate the asset. Normally the collateral would be sold immediately after default. This would cause a market impact. To reduce the market impact of a large sale, the collateral taker would need to sell over a longer period (i.e. x days) in smaller quantities. This extra time to dispose the assets leads to extra market risk and needs to be considered in the haircut calculation.

It is assumed that the market risk encountered over the longer realization horizon can be used as a proxy for the endogenous liquidity risk associated with the sale of a large position. Thus the realization horizon provides a measure of the impact on liquidity due to an immediate sale of a large position.

Traditional holding periods range from one week to one month. If the holding period is one month and one month volatilities are used no conversion is needed. However, if the holding period and the volatility estimation refer to different periods, there is a need to adjust the volatility by the square root of time.[13]

Additional margins could be added to the resulting haircut to cover for non-market related risks, such as for example legal or operational risks. There could be concerns about how quickly collateral could be captured after default due to legal or operational uncertainties. In addition to non-market related risks, cross-currency haircuts are often added when there is a mismatch between the currency of the exposure and the currency in which the collateral is denominated. For example, in local central bank repurchase operations collateral could be denominated in foreign currency.[14] Additional cross currency margins would typically be based on VaR calculations. They also need to be adjusted for the holding period.

[13] If for example, volatility is calculated on an annual basis, then the one month volatility is approximately equal to the annual volatility times the square root of 1/12.

[14] This is sometimes the case in so-called 'emergency collateral arrangements' between central banks in which foreign assets are allowed as eligible collateral in cases of emergency situations in which access to domestic collateral is not available or collateral is scarce due to a major disruptive event.

4.1 Basic VaR-related haircuts

The basic haircut estimation is built on the concept of VaR with a given confidence level in which the holding period (or the time necessary to orderly execute a sale) plays the key role. Let's for example take the typical case of a central bank receiving a debt instrument as collateral. One can assume that the asset price change v of a debt instrument i can be approximated by applying the following expression:

$$\nu_{i,t+\tau} = -D\Delta y \tag{8.1}$$

where D denotes the Macaulay Duration and y denotes the yield to maturity.[15] The changes of the yield to maturity are assumed to be approximately following a normal distribution. The VaR, which is the ξ percentile (ξ can be chosen, for example $\xi = 1$ per cent) of the distribution of ν, is related in the following manner to the standard deviation of changes of the yield to maturity:

$$h_i = QD\sigma_t y \tag{8.2}$$

where Q is a factor associated to the ξth percentile of the standard normal distribution and σ is the standard deviation associated to changes in the yield to maturity. In our analysis we choose the significance level ξ to be 1 per cent which implies that Q is equal to 2.33. Haircuts are given in percentage points, so to translate back into prices we define the 1 per cent worst price P' as

$$P' = P(1 - h_i) \tag{8.3}$$

The probability that the value falls below this price is only 1 per cent.
The holding period enters the expression through the time reference used to compute the standard deviation of changes in yield to maturity (e.g. one day, one week or ten days). If the volatility estimate for changes in yields is given in annual terms and the time to liquidation is one week, the volatility estimate would have to be divided by the square root of fifty-two (since there are fifty-two weeks in a year). In general, the standard deviation of changes in yield over the time required to liquidate would be given by the

[15] The Macaulay duration is a simplification of the total price volatility of the asset due to changes in interest rates. The fact that the required time to sell the collateral is usually not very long makes this assumption appropriate. With longer time horizons Macaulay duration distorts the results.

following expression ('*holding period*' expressed in years, e.g. 1/52 if the holding period is one week):

$$\sigma_{holding\ period} = \sigma_{annual} / \sqrt{holding\ period} \tag{8.4}$$

Let us look for example at a repo operation backed by a government bond with five-year maturity. The bond pays two semi-annual coupons at 4 per cent and the current market yield is 4.39 per cent, so the bond is quoted at a discount. The duration of this bond is 4.58 and the yield annual volatility is 31.70 per cent. To calculate the basic VaR haircut for this risk-free five-year bond, we need to translate the annual volatility into holding period adjusted volatility. We assume a one-week holding period, so the annual volatility needs to be divided by the square root of fifty-two, the number of weeks in a year, which yields an adjusted volatility of 4.4 per cent. We have now estimated all input values needed to compute the basic VaR haircut in (8.2), the duration, the yield to maturity of the bond, and the volatility adjusted by the holding period. Assuming a 1 per cent significance level, we can proceed to compute the haircut for this bond, which turns out to be equal to 2.06 per cent. If instead of a semi-annual coupon bond the collateral was a five-year zero coupon bond, the haircut would be higher since the duration of the zero coupon bond equals the maturity of the bond, i.e. five years. So for a zero coupon bond with five-year maturity the haircut would increase to 2.25 per cent.

In case the central bank would accept equities to back the repo operation the calculation of the basic VaR haircut would be simpler than in the case of risk-free bonds, as we do not need to translate changes in yields into bond prices. Assuming that equity returns follow a standard normal distribution, the basic VaR haircut h is given by

$$h_i = Q\sigma_t \tag{8.5}$$

where Q is a factor associated to the ξ percentile of the standard normal distribution and σ is the standard deviation associated to equity returns. Assuming a 1 per cent significance level ξ, the resulting VaR haircut estimate would be the percentage loss in equity value that would only be exceeded 1 per cent of the time. As with bonds, we can translate from haircuts that are given in percentage points into prices by using equation (8.3), i.e. $P' = P(1 - h)$.

Let us assume that the annual standard deviation of an equity stock that has been pledged as collateral is 35 per cent. With a one-week holding

period, and a 1 per cent significance level, the basic VaR haircut estimate will be equal to 11.31 per cent.

4.2 Liquidity risk adjusted haircuts

The central banks in the conduct of its activities in the collateral framework (e.g. sale of collateral after default of a counterparty) would be generally affected by market liquidity and not by specific asset liquidity.

A liquid market is defined as one in which trading is immediate, and where large trades have little impact on current and subsequent prices or bid–ask spreads.[16,17] It refers to what in the jargon is called the 'endogenous liquidity risk', the loss due to a large sale in a given liquidation time period. Endogenous liquidity risk is mainly driven by the size of the position: the larger the size, the greater the endogenous illiquidity. A good way to understand the implications of the position size is to consider the relationship between the liquidation price and the total position size held. This relationship is depicted in Figure 8.4.

If the market order to buy/sell is smaller than the volume available in the market at the quote, then the order transacts at the quote. In this case the market impact cost, defined as the cost of immediate execution, will be half of the bid–ask spread. In this framework, such a position only possesses *exogenous* liquidity risk and no *endogenous* risk.[18] Conversely, if the size of the order exceeds the quote depth, the cost of market impact will be higher than half-spread. In such a situation the difference between the market impact and half the spread is the endogenous liquidity risk.

4.2.1 Exogenous liquidity risk

When it comes to haircut determination taking into account liquidity risks one need to be aware of the relevant type of liquidity risk, whether exogenous, endogenous or both, to be measured in the haircut calculation.

[16] Market liquidity is distinct from the monetary or aggregate liquidity definition used in the conduct of the central bank's monetary policy.

[17] Market liquidity can be defined over four dimensions: Immediacy, depth, width and resiliency. Immediacy refers to the speed with which a trade of a given size at a given cost is completed. Depth refers to the maximal size of a trade for any given bid–ask spread. Width refers to the costs of providing liquidity (i.e. bid–ask spreads). Resiliency refers to how quickly prices revert to original (or more 'fundamental) levels after a large transaction. The various dimensions of liquidity interact with each other (e.g. for a given (immediate) trade, width will generally increase with size or for a given bid–ask spread, all transaction under a given size can be executed (immediately) without price or spread movement).

[18] Exogenous illiquidity is the result of market characteristics; it is common to all market players and unaffected by the actions of any one participant.

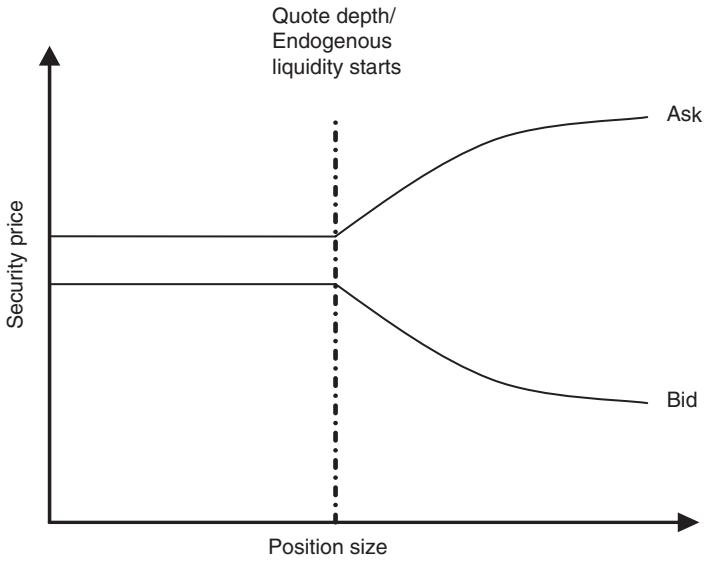

Figure 8.4 Relationship between position size and liquidation value.

We start with addressing exogenous liquidity risk. We assume a situation where the market offers high liquidity, with sufficient depth at both bid and ask quotes. Then the simplest way to incorporate exogenous liquidity risk into an adjusted VaR-based haircut is in terms of the bid–ask spread that is assumed to be constant. The liquidity risk adjusted haircut would incorporate a liquidity cost LC in addition to the basic VaR haircut:

$$LC = \frac{1}{2}\ relative\ spread \tag{8.6}$$

where the *relative spread* is equal to the actual spread divided by the midpoint of the spread. The liquidity adjusted haircut, lh, would then be equal to the basic VaR-calculated haircut, h, as presented above plus the liquidity cost, LC:

$$lh = h + LC \tag{8.7}$$

Assume for example the case of an equity haircut as in (8.5) calculated over a one-week holding period and a significance level ξ equal to 1 per cent with an estimated annual volatility of 25 per cent and a spread of 0.20 per

cent. The ratio of the liquidity adjusted haircut lh to the basic VaR based haircut h is

$$\frac{lh}{h} = 1 + \frac{relative\ spread}{2(Q\sigma)} = 1 + \frac{0.02}{2(2.33 \times 0.25/\sqrt{52})} = 1.049$$

The constant spread liquidity adjustment increases the basic haircut by approximately 5 per cent. Following this example, it is easy to show that the liquidity adjustment (a) increases with the spread, (b) decreases as the significance level decreases and (c) decreases as the holding period increases. Out of these three results, the first and the third correspond to what would be expected but the second result is not.

This approach is easy to implement and requires few inputs, but the assumption of a constant spread is highly improbable and it takes no regard for any other liquidity factors. A more plausible approach is to assume that the spreads are randomly distributed as suggested by Bangia *et al.* (1999). Assume for example that the bid–ask spread is normally distributed:

$$spread \approx N(\mu_{spread}, \sigma^2_{spread}) \tag{8.8}$$

where μ_{spread} is the mean of the spread and σ^2_{spread} is the spread volatility. The use of the normal distribution is entirely discretionary. Alternative distributional assumptions could be used, for example heavy-tailed distributions to take into account the well known feature of excess kurtosis in the spread. The liquidity cost LC will then be given by

$$LC = \frac{1}{2}(\mu_{spread} + k\sigma_{spread}) \tag{8.9}$$

where k is a parameter to be determined by for example Monte Carlo simulation. Bangia *et al.* (1999) suggest that $k = 3$ is a reasonable assumption as it reflects the empirical fact that spreads show excess kurtosis. The liquidity adjusted haircut lh would then be calculated as in (8.7), but with the liquidation cost now defined as in (8.9).

4.2.2 Endogenous liquidity risk

The previous two approaches assume that prices are exogenous and therefore ignore the possibility of the market price responding to the trading of the collateral by the central bank. In most situations, this is unreasonable, in particular in situations in which the central bank is forced to liquidate a

large amount of collateral, possibly from one single issue. In those cases, the liquidity adjustment of basic haircuts needs to take into account endogenous liquidity risk considerations rather than just exogenous ones as in the last two approaches.

Some models have been proposed for modelling endogenous liquidity risk by Jarrow and Subramanian (1997), Bertsimas and Lo (1998) and Almgren and Chriss (1999). These approaches, however, typically rely on models whose key parameters are unknown and extremely difficult to gauge due to a lack of available data. For example, in Jarrow and Subramanian an optimal liquidation of an investment portfolio over a fixed horizon is analysed. They characterize the costs and benefits of block sale vs. slow liquidation and propose a liquidity adjustment to the standard VaR measure. The adjustment, however, requires knowledge of the relationship between the trade size and both the quantity discount and the execution lag. Normally, there is no available data source for quantifying those relationships, so one is forced to rely on subjective estimates.

In the framework presented in this chapter, a more practical approach is proposed to estimate the endogenous liquidity risk. The approach is based on the definition of the *relevant liquidation horizon*, which is the expected average liquidation time needed to liquidate the position without depressing the market price.

To calculate the required (endogenous) liquidity risk-adjusted haircuts, it is easiest to group the varied collateral assets that are eligible by the central bank into collateral groups. For example, the Eurosystem classifies the eligible collateral pool into nine groups: sovereign government debt, local and regional government debt, Jumbo covered bonds, traditional covered bonds, supranational debt, agency debt, bank bonds, corporate bonds and asset backed debt (ECB 2006b). This type of classification streamlines the haircut schedule since haircuts are calculated for broad collateral groups instead of individual assets.[19]

Once all assets eligible to be used as collateral are classified into homogenous groups, the liquidity risk indicators that would define the liquidity risk profile of each of these groups have to be identified. These liquidity indicators are then combined into a so-called 'liquidity risk score card table' which is ultimately the piece of information needed to assign a *liquidation horizon* to each of the collateral groups. The higher the liquidity risk of a

[19] In the case of the ECB with over 25,000 eligible securities that can be used as collateral in its monetary policy operations, the grouping of collateral into few broad groups greatly facilitates the calculation of haircuts.

collateral group based on the indicators, the lower the market liquidity quality of the group. Therefore, a higher liquidation horizon is required to perform a sale without depressing the market price. As discussed earlier, higher liquidation horizons mean higher haircut levels. The Eurosystem currently uses a risk control system for its eligible collateral based on this strategy.

The choice of liquidity risk indicators depends on the level of depth and sophistication that the collateral taker would like to have in the measurement of liquidity risk. In the case of the Eurosystem, three variables have been identified as relevant proxies of liquidity risk: (a) yield-curve differentials, (b) average issue size and (c) bid–ask spreads. All of these measures provide a statement on exogenous liquidity risk.

A crucial assumption in the application of the strategy is that the (exogenous) liquidity risk priced either by the yield-curve differential, the average issue size or the bid–ask spread is a good proxy for (endogenous) liquidity risk. In other words, the ranking obtained by analysing the exogenous liquidity risk of collateral groups would be equal to the ranking that one would obtain by looking at endogenous liquidity risk.

The three above-mentioned liquidity risk proxies will now be discussed one by one.

4.2.3 Yield-curve differentials

Since investors do not in general buy and hold assets until maturity, less liquid assets will trade at a discount because buyers require a compensation for the loss of degrees of freedom about the timing of a possible future resale of that asset.[20] For fixed-income securities, Amihud and Mendelson (1991) formalized this concept suggesting that (exogenous) liquidity risk can be seen as the discounted value of future transaction costs incurred by future owners of the asset. Hence the current price of an illiquid asset can be calculated as the price of a comparable liquid asset minus the net present value (NPV) of future transaction costs. Within this frame of thought Amihud and Mendelson suggest to measure liquidity by yield differentials between liquid and illiquid bonds of the same credit quality. Naturally, government bonds would be the benchmark of liquidity as they represent the most liquid asset class.[21] The approach for liquidity measurement is

[20] The investors who use buy-and-hold strategies can profit from this and obtain additional yield pickup if they overrepresent illiquid bonds in their portfolios.

[21] In the Amihud and Mendelson (1991) paper a comparison is made between U.S. bills and notes having identical maturities.

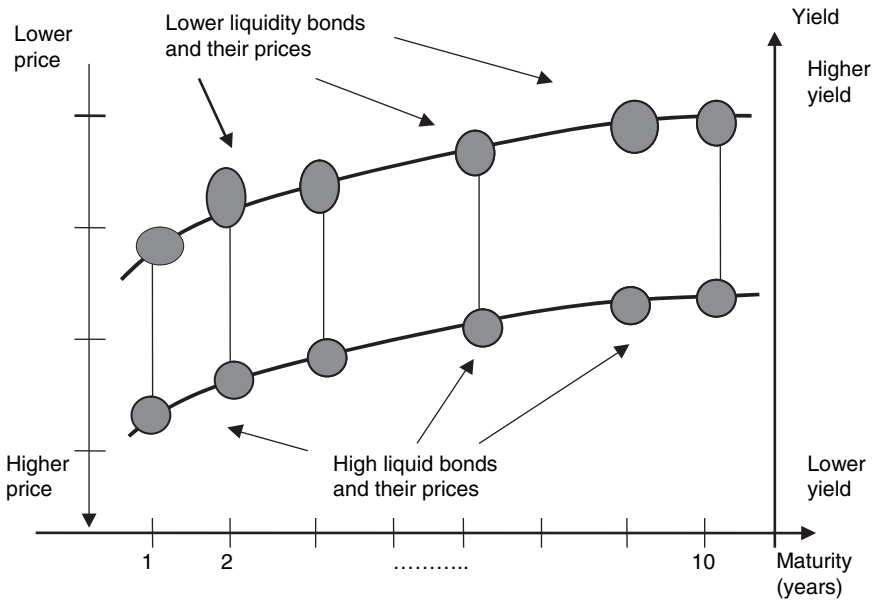

Figure 8.5 Yield-curve differentials.

then based on the difference in spread between the benchmark yield curves and the market segment yield curves with the same credit quality. The benchmark yield curve represents the market segment with lowest liquidity risk within each credit quality category.

Figure 8.5 illustrates this methodology.

Two distinct types of bonds are plotted (for illustrative purposes): highly liquid bonds selling at a relative high price and low liquidity bonds selling at a relative low price (note that the price axis in the figure is inverted). Pricing errors occur because not all bonds sell at prices that match the implied yield curve. These errors are illustrated in the figure as the differences between the solid lines drawn and the individual points in the bond price scatters. The solid lines represent the estimated (implied) yield curves valid for each of the two groups of bonds. One curve is located around low yields and corresponds to highly liquid bonds and another high-yield curve corresponding to low liquidity bonds. The area between these two curves is the liquidity measure used to rank the different collateral groups.

It is important that a 'clean' measure of liquidity risk is obtained, i.e. that the credit-risk component of the yield differentials between collateral groups is filtered out of the results. This is done by constructing benchmark curves defined on the basis of credit rating and subsequently measuring

liquidity risk for each credit grade separately, within each collateral group. The area between the estimated yield curves for each segment is used as the quantitative measure of liquidity risk. In effect, for each group several liquidity-risk indicators are calculated (e.g. one for each credit rating, AAA, AA, A, ...).

The credit-risk adjusted yield differential liquidity indicator L is obtained in the following way:

$$L_{c,s} = \int_a^b [y_{c,s}(\tau) - y_B(\tau)]d\tau \tag{8.10}$$

where, $c = [AAA, AA, A, \ldots]$ refers to the credit rating, s refers to the particular market segment being analysed (e.g. local and regional debt, supranational, bank bonds, Pfandbriefe, ...), $B = [AAA, AA, A, \ldots]$ refers to the relevant benchmark curve, $[a, b]$ are the limits for which the integral is calculated, $y(\tau)$ is the yield curve.[22] In order to obtain one yield differential liquidity indicator for each group, a volume-weighted average can be calculated, where the w's are intra-market volume weights:[23]

$$L_s = w_1 L_{c=AAA,s} + w_2 L_{c=AA,s} + w_3 L_{c=A,s} + \cdots \tag{8.11}$$

4.2.4 Effective supply and average issue size

The overall size of the market has a positive effect on market liquidity: the higher the total outstanding amount of securities traded in the market, the higher the liquidity in that market. Going deeper into the examination of market size as a liquidity proxy variable, the maturity distribution across the yield curve provides an additional element for liquidity assessment. The number of original maturities employed by the issuers of the securities in a market segment relates to the degree of fragmentation and hence to the liquidity of the market. On the one hand, a large number of original maturities fragment the market, because several securities with different coupon rates and the same remaining maturity would coexist. On the other hand, investors may not be able to find on-the-run securities to fit their

[22] For example, to parameterize the yield curve needed to calculate the yield spreads a three-factor model suggested by Nelson and Siegel (1987) could be used.

[23] Liquidity scores for the defined collateral groups are calculated using numerical integration for maturities between one and ten years.

needs, if too few original maturities are available. In order to keep a good balance, markets usually range from five to twelve original maturities with an even distribution of outstanding volume across different maturities.

In addition to total outstanding volume and its balanced distribution across different maturity buckets, average issue size is important. In general, liquid bonds are mostly large issues. Average issue size would provide a measure of issue size. It is also a measure of market fragmentation and therefore indicative of market liquidity.[24] In general, liquid markets are those that commit to issue large issues, at a regular and transparent issuance calendar across the main maturity buckets (say 2, 5, 10, and 20 or 30 years).[25]

4.2.5 Bid–ask spread

The bid–ask spread is a reflection of the level of trading intensity and therefore a good proxy for liquidity risk in a broad sense. Since inventory-control or rebalancing risks diminish as trading intensity increases, so does the inventory-control component of the spread. The spread not only reflects trading intensity, but other factors such as adverse selection, transparency regimes, asset price volatility, dealer competition as well as other factors influencing market making costs.[26]

Once the relevant indicators are computed, they are presented in a liquidity risk score card as in Table 8.2. The objective of the score card is to facilitate a ranking of the different collateral groups in terms of their liquidity profile. The ultimate ranking may be based also on qualitative criteria (for example institutional elements).[27]

[24] A related and complementary measure to average issue size is the frequency of new issues. For a given amount of overall issuance, the average issue size and frequency of new issues will be negatively correlated. On the one hand, when issue frequency is low, i.e. particular issues remain on-the-run for a long time, the average issue size is larger and the degree of fragmentation is low. However, prices of on-the-run issues tend to deviate from par value, which some investors may not like. On the other hand, when issue frequency is high, prices of on-the-run issues are close to the par value. However, the average issue size is smaller thus the degree of market fragmentation is higher.

[25] Other sources of market fragmentation affecting market liquidity are the possibility of reopening issues, the difference between on-the-run and off-the-run issues, the profile of products (e.g. strips, hybrids, . . .), the profile of holders (e.g. buy and hold, non-resident, . . .) and the institutional framework (e.g. tax conditions, accounting treatments). These factors may provide an additional qualitative assessment if needed.

[26] Bid–ask spread is seen as a superior proxy for liquidity compared to turnover ratio (or volume traded) as the latter only reflects trading intensity and the former comprises trading intensity and other factors.

[27] These factors should include considerations on the operational problems that may be encountered in the eventual implementation and communication strategy to the banking and issuer communities on the final classification decision. In this regard, it would be advantageous, for example, to consider liquidity groups that are homogeneous not only in their liquidity but also in their institutional characteristics.

Table 8.2. Liquidity score card

Collateral group	Yield differential liquidity score	Bid-ask spread in bps.	Avg. issue size (EUR million)
Government debt	0	5	2674
Jumbo Pfandbriefe	2.13	6	1205
Local & regional	1.35	9	150
Supranationals	2.70	20	271
Pfandbriefe	2.44	n.a.	88
Bank bonds	2.66	18	168
Corporate bonds	4.27	36	337

Source: European Central Bank. 2003. 'Liquidity risk in the collateral framework', internal mimeo.

4.2.6 Defining liquidity categories

Once the collateral groups are ranked by examining the different liquidity indicators and possibly other criteria of more qualitative nature, they are mapped into liquidity classes or categories. For practical reasons, the number of liquidity groups should be low but still sufficient to guarantee a certain level of homogeneity within liquidity groups. If the number of liquidity groups were too small, there would be a risk of lumping together assets with unequal liquidity profiles. If, on the contrary, the number of classes were high, the risk control framework might become unmanageable.

The analysis of the empirical quantitative results, qualitative considerations and the trade-off between homogeneity of liquidity categories and complexity of the framework lead to a decision on the optimal number of liquidity categories. In the case of the Eurosystem collateral framework, it was decided on four categories. The general content of these four liquidity categories is expressed in the following manner:

- **Category I**: Assets with outstanding liquidity. The assets present unequivocal and unambiguous top liquidity characteristics. Assets in this category would score the highest marks in the three different liquidity measurement tools. From an institutional point of view, the assets would in general be issued by sovereigns.
- **Category II**: Assets with good liquidity. These are assets that rank second to category I assets in the three liquidity measurement methods. The assets are normally issued by public or semi-public entities or have institutional features that confer them very high quality.

Table 8.3 Eurosystem liquidity categories for marketable assets

Category I	Category II	Category III	Category IV
Central government debt instruments	Local and regional government debt instruments	Traditional covered bank bonds	Asset-backed securities
Debt instruments issued by central banks	Jumbo covered bank bonds Agency debt instruments	Credit institution debt instruments Debt instruments issued by corporate and other issuers	
	Supranational debt instruments		

Source: ECB (2006b).

- **Category III**: Assets with average liquidity. These are assets that rank third to categories I and II in the liquidity measurement methods. The assets are normally issued by private entities.
- **Category IV**: Asset with below-average liquidity. Assets included in this category would represent a marginal share of the total outstanding amount of eligible assets. These are assets normally issued by private entities.

The classification of collateral assets into the different liquidity risk categories proposed in Table 8.3 does not lend itself to mechanistic translation into a haircut level. The classification is a reflection of relative liquidity and not of absolute liquidity. Therefore, some assumptions are necessary to map the assets in the different liquidity categories with haircut levels that incorporate both market and liquidity risk.[28]

The haircut determination model applies different assumptions on the liquidation horizon or the holding period depending on the liquidity category considered. The market impact is higher for those assets classified in the lower quality liquidity categories. In order to have a similar market impact across liquidity categories, those assets in the lower liquidity categories require more time for an orderly liquidation. Such an expanded sale period originates extra market risk. It is assumed that this extra market risk would proxy the liquidity risk that would be experienced if the sale were

[28] Credit risk is not accounted for in the haircut level. The haircut levels aim at protecting against an adverse market move and market impact due to a large sale. It is assumed for explanatory purposes that eligible assets enjoy high credit quality standards and that therefore credit risk considerations can be disregarded in the calculation of haircuts. Section 5 presents a method for haircut calculation when the collateral asset presents a non-negligible amount of credit risk.

Table 8.4. Eurosystem levels of valuation haircuts applied to eligible marketable assets in relation to fixed coupon and zero-coupon instruments (percentages)

	Liquidity categories							
Residual maturity (years)	Category I		Category II		Category III		Category IV	
	fixed coupon	zero coupon	fixed coupon	zero coupon	fixed coupon	zero coupon	fixed coupon	zero coupon
0–1	0.5	0.5	1	1	1.5	1.5	2	2
1–3	1.5	1.5	2.5	2.5	3	3	3.5	3.5
3–5	2.5	3	3.5	4	4.5	5	5.5	6
5–7	3	3.5	4.5	5	5.5	6	6.5	7
7–10	4	4.5	5.5	6.5	6.5	8	8	10
>10	5.5	8.5	7.5	12	9	15	12	18

Source: ECB (2006b).

done immediately. This time measure is the key parameter to feed the level of haircut to be applied.

The actual liquidation horizon varies depending on the liquidity category in which the asset is classified. For example, it can be assumed that category I assets require 1–2 trading days, category II assets 3–5 trading days, category III assets 7–10 trading days and category IV assets 15–20 trading days for liquidation. The assumed liquidation horizon needs to be added to the grace period to come up with the total holding period as depicted in Figure 8.2. The holding period is then used in the calculation of a haircut level as in equation (8.2). In this manner, the total holding period required can be assumed to be approximately equal to five days for category I, ten days for category II, fifteen days for category III and twenty days for category IV.

With this holding period information and an assumption on volatilities for the different collateral classes, it is possible to compute haircut levels. Table 8.4 presents the Eurosystem haircut schedule for fixed-income eligible collateral following the assumptions on liquidation horizons that were described earlier for each of the four different liquidity categories identified.

4.3 Credit risk-adjusted haircuts

This section illustrates a method for incorporating credit risk in the level of haircuts for a single bond. The basic VaR haircut that accounts for market risk is supplemented by an additional haircut accounting for credit risk.

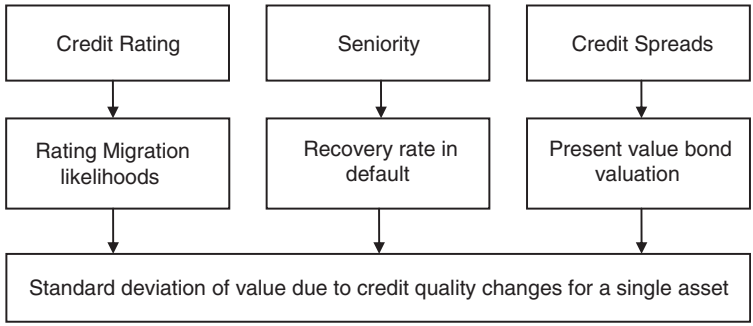

Figure 8.6 Value-at-Risk due to credit risk for a single exposure.

This additional haircut for credit risk can be estimated using the CreditMetrics methodology for calculating the credit risk for a stand-alone exposure (Gupton *et al.* 1997).

Credit risk implies a potential loss in value due to both the likelihood of default and the likelihood for possible credit quality migrations. The CreditMetrics methodology estimates the volatility of asset value due to both events, i.e. default and credit quality migration. This volatility estimate is then used to calculate a VaR due to credit risk. The Value-at-Risk methodology due to credit risk can be summarized as in Figure 8.6.

In essence, there are three steps to calculating the credit risk associated to a bond. The first step starts with assigning the senior unsecured bond's issuer to a particular credit rating. Credit events are then defined by rating migrations which include default, though a matrix of migration probabilities. The second step determines the seniority of the bond which in turn determines its recovery rate in the case of default. The forward zero curve for each credit rating category determines the value of the bond upon up/downgrade. In the third step the migration probabilities of step 1 and the values obtained for the bond in step 2 are then combined to estimate the volatility due to credit quality changes.

This process is illustrated in Table 8.5. We assume a five-year bond or credit instrument with an initial rating of single A. Over the horizon, which is assumed here to be one year, the rating can jump to seven new values, including default. For each rating, the value of the instrument is recomputed using the forward zero curves by credit rating category. For example, the bond value increases to 108.41 if the rating migrates to AAA, or to the recovery value of 50 in case of default. Given the state probabilities and associated values, we can compute an expected bond value of 107.71 and a standard deviation of 1.36.

Table 8.5 The distribution of bond values of an A rated bond

		Probability (p_i)	F. Bond Value V_i	Mean, μ $\Sigma(p_iV_i)$	Var.σ^2 $\Sigma p_i(V_i-\mu)^2$
	AAA	0.08%	108.41	0.09	0.00
	AA	2.42%	108.05	2.61	0.00
A	A	91.30%	107.79	98.42	0.01
	BBB	5.23%	107.37	5.62	0.01
	BB	0.68%	104.34	0.71	0.08
	B	0.23%	101.90	0.23	0.08
	C	0.01%	98.96	0.01	0.01
	Default	0.05%	50.00	0.03	1.67
				Variance (σ^2)	1.84
		Mean (μ)	107.71	Std. deviation	1.36

Given an expected bond value and a standard deviation, and assuming a normal distribution in asset returns, we would be able to compute a haircut based on credit risk changes.[29]

A critical input in the calculation of bond values in Table 8.5 is given by the credit transition matrix that provides an estimate not only of the likelihood of default but also of the chance of migrating to any possible credit quality step at the risk horizon. Two main paradigms underline the estimation of credit transition matrices: the 'through-the-cycle' approach representative of rating transitions as observed by the rating actions of major international rating agencies and the 'point-in-time' approach obtained from rating changes produced by Merton-type models such as Moody's KMV. An example of typical credit transition matrices of both paradigms is given in Tables 8.6 and 8.7. Notice that the approach followed by rating agencies is designed to give less variability in ratings if economic conditions change (Catarineu-Rabell *et al.* 2003) whereas the approach followed by for example Merton-based models present rating migrations that are more volatile. The volatility in rating migrations has an important bearing in the final estimated volatility of the bond value due to credit rating changes.

Another important element in the calculation of credit risk-related haircuts refers to the uncertainty associated to the recovery rate in default.

[29] Standard deviation is one credit risk measure. Percentile levels can be used alternatively to obtain the risk measure. Assuming that per cent level is the measure of choice, this is the level below which the bond value will fall with probability 1 per cent.

Table 8.6 'Through-the-cycle' credit migration matrix

From/to	AAA	AA	A	BBB	BB	B	CCC	D
AAA	90.82%	8.26%	0.74%	0.06%	0.11%	0.00%	0.00%	0.00%
AA	0.65%	90.88%	7.69%	0.58%	0.05%	0.13%	0.02%	0.00%
A	0.08%	2.42%	91.30%	5.23%	0.68%	0.23%	0.01%	0.05%
BBB	0.03%	0.31%	5.87%	87.46%	4.96%	1.08%	0.12%	0.17%
BB	0.02%	0.12%	0.64%	7.71%	81.16%	8.40%	0.98%	0.98%
B	0.00%	0.10%	0.24%	0.45%	6.86%	83.50%	3.92%	4.92%
CCC	0.21%	0.00%	0.41%	1.24%	2.67%	11.70%	64.48%	19.29%

Note: Typical rating agency migration over a one-year horizon.
Sources: Moody's; Fitch; Standard & Poor's; ECB's own calculations.

Table 8.7 'Point-in-time' credit migration matrix

From/to	AAA	AA	A	BBB	BB	B	CCC	D
AAA	66.3%	22.2%	7.4%	2.5%	0.9%	0.7%	0.1%	0.0%
AA	21.7%	43.0%	25.8%	6.6%	2.0%	0.7%	0.2%	0.0%
A	2.8%	20.3%	44.2%	22.9%	7.4%	2.0%	0.3%	0.1%
BBB	0.3%	2.8%	22.6%	42.5%	23.5%	7.0%	1.0%	0.3%
BB	0.1%	0.2%	3.7%	22.9%	44.4%	24.5%	3.4%	0.7%
B	0.0%	0.1%	0.4%	3.5%	20.5%	53.0%	20.6%	2.0%
CCC	0.0%	0.0%	0.1%	0.3%	1.8%	17.8%	69.9%	10.1%

Note: Typical Merton based rating migration over a one-year horizon.
Sources: Moody's KMV; ECB's own calculations.

Recovery rates are best characterized not by the distributional mean but rather by their consistently wide uncertainty.[30] There should be a direct relationship between this uncertainty and the estimate of volatility of price changes due to credit risk. This uncertainty can be incorporated in the calculation of price volatility by adjusting the variance estimate in Table 8.5 (see Gupton *et al.* 1997).

Finally, the selection of an appropriate time horizon is also important. Much of the academic and credit risk analysis and credit data are stated on an annual basis. However, we are interested in a haircut that would mitigate the credit risk that could be experienced in the time span between the default of the counterparty and the actual liquidation of the collateral.

[30] In case we were unable to infer from historical data or by other means the distribution of recovery rates, we could capture the wide uncertainty and the general shape of the recovery rate distribution by using the Beta distribution.

Table 8.8 99 per cent credit risk haircut for a five-year fixed coupon bond

	Holding Period (liquidation horizon)			
Rating	1 week	2 weeks	3 weeks	4 weeks
AAA	0.48%	0.67%	0.83%	0.99%
AA	0.46%	0.65%	0.80%	0.95%
A	0.71%	1.00%	1.24%	1.48%
BBB	1.13%	1.59%	1.97%	2.34%
BB	1.71%	2.42%	2.99%	3.56%
B	3.66%	5.18%	6.41%	7.63%

As discussed earlier, this holding period would normally be below one year, typically several weeks. The annual volatility estimate would need to be adjusted for the relevant holding period as in equation (8.4).

Table 8.8 illustrates typical credit-risk haircut levels for a fixed-income bond with five-year maturity and different holding periods. Notice the exponential behaviour of credit-risk haircuts, i.e. as credit quality decreases, the haircut level increases on an exponential basis. Haircuts with different holding periods are scaled using the square root of time as in equation (8.4). The ultimate credit-risk haircut for a given bond does not only depend on the degree of risk aversion of the institution measured by the confidence level of the credit VaR, but also and most crucially on the different assumptions taken as regards credit-risk migration, recovery rate level and associated volatility, credit spreads and holding period.

5. Limits as a risk mitigation tool

Collateral limits are the third main risk mitigation tool at the disposal of the collateral taker. The other two are mark-to-market policy and haircut setting. If the collateral received by the collateral taker is not well diversified, it may be helpful limiting the collateral exposure of the issuer, sector or asset class to for example a maximum percentage of the total collateral portfolio.

There are also haircut implications to consider when diversification in the collateral portfolio is not achieved. For example, consider the case of a counterparty that pledges the entire issue of an asset-backed security as the sole collateral to guarantee a repo operation with the central bank. The average volatility assumptions used to compute haircuts for asset-backed

securities may not hold for this particular bond, so the haircut will not cover its potential price movements. In this case, the collateral taker may decide to supplement the haircut level with an additional margin or to limit the collateral exposure to this bond and in effect forcing the counterparty to provide a more 'diversified' collateral pool.

Collateral limit setting can vary widely depending on the sophistication that the collateral taker would like to introduce. Ultimately, limit setting is a management decision that needs to consider three aspects: the type of limits, the risk measure to use and the type of action or policy that the limits imply in case of breach.

The collateral taker may consider three main types of limits: a) limits based on a pre-defined risk measure such as applying limits based on credit quality thresholds, for example only accepting collateral rated single A or higher,[31] b) limits based on exposure size so as to restrict collateral exposures above a given size, and c) limits based on marginal additional risk so as to limit the addition of a collateral to a collateral portfolio that increases portfolio risk above a certain level. Obviously, the collateral taker could implement a limit framework that combines these three types. Limits based on additional marginal risk need a portfolio risk measurement system. The concept of a portfolio risk measurement approach is appealing as it moves beyond the risk control of individual assets and treats the portfolio of collateral as the main subject of risk control. It is in this approach where diversification and interaction of collateral types can be treated in a consistent manner allowing the risk manager to control for risk using only one tool in the palette of tools. For example, instead of applying limits and haircuts to individual assets, a haircut for the entire collateral portfolio could be applied taking into account the diversification level of the collateral pool. Such portfolio haircut would penalize the collateral pools with little or no diversification.

6. Conclusions

This chapter has reviewed important elements of any central bank collateral management system: the credit quality assessment of eligible collateral and the risk mitigation framework.

[31] As regards risk measures that drive limit setting, it is important to keep in mind their application. The risk estimates underlying limits need to provide an accurate view of the relative riskiness of the various collateral exposures. Typical risk measures that can be used in limits include issuer rating information, credit risk induced standard deviation, average shortfall and/or correlation.

Collateral credit quality assessment is a vital part of the risk control arsenal of the central bank. Fortunately, recent regulatory developments have provided the possibility for the central bank to make use of additional credit assessment systems other than public rating agencies or own in-house credit systems. Internal rating based (IRB) systems supervised by national supervisors following the capital requirements regime of Basel II are a strong credit quality assessment source option. However, the possibilities of multiple credit assessment sources also come with challenges. In particular, central banks should lay down clear credit quality assessment frameworks, in which comparability and accuracy of rating outputs is ensured among the range of heterogeneous rating sources.

Risk control frameworks should be built around three main pillars: mark-to-market collateral valuation, haircuts and limits. Mark-to-market valuation is the first element of such a framework. Frequent mark-to-market valuation of collateral reduces risk. Unfortunately, not all assets, in particular those fixed debt assets that normally form part of the eligible set of central bank collateral, have valid or representative market prices. When mark to market is not possible, the central bank needs to resort to mark-to-model or theoretical pricing making sure that sufficient resources are allocated to this task. The implementation of haircuts or margins to collateral is crucial as well. Collateral assets vary in their risk profile. Ideally, haircuts mitigate or compensate for those additional risk factors associated to the collateral asset under question (e.g. additional liquidity or credit risk). The chapter describes different methods to estimate the level of haircuts based on different market, liquidity and credit risk profiles. The ultimate definition of the levels of haircuts depends on the degree of risk aversion of the central bank. Finally, collateral limits, the last tool in such a risk control framework, allow mitigating the risk of too much concentration in the collateral portfolio.

The type of risk control framework presented in this chapter relates to individual assets or generic asset types. Future developments in the risk control framework design would focus more on the portfolio concept, looking at the interaction of new collateral assets in a collateral portfolio and adjusting haircuts and limits based on the risk contribution to the portfolio.

9 Collateral and risk mitigation frameworks of central bank policy operations – a comparison across central banks

Evangelos Tabakis and Benedict Weller

1. Introduction[1]

Chapter 7 has presented a theoretical approach for deducing the optimal collateral framework of the central bank based on a cost–benefit analysis. The analysis was based on the assumption that the central bank needs to cover exogenously determined refinancing needs of the banking system vis-à-vis the central bank without exposing itself to risks beyond a certain level deemed acceptable. Furthermore, the central bank will always try to do so in the most cost-efficient way. Therefore, it will rank assets that are potentially eligible for collateral according to handling costs and the cost of the risk mitigation tools that would be needed to reduce the risk of these assets to the level set as acceptable by the central bank. These will then be included in the list of eligible assets in order of increasing cost until the aggregate outstanding volume of the eligible assets covers the corresponding needs of the system in cash.

Applying this basic framework in different central banks that act mainly as liquidity providers would normally lead to the implementation of collateral frameworks that are similar in their core features. Indeed when comparing the frameworks of the leading central banks, these exhibit many similarities. At the same time important differences also exist. This is also a result of the simplicity of the model described in Chapter 7 that does not

[1] The authors are indebted to Tonu Palm, Gergely Koczan and Joao Mineiro for their input to this chapter. Any mistakes and omissions are, of course, the sole responsibility of the authors. Parts of this chapter draw on the ECB monthly bulletin article published in October 2007, under the title 'The collateral frameworks of the Federal Reserve System, the Bank of Japan and the Eurosystem', pp. 85–100 (ECB 2007b).

fully capture the diversity of 'liquidity needs' in practice. These differences between central banks can be attributed to any of the following factors:

- Amount of liquidity deficit covered by collateralized lending: Generally speaking, the higher the amounts covered by such operations, the more extensive the range of collateral accepted would tend to be.
- Counterparty policy: The central bank may choose to establish a limited number of counterparties or open its operations to all financial institutions fulfilling some basic criteria.
- Financial markets: Volume and liquidity in government bond and corporate bond markets can affect the distribution of eligible assets between private and public issuers.[2]
- Differentiating according to type of market operations: A central bank could establish one list of collateral for all types of collateralized lending but it can also differentiate its collateral framework depending on the operations.
- Historical reasons: The cost–benefit analysis of Chapter 7 does not take into account the fact that costs are also incurred by *changes* in the collateral framework so that older practices, even if no longer optimal, may still be maintained at least for some time.

To make these differences concrete, this chapter looks at the collateral frameworks of the three major central banks: the Federal Reserve Board (FED), Bank of Japan (BoJ) and the European Central Bank (ECB).[3] It attempts to compare practices and, where possible, to explain the differences. Reference is made also to the collateral practices of a larger group of central banks which have been surveyed in 2005 in the context of a central bank risk management seminar organized by Bank of Canada.

Section 2 of this chapter introduces the main features of the collateral frameworks in the FED, BoJ and ECB linking them to basic principles and economic factors. Section 3 compares the concrete eligibility criteria applied by the three central banks and Section 4 treats the risk mitigation tools applied to collateral. Section 5 draws some general conclusions from these comparisons.

[2] The impact of the maturity of financial markets on monetary policy implementation is the focus of Laurens (2005).

[3] For a comparison of a more general scope between the monetary policy frameworks of the Eurosystem, the Federal Reserve and Bank of Japan, not focusing on collateral, the reader is referred to Borio (1997, 2001) and Blenck *et al.* (2001). Another such general comparison of the institutional framework, the monetary policy strategies and the operational mechanisms of the ECB, the FED and the pre-euro Bundesbank is provided in Apel 2003. Finally, Bindseil (2004) provides a general account of monetary policy implementation theory and practice.

2. General comparison of the three collateral frameworks

2.1 Types of operations

Open market operations represent the key instrument used by all three central banks for supplying liquidity to the banking sector. Open market operations can be conducted on either an outright or a temporary basis. Outright purchases result in assets being bought in the open market and remaining on the balance sheet of the central bank, leading to a permanent increase in banks' holdings of central bank money. Temporary open market operations, on the other hand, involve lending central bank money to banks with a fixed and usually short maturity. These operations allow the central bank to manage marginal liquidity conditions in the interbank market for overnight reserves and thus to steer very short-term money market interest rates so as to implement monetary policy decisions.[4]

In addition to temporary open market operations, all three central banks also conduct two other main types of credit operations, i.e. the borrowing facility and intraday credit. The borrowing (Lombard) facility – known as the marginal lending facility in the Eurosystem, the primary credit facility in the Federal Reserve System and the complementary lending facility in the Bank of Japan – aims to provide a safety valve for the interbank market, so that, when the market cannot provide the necessary liquidity, a bank can still obtain it from the central bank, albeit at a higher rate.[5] Moreover, central banks provide, on an intraday basis, the working balances which banks need to carry out payments.

For all these different types of credit operations – open market operations, the borrowing facility and intraday credit – the central bank generally requires counterparties to pledge collateral as security. An exception is the Federal Reserve System which does not require intraday credit to be collateralized except in certain circumstances (e.g. if the counterparty needs additional daylight capacity beyond its net debit cap, or if there are

[4] The importance of understanding the economic and policy issues related to the functioning of repo markets for conducting temporary open market operations was emphasized in BIS 1999. The ECB regularly publishes the results of studies on the structure and functioning of the euro money market (ECB 2007a).

[5] In the United States, until the reform of the Federal Reserve System's discount window in 2003, lending was only made on a discretionary basis at below-market rates. There were, however, certain exceptions, such as a special liquidity facility with an above-market rate that was put in place in late 1999 to ease liquidity pressures during the changeover to the new century. The complementary lending facility was introduced in 2001 in Japan.

Table 9.1 Differentiation of collateral policy depending on type of operation

	Federal Reserve System	Eurosystem	Bank of Japan
Temporary operations	Treasuries, Agencies, MBSs, separate auctions with different marginal rates	The same broad set of collateral accepted for all operations	The same broad range of collateral accepted for open market operations and the complementary lending facility
Borrowing facility	Wide set beyond the set for temporary operations		
Intraday credit	No collateralization as long as credit remains below cap		Mainly JGBs; other securities accepted under conditions

concerns about the counterparty's financial condition). Table 9.1 shows how the type of operation affects the collateral accepted in the three central banks.

2.2 Common principles

Of course, assuming that the collateral can be legally transferred to the central bank and that adequate valuation and risk control measures can be designed, there is, in theory, an almost infinitely wide range of assets which could potentially perform the role of collateral. This may cover liquid marketable fixed-income securities, such as government and corporate bonds, equity-style instruments, loans to the public sector, corporations or consumers, and even exotic assets such as real estate and commodities. Therefore, in order to guide decision making on what types of assets to accept as collateral, each central bank has established some guidelines or principles for its collateral framework. These principles can be distilled down to a rather similar set of elements:

- All three central banks require eligible collateral to be creditworthy in order to maintain the soundness of the bank's assets.
- The type and quantity of eligible collateral must allow the central bank to conduct its open market operations smoothly, even for large amounts at very short notice. In addition, the choice and quantity of collateral available must also allow the payment systems to function efficiently.
- All three central banks strive for efficiency. Thus, the collateral ideally should not cause costs in its mobilization to both the counterparty and the central bank which exceed the actual benefits to counterparties.

- All three central banks aim for a high degree of transparency and accountability. These principles ensure that the public trusts that the institution is behaving objectively, responsibly and with integrity, and that it is not favouring any special interests. For the collateral framework, this would imply selecting assets for eligibility based on objective and publicly available principles and criteria, while avoiding unnecessary discretion.
- All three central banks, albeit in rather different ways, strive to avoid distortions to asset prices or to market participants' behaviour which would lead to an overall loss in welfare.[6]

One of the asset classes which would normally most readily comply with these principles is marketable securities issued by the central government. Government securities are generally the asset class which is most available on banks' balance sheets and thus they ensure that operations of a sufficient size can be conducted without disrupting financial markets. Furthermore, government bonds have a low cost of mobilization, as they can be easily transferred and handled through securities settlement systems, and the information required for pricing and evaluating their credit risk is publicly available. Third, accepting government bonds would also not conflict with the central bank's objectives of being transparent, accountable, and avoiding the creation of market distortions.

Having said this, there are other types of assets that also clearly fulfill these principles. In fact, all three central banks have expanded the eligibility beyond central government debt securities, although to different degrees. The Federal Reserve System, in its temporary open market operations, accepts not only government securities, but also securities issued by the government-sponsored agencies and mortgage-backed securities guaranteed by the agencies; in its primary credit facility operations, the Federal Reserve System accepts a very wide range of assets, such as corporate and consumer loans and cross-border collateral. The Bank of Japan and the Eurosystem accept as collateral for temporary lending operations a very wide range of private-sector fixed-income securities, as well as loans to the public and private sector. For each central bank, the decision to expand eligibility beyond government securities can be explained by several factors related to the overall design of the operational framework, such as the size of the temporary operations and the decision on how many counterparties can

[6] The potential impact of collateral use on markets has been studied by the Committee on the Global Financial System, see CGFS (2001).

participate, and also by the financial environment in which the central bank operates, in particular, the depth and integration of non-government securities markets. These factors are explored in detail in the following two subsections.

2.3 Choices of the overall operational framework

One of the key aspects of the operational framework which impacts on the collateral framework is how the central bank supplies liquidity to the banking sector. Table 9.2 compares the size of central bank temporary operations, both in terms of amounts outstanding and as a proportion of their total balance sheet.

The table raises a number of interesting observations. First, the size of the Federal Reserve System's temporary open market operations is significantly lower than that of the Eurosystem and the Bank of Japan, both in absolute amounts and as a proportion of the balance sheet. This is because the Federal Reserve System primarily supplies funds to the banking sector via outright operations, which accounted for 90 per cent of its balance sheet at the end of 2006. The Fed's temporary operations play the role of smoothing short- to medium-term fluctuations in liquidity needs at the margin. Second, for all three central banks, the size of the Lombard facility is negligible, in line with its role of providing funds when the market cannot provide them and putting a ceiling on overnight interest rates. Third, the Eurosystem issues by far the largest volume of intraday credit, both in absolute terms and as a proportion of its balance sheet.

The size of the temporary operations clearly has an impact on the choice of collateral: all other things being equal, the larger the size of the operations, the greater the need to expand the type of collateral accepted to a wider set of instruments in order to ensure that the central bank can comply, in particular, with one of the principles identified in Section 2.2: the ability to conduct monetary policy and ensure the smooth operation of the payment systems.

A second important aspect of the overall operational set-up, which impacts on the design of the collateral frameworks, is the choice of counterparties which can participate in the various central bank operations. To ensure that its open market operations can be conducted efficiently on a daily basis and also at very short notice, the Federal Reserve System uses only a small group of currently twenty-one 'primary dealers'. These primary dealers are relied upon to re-distribute liquidity to the rest of the banking

Table 9.2 Comparison of sizes of credit operations (averages for 2006, in EUR billions)

	Federal Reserve System		Eurosystem		Bank of Japan	
	Average size of operations	% of balance sheet	Average size of operations	% of balance sheet	Average size of operations	% of balance sheet
Temporary operations	19	3%	422.4	38%	274	34%
Lombard facility	0.2	0%	0.1	0%	0.6	0.1%
Intraday credit	102	15%	260	24%	124.3	15.5%
Total	121	18%	682.5	62%	398.9	49.7%

Note: Converted to euro using end-2006 exchange rates.
Sources: Federal Reserve System; ECB; Bank of Japan.

sector. For the primary credit facility, the approach is different: all 7,000 credit institutions which have a reserve account with the Federal Reserve Bank and an adequate supervisory rating are allowed access. The Eurosystem's operational framework has been guided, instead, by the principle of ensuring access to its refinancing operations to any counterparty which so desires. All credit institutions subject to minimum reserve requirements can thus participate in the main temporary operations, provided they meet some basic requirements. Currently, about 1,700 are eligible to participate in regular open market operations, although in practice fewer than 500 participate regularly in such operations; whereas 2,150 have access to the Lombard facility and a similar number can use intraday credit. The Bank of Japan takes an intermediate approach in order to ensure that it can operate in a wide range of different markets and instruments, but at the same time also maintains operational efficiency: around 150 counterparties are eligible to participate in the fund-supplying operations against pooled collateral, but they must also fulfill certain criteria.

The selection of counterparties has certain implications: the wider their range, all other things being equal, the more heterogeneous is the type of collateral assets held on their balance sheets. In the case of the Eurosystem, this heterogeneity of counterparties' balance sheets was even greater – relative to the other two central banks – due to the fragmented nature of national financial markets at the inception of the euro in 1999. The Eurosystem has therefore considered it especially important to take into account this heterogeneity when designing its collateral framework, in order to

ensure that banks in the (by now fifteen) different countries of the euro area can participate in central bank operations with relatively similar costs of collateral and without requiring a significant restructuring of their balance sheets. In the case of the Federal Reserve System, instead, the relatively few counterparties participating in open market operations are very active in the government securities markets, so the Federal Reserve System can be fairly confident that these banks have large holdings of the same type of collateral. In contrast, for its primary credit facility operations, it has chosen a very diverse range of counterparties – even broader than for the Eurosystem open market operations.

2.4 External constraints

In addition to the design of the overall operational framework, the central bank also needs to take into account its specific financial environment, in particular the size of the government and private bond markets relative to the demand for collateral. In the United States, there are three types of fixed-income assets – the US Treasury paper, the agency bond securities and mortgage-backed securities – which have large outstanding amounts, are highly liquid and standardized, have a high credit quality and are widely held on the primary dealers' balance sheets. The large size and liquidity of the markets for these assets ensure that the central bank can intervene at short notice and for large amounts without disturbing financial markets. The high credit rating of the issuers ensures that the Federal Reserve System faces little risk; in addition, the fact that all these securities are book-entry format and can be easily priced and settled ensures operational efficiency; lastly, operating in highly standardized markets of a limited number of public or quasi-public entities ensures transparency. Given the relatively small size of the Federal Reserve System's temporary operations (and the fact that the majority of these are already collateralized with US treasuries), it would probably be feasible to implement monetary policy only with government bonds. But given that two other markets exist, which also obviously fulfill the Federal Reserve System's principles, granting eligibility to them provides even more flexibility to counterparties with relatively limited additional costs.

In the euro area, private-sector bond markets have not yet reached the same scale as in the United States, where the vast majority of residential mortgages are funded through the capital markets, in which the government-sponsored agencies have played a critical role. In Europe, instead, the

funding of residential mortgages is still predominantly done through retail deposits. It is estimated that retail deposits accounted for approximately 60 per cent of Europe's EUR 5.1 trillion of outstanding residential mortgage balances in 2005, with only 27.5 per cent funded using securities, such as covered bonds and mortgage-backed securities, and the remainder through unsecured borrowing. In addition, in Europe, the corporate bond market is less developed than in the United States, as firms have traditionally tended to obtain financing directly from banks rather than the capital markets. This is reflected in the composition of banks' balance sheets: loans to euro area residents accounted for EUR 15 trillion or 57 per cent of euro area banks' balance sheets at the end of 2006. The fact that loans still form a major part of the assets of Eurosystem counterparties, and will likely continue to do so for the foreseeable future, was one of the reasons why the Eurosystem developed a euro area-wide eligibility framework that includes loans to the corporate sector, which was launched at the start of 2007.

In Japan, private sector bond markets are also less developed than in the United States, with only a very small proportion of mortgages being financed through mortgage-related securities, and corporations mainly obtaining financing from banks rather than the capital markets. However, given that the government bond market is extremely deep, with higher outstanding issuance volume than both the US and euro area government bond markets, the lack of alternative private-sector bond markets has posed fewer difficulties for the Bank of Japan than for the Eurosystem. Nevertheless, the Bank has modified its collateral framework as the economic and financial environment has changed. It has also broadened the range of eligible collateral to include relatively new instruments such as asset-backed securities as the marketability of these instruments increased. Furthermore, it has made loans to the Deposit Insurance Corporation as well as to the Government's 'Special Account for the Allotment of Local Allocation Tax and Local Transfer Tax' eligible in early 2002. These actions noticeably increased the amount of eligible collateral and hence contributed to the smooth provision of liquidity under the quantitative easing policy.

3. Eligibility criteria

This section describes how the three central banks have translated their principles into eligibility criteria, while also taking into account the various

external constraints that they face. The precise eligibility criteria are summarized very broadly in Table 9.3.

There are a number of interesting similarities and differences. First, for the Federal Reserve System's open market operations, the eligibility criteria are fundamentally issuer-based: all debt securities issued by the US Treasury are eligible, plus all senior debt issued by the government-sponsored agencies (the largest of which are Fannie Mae, Freddie Mac and the Federal Home Loan Bank), plus all the mortgage-backed securities which are fully guaranteed by the same agencies. For the Eurosystem and the Bank of Japan's refinancing operations against pooled collateral, the eligibility criteria are more general and not issuer-based, so as to encompass a broader range of assets.

Second, the Federal Reserve System accepts a substantially wider range of collateral at its primary credit facility than in its open market operations; furthermore, the range of collateral accepted for its primary credit facility is also broader than that accepted in the borrowing facility at the Eurosystem and the Bank of Japan. For example, foreign currency-denominated securities, securities issued abroad, and mortgage loans to households are eligible for the Fed's primary credit facility, but would not be eligible in Japan or the euro area.

Third, the Eurosystem is the only central bank which accepts unsecured bonds issued by credit institutions as collateral in its main open market operations, although these are eligible in the Fed's primary credit facility. The Bank of Japan does not accept unsecured bonds issued by counterparties of the Bank, to avoid disclosing the Bank's judgement on any particular counterparty's creditworthiness and collateralizing credit to the counterparties with liabilities of the counterparties which may be redeemed by proceeds from the central bank's credit itself.

Fourth, asset-backed securities (ABS) are generally eligible for use in the main open market operations of all three central banks, although in the case of the United States they must be guaranteed by a government agency. The Eurosystem has established in 2006 some additional specific criteria that must be fulfilled by ABS and asset-backed commercial paper (ABCP)[8]: as well as fulfilling the other general eligibility criteria such as being denominated in euro and settled in the euro area etc., there must be a true sale of the underlying assets to the special purpose vehicle (SPV)[9] and SPV must be

[8] Only a very small number of ABCP are currently eligible, mainly because they do not fulfill one of the general eligibility criteria, in particular the requirement to be traded on a non-regulated market that is accepted by the ECB.

[9] A true sale is the legal sale of an underlying portfolio of securities from the originator to the special purpose vehicle, implying that investors in the issued notes are not vulnerable to claims against the originator of the assets.

Table 9.3 Comparison of eligibility criteria

		Federal Reserve System (temporary open market operations)	Federal Reserve System (primary credit facility)	Eurosystem	Bank of Japan
Type of assets	Marketable debt securities	√	√	√	√ Debtor must not be a counterparty
	Equities	–	√ Government agency stocks only	–	–
	Bank loans	–	√	√ Debtor must be a non-financial corporation or public-sector entity	√ Debtor must not be a counterparty
Type of issuer/debtor	Central government	√	√	√	√
	Government agency	√	√	√	√
	Regional, local government	–	√	√	√
	Corporate	–	√	√	√ Debtor must not be a counterparty
	Bank	–	√	√	√ Debtor must not be a counterparty
	Supranational	–	√	√	√ International financial institutions
	Asset-backed securities	√ Only if guaranteed by an agency	√	√ Only if there is a true sale of assets and SPV is bankruptcy remote from originator	√ Only if there is a true sale of assets and SPV is bankruptcy remote from originator

	Household	–	√ Residential property and consumer loans	–	–
Issuer residence	Domestic	√	√	√	√
	Foreign	–	√ Includes foreign governments, supranationals and European Pfandbriefe issuers	√ For marketable securities, it includes all 30 countries of the European Economic Area (EEA), the four non-EEA G10 countries and supranationals.	√ Valid only for commercial paper that is guaranteed by a domestic resident, certain foreign governments and supranationals
Seniority	Senior	√	√	√	√
	Subordinated	–	–	–	–
Credit standards	Minimum credit threshold for issuer or asset	Not applicable	Minimum rating of BBB or equivalent, but AAA for some complex or foreign currency assets	Minimum single A or equivalent	Minimum rating varies from single A to AAA depending on issuer group and asset class[7] JGB, government guaranteed bond and municipal bonds are eligible regardless of the ratings
Settlement	Domestic	√	√	√	√
	Foreign	–	√ Euroclear, Clearstream and third party custodians	–	–
Currency	Domestic	√	√	√	√
	Foreign	–	√ Usually only the major currencies	–	–

[7] For bills, commercial paper, loans on deeds to companies and other corporate debt, the Bank of Japan evaluates collateral eligibility based on its own criteria for assessing a firm's creditworthiness. Additionally, for some assets, the Bank of Japan requires debtors to have at least a certain credit rating level from credit rating agencies.

bankruptcy remote; the underlying assets must also not consist of credit-linked notes or similar claims resulting from the transfer of credit risk by means of credit derivatives. One of the clearest consequences of these criteria is that synthetic securitizations,[10] as well as collateralized bond obligations which include tranches of synthetic ABS as underlying assets, are not eligible. However, despite introducing these additional criteria, the volume of ABS that is potentially eligible is still very large, amounting to EUR 746 billion at the end of August 2007. The Bank of Japan has also established specific eligibility criteria for ABS and ABCP which are similar to the Eurosystem's; there must be a true sale (i.e. no synthetic securitization) and the SPV must be bankruptcy remote; there must also be alternative measures set up for the collection of receivables and the securities must be rated AAA by a rating agency. In its open market operations, the Federal Reserve only accepts mortgage-backed securities which are guaranteed by one of the government agencies (which are incidentally also only true sale securitization), but in its primary credit facility operations it would accept a wide range of ABS, ABCP and collateral debt obligations, including synthetic securitization. Furthermore, in August 2007, there was also a minor change in the primary credit facility collateral policy which implied that a bank could pledge ABCP of issuers to whom that bank also provides liquidity enhancements such as a line of credit.

Fifth, the Eurosystem and the Bank of Japan (as well as the Fed in its primary credit facility) accept bank loans to corporations and the public sector as collateral.

Sixth, in terms of foreign collateral,[11] there are both similarities and differences. In their open market operations, all three central banks only accept collateral in local currency, which is also issued and settled domestically. However, unlike the two other central banks, the Eurosystem also accepts assets denominated in euros but issued by entities from some countries outside the European Economic Area in its operations.

Lastly, all three central banks have somewhat different approaches regarding the assessment of compliance with the eligibility criteria and the disclosure to the banks of which assets are eligible. The Federal Reserve System, in its open market operations, publishes its eligibility criteria in

[10] A synthetic securitization uses credit derivatives to achieve the same credit-risk transfer as a true sale structure, but without physically transferring the assets.

[11] The Committee on Payment and Settlement Systems (CPSS) has studied the advantages but also the challenges of accepting cross-border collateral (CPSS 2006).

several documents and on its website (see Federal Reserve System 2002 and Federal Reserve Bank of New York 2007). Because of the simplicity of assets it accepts, there is no need to publish a list of eligible assets on its website. For its primary credit facility, the Federal Reserve System publishes a general guide regarding the eligibility criteria, and suggests that the counterparty contact its local Federal Reserve Bank regarding specific questions on the details of eligibility. The Bank of Japan publishes a general guideline on eligibility on its website,[12] which for most assets is sufficient to clarify to banks whether a specific asset is eligible or not. For some assets, in most cases whose obligors are private companies, the Bank of Japan only assesses eligibility at a counterparty's request. For the Eurosystem, the ECB publishes daily a definitive list of all eligible assets.[13] Because of the Eurosystem's very large and diverse collateral framework (about 26,000 securities are listed in the eligible asset database), as well as the decentralized settlement of transactions at the level of the Eurosystem NCBs, this is important both for transparency to counterparties and operational efficiency. For obvious reasons, the eligibility of bank loans can only be assessed on request and a list cannot be published.

4. Credit risk assessment and risk control framework

Once a central bank determines the level of risk that it will normally accept in collateralized lending, it has a number of tools to achieve that level of risk: counterparty borrowing limits; credit standards for collateral; limits on collateral issuers or sectors; collateral valuation procedures; initial haircuts; margin calls; and close links prohibitions. Chapter 7 of this book described these tools in detail drawing also on ECB (2004a). All three central banks use a combination of these tools and, unlike in the choice of eligible collateral, the underlying methodologies and practices of the risk control frameworks are relatively similar.

4.1 Credit risk assessment framework

The Eurosystem, Bank of Japan and FED (in its primary credit facility operations) consider external credit ratings by rating agencies as a main source of reference for determining whether assets have sufficiently *high*

[12] See Bank of Japan (2004) for details. [13] The general eligibility criteria can be found in ECB (2006b).

credit quality. The general threshold for a minimum rating is A– for the BoJ[14] and the Eurosystem. For the Fed's primary credit facility operations, the minimum rating is generally BBB−, but, like the BoJ, the Fed requires a higher rating for some complex assets (e.g. ABS). In addition to external ratings, the three central banks use a number of alternative sources of credit assessment. The BoJ uses its own in-house credit assessment system for corporate bonds, commercial paper and bills and requires these assets to exceed both the external and the internal rating thresholds. For its primary credit facility collateral, the Fed can also rely on counterparties' internal rating systems if these are accepted by the regulator. The Eurosystem uses all types of alternative credit assessments: in-house credit assessment systems, counterparties' internal rating systems as well as third-party rating tools.

4.2 Valuation

Regarding the *valuation* of collateral, there are only some minor differences in the practices of the three central banks. For the Federal Reserve System's repo operations, valuation is carried out daily using prices from a variety of private vendors. For its primary credit facility operations, revaluation takes place at least weekly, based on market prices if available. For the Eurosystem, valuation is carried out daily using the most representative price source, and, if no up-to-date price exists, theoretical valuation is used. For the Bank of Japan, daily valuation is used for the Japanese government bond repos, but weekly revaluation is used for the standing pool of collateral. For the valuation of bank loans, all three central banks generally use face value with the application of higher haircuts, generally depending on the maturity of the loan.

4.3 Risk control measures

All three central banks use *haircuts* to take account of liquidity and market risk. The haircuts depend on the liquidity characteristics of the asset, issuer group, asset type, the residual maturity of the asset and the coupon type. For the primary credit facility, if a market price does not exist, the Federal Reserve System uses the face value and applies higher haircuts.

A detailed comparison of haircut schedules of the three central banks would be difficult due to the differences in the set of eligible assets. In

[14] For some special asset types (e.g. asset-backed securities, agency bonds, foreign government bonds), the BoJ requires a higher rating and/or ratings from more than one rating agency.

Table 9.4 Comparison of haircuts applied to government bonds

	Federal Reserve System[15]	Eurosystem	Bank of Japan
Up to 1 year	2%	0.5%	1%
1–3 years	2%	1.5%	2%
3–5 years	2%	2.5%	2%
5–7 years	3%	3.0%	4%
7–10 years	3%	4.0%	4%
10–20 years	7%	5.5%	7%
20–30 years	7%	5.5%	10%
>30 years	7%	5.5%	13%

Table 9.5 Comparison of haircuts of assets with a residual maturity of five years

	Federal Reserve System	Eurosystem	Bank of Japan
Government bonds	2%	2.5%	2%
Regional/Local Government bonds	3%	3.5%	3%
Corporate bonds	3%	4.5%	4%
ABS	2–3%	5.5%	4%
Loans to corporates	10–13%[16]	11%/20%[17]	15%

particular the haircuts applied by the Fed in its open market operations are not public. Therefore tables 9.4 and 9.5 compare the haircuts applied by the Fed in its primary credit facility to those applied by the Eurosystem and Bank of Japan in their main open market operations.

Table 9.4 compares the haircuts applied to debt instruments issued by central governments for different residual maturities.

Table 9.5 compares the haircuts applied to various asset types accepted by both central banks by fixing the residual maturity to five years.

All three central banks use global margin calls in case the aggregate value of the collateral pool falls below the total borrowing by the counterparty in a

[15] Haircuts apply to the primary credit facility. If the market price of the securities is not available, a 10 per cent haircut is applied independently of maturity.

[16] These haircuts apply to individually deposited loans. Group deposited loans are subject to higher haircuts.

[17] The Eurosystem haircut for loans to corporates in this maturity bucket is 11 per cent if the value of the loan is computed by a theoretical method (discounting cash flows). In most cases however, the value of the loan is computed on the basis of the outstanding amount in which case the haircut is 20 per cent.

particular operation, i.e. margin calls are not calculated on an asset-by-asset basis. All three central banks apply daily valuation and execute margin calls for their open market operations.

None of the central banks currently uses *counterparty borrowing limits*[18] for their temporary operations, and no predetermined limits are placed on exposure to certain individual collateral issuers or guarantors.

Finally, all three central banks prohibit counterparties from using assets where they may have a close financial link with the issuer, which would negate the protection from the collateral. This minimizes the risk of a double default scenario. The Bank of Japan does not generally accept any asset issued by its counterparties thus significantly decreasing the concentration of its exposures. The Federal Reserve does not accept any bank bonds in its open market operations.

Box 9.1. Survey of credit and market risk mitigation in a collateral management in central banks[19]

On 6 and 7 June 2005, a central bank risk management seminar attended by twenty-two central banks took place in Ottawa, hosted by Bank of Canada. The seminar focused on credit risk and part of it was dedicated to the management of collateral. The twenty-two participating central banks replied to a survey of their collateral practices. All of these central banks accepted government securities as assets. Securities issued by private entities were accepted by seventeen central banks, non-marketable assets by eleven central banks and equity by one central bank. About 55 per cent of the participating central banks mentioned the use of cross-border collateral for some of their operations. Three main issues on the risk management of collateral were covered in the survey: credit risk assessment, risk control measures and asset valuation.

The credit assessment of the collateral was primarily based on ratings provided by recognized rating agencies. For those operations where a wide range of assets was accepted the rating threshold was set lower, typically at A or in one case at BBB. A few central banks reported that domestic government paper was accepted as a rule regardless of rating. As far as a rating threshold was applied, it was set to a single A level by seven central banks, to A- by four, and to AAA, AA-, and BBB- by one central bank, respectively.

One available agency rating was usually enough but three central banks mentioned they require two ratings. About one-third of the respondents mentioned the use of some form of an in-house credit assessment for some assets and two central banks mentioned the use of the assessment of commercial banks.

[18] Counterparty limits are, instead, a typical risk control measure in transactions between private institutions (see, for example, Counterparty Risk Management Policy Group II 2005).

[19] Source: Bank of Canada.

About half of the central banks surveyed reported using a rather simple haircut policy with a limited number of different haircut values in the range of 1–5 per cent applied to all collateral which was based on standard market practices rather than a specific model-based methodology. Central banks with a wider range of eligible collateral tended to develop also more complex risk control frameworks often based on a VaR calculation, using historical asset volatilities and an estimation of the assets liquidity and distinguishing among different residual maturities. Seven central banks reported the use of some form of concentration limits at least for some type of collateral and seven central banks used pooling of collateral across some of their operations.

Daily valuation was the norm for all collateral accepted. In rare cases and for some operations a weekly valuation was applied and one central bank mentioned valuation of assets twice a day. The use of margin calls was linked to the complexity of the overall risk control framework. In general, a threshold is agreed (either in percentage or absolute value) beyond which a call for additional collateral is triggered. Valuation problems because of lack of market prices arose in those central banks that accepted a wide range of assets which include also illiquid securities or loans. In these cases one central bank used the face value of the asset while others computed the present value by discounting future cash-flows. One central bank made use of ISMA (international securities market association) prices and another central bank mentioned the use of vendor tools.

5. Conclusions

This chapter's main focus was a comparison of collateral policies and related risk management practices of three major central banks (the Federal Reserve Board, Bank of Japan and the European Central Bank) supplemented by less detailed information on a larger group of central banks. This comparison could serve also as an informal test of the model of collateral management policy presented in Chapter 7. Two general facts distilled from the comparison seem to suggest that the model does capture the 'way of thinking' of central banks when developing their collateral policy.

First, central banks that implement monetary policy mainly or partly by lending to the banking system collateralize their exposure. This implies that protection against financial loss in such operations, even if these have a policy objective, ranks high in the priorities of central banks' policies.

Second, the first assets to be accepted as eligible collateral are invariably government securities. This seems to confirm the prediction of the model that assets are included in the list of eligible collateral in the order of increasing risk mitigation costs. Government securities, arguably the least risky assets to be accepted as collateral, carry a minimum such cost.

At the same time, it becomes clear that the model is too simple to capture and explain the variability of collateral policies among central banks even if these implement monetary policy in broadly similar ways. Both differences in the fundamental principles chosen as the basis for the collateral policy of the central bank as well as the differences in the financial markets in which central banks operate are important determinants of the ultimate form that the collateral framework will take. Finally, the fact that collateral management is a cost-intensive function in a central bank suggests that decisions to change it could be difficult and slow explaining also why practices may remain different despite converging tendencies.

10 Risk measurement for a repo portfolio – an application to the Eurosystem's collateralized lending operations

Elke Heinle and Matti Koivu

1. Introduction

This chapter presents an approach to estimate tail risk measures for a portfolio of collateralized lending operations. While the general method is applicable to any repo portfolio, this chapter presents an application of the approach to the estimation of the residual risks of the Eurosystem's collateralized lending operations (which on average exceeded half a trillion euro during 2006).

This chapter can be viewed as extending one of the specific steps constituting any collateralization framework as described in Chapter 7 Section 2.4 ('Monitoring the use of the collateral framework and related risk taking'). Any efficient collateralization framework will provide some discretion to counterparties on what types of collateral to use, and to what extent. This discretion implies that the actual risk taking, for instance driven by concentration risks, cannot be anticipated. The central bank only can ensure that the outcome is actually acceptable by closely monitoring the actual use of the collateralization framework by counterparties, and establishing a sound methodology to measure residual risks. If it is not acceptable, specific changes to the framework are necessary to address the non-anticipated (concentration) risks that arose. The thorough monitoring is the precondition for a collateralization framework that provides leeway to counterparties, and therefore also for an efficient framework.

For the implementation of monetary policy, the Eurosystem has a number of instruments available of which liquidity-providing reverse transactions have so far been the most important. In these transactions, the Eurosystem buys specific types of assets under repurchase agreements or conducts credit operations collateralized by such assets. In these reverse transactions the

Eurosystem incurs a counterparty risk, since the counterparty may be unable to meet its credit obligations. This type of credit risk is mitigated by the requirement of adequate collateral to guarantee the credit provided. Article 18.1 of the Statute of the European System of Central Banks requires that all Eurosystem credit operations have to be based on adequate collateral.

The Eurosystem's collateral framework[1] translates the statutory requirement of adequate collateralization into concrete tools and procedures that guarantee sufficient mitigation of the financial risks in a reverse transaction. However, the collateral framework cannot provide absolute security and therefore there remain some risks for the Eurosystem. These risks that only arise in case of a counterparty default can be grouped into two categories: (1) the credit risk associated with the collateral accepted; (2) the liquidity-related risk associated with a drop in the market value of the collateral accepted before its liquidation.

To assess the adequacy of the risk control framework, it is necessary to measure the residual risks of the Eurosystem's credit operations. The risk measure used is expected shortfall (ES) at a 99 per cent confidence level. Since risk is approximated by simulation, it is not necessary to enforce distributional restrictions on the calculation of ES. Defaults for the counterparties and issuers are simulated by using Monte Carlo simulations with variance reduction techniques. With these techniques, the number of required simulations can be largely reduced and at the same time the accuracy of the resulting estimates can be improved.

This chapter is structured as follows: Sections 2 and 3 describe the data set and the assumptions used for the estimation of the residual risks, split up into credit risk (Section 2) and liquidity-related risk (Section 3). Section 4 addresses issues related to concentration risks. Sections 5 and 6 describe the risk measure used for the estimation of residual risks and explain the applied Monte Carlo simulation techniques. In Section 7 the results of the residual risk estimations for the Eurosystem's monetary policy operations are presented. Section 8 concludes.

2. Simulating credit risk

Credit risk in the Eurosystem's monetary policy operations is limited to the so-called double default events. Only if the counterparty who has submitted

[1] For further details see ECB (2006b).

the collateral and the collateral issuer default at the same time, losses due to credit risk may arise for the Eurosystem. This probability of a joint default mainly depends on the following parameters:

- The counterparty's probability of default (PD);
- The collateral issuer's PD;
- The default correlation between the counterparty and the collateral issuer.

The Eurosystem has put in place some risk mitigation measures to limit the probability of a joint default. As regards the collateral issuer's PD, the collateral issuer's minimum credit quality must at least correspond to a single-A rating based on a first-best rating. A PD over a one-year horizon of ten basis points is considered as equivalent to a single-A credit assessment. Moreover, in order to limit the default correlation between the counterparty and the collateral issuer, the Eurosystem collateral framework does in principle not foresee that a counterparty submits as collateral any asset issued or guaranteed by itself or by any other entity with which it has close links.

However, the Eurosystem has defined some exceptions to this no-close-link provision, like for example in the case of covered bonds. Moreover, the Eurosystem opted to give a broad range of institutions access to its monetary policy operations and therefore sets no restrictions on the counterparty's credit quality and hence its PD. Additionally, the Eurosystem sets so far no limits on the use of collateral from certain issuers or on the use of certain types of collateral. All these factors are potential risk sources in the Eurosystem's monetary policy operations that may especially materialize in phases of financial stress.

For the estimation of the credit risk arising from the Eurosystem's monetary policy operations, the expected shortfall / credit value-at-risk is estimated by using simulation techniques (see Section 6) that broadly rely on the CreditMetrics approach. The data set used for these estimations is a snapshot taken in November 2006 on the assets submitted by the Eurosystem's counterparties. The total amount of submitted collateral adds up to around EUR 928 billion which is spread among more than 18,000 different counterparty-issuer pairs.

In order to make this high dimensional problem operationally workable, some few basic assumptions need to be made. These assumptions refer mainly to the PDs, the recovery rates in the case of defaults and the dependencies between the defaults of issuers and counterparties. They are discussed in the following two subsections.

2.1 Default probabilities and recovery rates

In this analysis, credit risk is estimated over an annual horizon – the underlying assumption being that the collateral portfolio remains fixed. PDs for the various entities (counterparties and issuers) are derived from their credit ratings. For the derivation of PDs for the different rating grades, historical default rate information from three international rating agencies (FitchRatings, Moody's, Standard & Poor's) is used.[2]

The credit rating information for counterparties and issuers is collected on a second-best rating basis. The average credit rating for counterparties is around AA-. And the average credit rating for collateral is around AA, while the average credit rating for bank bonds and corporate bonds is lower than average and the average credit rating of asset-backed securities, government bonds and covered bonds is above this average.

The historical default rate information captures default statistics for the corporate sector that includes a wide variety of industries, including banks and real estate. Since it is assumed that rating agencies produce unconditional PDs, it can be expected that e.g. a BBB-rating from an industrial sector to be the same as a BBB-rating from the financial industry in terms of PDs.

The benchmark PD for each rating grade is derived by applying the central limit theorem to the arithmetic averages of the default frequency over the respective time period. As such, it is possible to construct confidence intervals for the true mean $\mu_{\bar{x}}$ of the population around this arithmetic average. The central limit theorem states that the arithmetic average $\bar{x}$ of n independent random variables x_i, each having mean μ_i and variance ${\sigma_i}^2$, is approximately normally distributed with parameters $\mu_{\bar{x}} = \frac{\sum_{i=1}^{n}\mu_i}{n}$ and $\sigma_{\bar{x}}^2 = \frac{\sum_{i=1}^{n}\sigma_i^2}{n^2}$. Applying this theorem to the rating agencies default frequencies, random variables with $\mu_i{=}p$ and $\sigma_i^2 = p(1-p)/N_i$, yields the result that the arithmetic average of the default frequencies is approximately normal with mean $\mu_{\bar{x}} = \frac{\sum_{i=1}^{n}p}{n} = p$ and variance $\sigma_{\bar{x}}^2 = \frac{\sum_{i=1}^{n}\frac{p(1-p)}{N_i}}{n^2}$. After estimating p and $\sigma_{\bar{x}}^2$ from the rating agencies' data, confidence intervals for the mean, i.e. the default probability p, can be constructed. These confidence intervals can then be used to derive estimates for annual PD thresholds for each credit quality step.

[2] For further details on the methodology used see Coppens *et al.* 2007; for further details on PD information, see Standard & Poor's 2006; Hamilton and Varma 2006; FitchRatings 2006.

Tabel 10.1 Default probabilities for different rating grades

Rating	Numerical rating	Annual PD
AAA/Aaa	1	1 basis point
AA+/Aa1	2	2 basis points
AA/Aa2	3	3 basis points
AA−/Aa3	4	4 basis points
A+/A1	5	6 basis points
A/A2	6	8 basis points
A−/A3	7	10 basis points
BBB+/Baa1	8	20 basis points
BBB/Baa2	9	30 basis points
BBB−/Baa3	10	42 basis points

Sources: Standard & Poor's (2006); Hamilton and Varma (2006); FitchRatings (2006); own calculations.

The results obtained using a (two-sided) 99.9 per cent confidence interval are summarized in Table 10.1. These figures are used as input parameters for the credit risk calculations.

The PDs of issuers are scaled down linearly from the annual PDs according to the liquidation time of the least liquid instrument that has been submitted from the issuer. This approach is based on the idea that whenever a counterparty defaults (which may be at any point in time during the year considered), a double default only occurs when an issuer from a counterparty's collateral pool also defaults during the time it takes for the liquidation of the asset. The scaling down of the issuer PDs to the liquidation period is therefore a possible way to consider the timing of defaults.

Linear scaling of PDs is used for example in CreditMetrics (see Gupton *et al.* 1997). It reflects a rather conservative approach (see Bindseil and Papadia 2006). In line with the CreditMetrics model, the one-year PDs are simply divided by fifty-two, if the liquidation time is one week.

In this analysis the same liquidation time assumptions are used as those applied for the derivation of haircut levels for eligible marketable assets. For this purpose, the different types of marketable assets are grouped into four different liquidity categories, arranged from most liquid to least liquid assets. The total liquidation time is largely based on assumptions regarding the so-called 'valuation period', 'grace period' and 'actual realization time' and their relation with the default event time. It is assumed that the valuation and grace period is the same for all asset classes (three to four working days). The realization time refers to the time necessary to orderly liquidate

Tabel 10.2 Liquidation time assumptions used for the different asset classes

Liquidation time	Category I 5 days	Category II 10 days	Category III 15 days	Category IV 20 days
Asset classes	Central government debt instruments Debt instruments issued by central banks	Local and regional government debt instruments Jumbo covered bank bonds Agency debt instruments Supranational debt instruments	Traditional covered bank bonds Credit institution debt instruments Debt instruments issued by corporate and other issuers	Asset-backed securities

the asset. This realization time is derived for each asset class separately by using a combination of quantitative and qualitative criteria. For an overview of the liquidation time assumptions used for this analysis, see Table 10.2.[3]

In this analysis the PDs of issuers are scaled down linearly from the annual PDs according to the liquidation time of the *least liquid instrument* submitted from the issuer. This is due to the constraint that a bond-specific analysis is operationally not possible on this level given the high dimension of the problem. To be conservative in the assumptions, the least liquid instrument from an issuer was chosen to fix the liquidation time applied. But currently there are only few issuers that have issued debt instruments that belong to different categories as listed in Table 10.2.

With regard to the recovery rates, the basic assumption is a constant recovery rate of 40 per cent for all bonds. This assumption is roughly in line with estimates for senior unsecured bonds reported by Altman *et al.* (2004). A constant recovery rate of 40 per cent is of course a simplifying assumption and in reality recovery rates depend on a number of factors, like the economic cycle (see Frye 2000), the conditions of supply and demand (see Altman *et al.* 2005a), the seniority of the assets within the capital structure (see Acharya *et al.* 2003) or the initial credit quality of the assets (see Varma *et al.* 2003). But since all the debt instruments considered in this analysis are of comparable high credit quality (single-A rating or above) and in principle no bonds are accepted that have subordinated structures, the application of one single recovery rate for all the assets seems acceptable.

[3] For further details on the haircut framework in the Eurosystem's monetary policy operations, see: ECB 2006b, 49 ff. Further information may also be found in Chapter 8 of this book.

2.2 Default correlation

A very important determinant of double defaults is the correlation between the defaults of counterparties and issuers. Since this is a difficult quantity to observe directly, asset correlations are used instead. These are easier to observe and could, for example, be estimated from financial statements by looking at the way the assets of different companies move together. The intuition behind the use of asset returns of the counterparty to estimate the probability of joint defaults lies in Merton's structural model for default. In Merton's model, the firm's default is driven by changes in the asset value of the firm. As a result, the correlation between the asset returns of the two obligors can be used to compute the default correlation between the two obligors.

For the estimation of credit risk in this context an equal asset correlation (between and across counterparties and issuers) is used. Such an approach is partly necessitated due to technical restrictions, since the correlation matrix needs to be positive definite. This would be extremely difficult to guarantee in case the individual correlations differ. The approach of using a fixed correlation level for all can be thought of as mainly modelling the common systematic elements driving the asset returns of all companies.

Under normal conditions, a fixed correlation level of 24 per cent is assumed. This assumption is based on academic studies[4] which have shown that the average asset correlation to be in the range of 22.5 per cent and 27.5 per cent, and on the Basel II accord which (approximately) assumes a 24 per cent asset correlation for highly rated assets. In Basel II, the asset correlation is determined by the following formula:[5]

$$Correlation = \frac{1}{1-\exp(-50)}\{0.12^{*}(1-\exp(-50^{*}PD)) + 0.24^{*}\exp(-50^{*}PD)\}$$

It is important to note that asset and default correlation are different concepts. A default correlation is defined as the correlation between two random variables that get a value of one when the corresponding company defaults and a value of zero otherwise (over a fixed time interval). This can roughly be interpreted as how often both companies default when one of them defaults. Therefore, default correlation is determined both by the asset

[4] See for example Lopez (2002); Ramaswamy (2005).
[5] See BCBS (2006b, 64).

correlation and the default probabilities. For a given level of asset correlation, default correlation is a (generally increasing) function of the individual PD.[6]

Another aspect to be considered is the 'nature' of the dependence. A common approach – which is also followed here – is to use a normal copula model, where the dependence is introduced through a multivariate normal vector $(x_1, \ldots, x_d)$. Each default indicator is represented by $Y_k = \mathbf{1}\{x_k > z_k\}$, $k = 1, \ldots, d$, with z_k chosen to match the marginal default probability p_k. Since the issuer defaults are assumed to follow a multivariate normal distribution, it follows that $z_k = \Phi^{-1}(1-p_k)$, where Φ^{-1} denotes the inverse of the standardized cumulative normal distribution.

The use of a normal copula model is widespread. Such an approach is for example also followed in Moody's KMV or in CreditMetrics. This frequent use of the multivariate normal distribution is certainly related to the simplicity of its dependence structure, which is fully characterized by the correlation matrix.

3. Simulating liquidity-related risks

Liquidity-related risks can arise if the value of the collateral falls in the period between the counterparty's default and the realization of the collateral. In the time between the last valuation of the collateral and the realization of the collateral in the market, the collateral price could decrease to the extent that only a fraction of the claim could be recovered by the borrower. Liquidity risk may be defined as the risk of financial loss arising from difficulties in liquidating a position quickly without this having a negative impact on the price of the asset. Market risk may be defined in this context as the risk of financial loss due to a fall of the market value of collateral caused by exogenous factors. In the following, these two different kinds of risk will be treated jointly as liquidity-related risks.

The Eurosystem's collateral framework foresees several risk mitigation measures in order to reduce considerably these liquidity-related risks. As regards the valuation of collateral, collateral needs to be valued on a daily basis using the most representative price on the business day preceding the valuation date. For non-marketable assets in general, and for marketable assets in case no sufficiently reliable market price is available, the Eurosystem uses a theoretical price valuation.

[6] For a more rigorous treatment of default correlation, see Hanson *et al.* (2005).

With respect to the risk control measures currently applied by the Eurosystem,[7] 'valuation haircuts' play the most important role. When accepting the collateral, the Eurosystem deducts a certain percentage of the collateral value in order to ensure that there are no losses at liquidation. This percentage depends on the price volatility of the relevant asset class and on the prospective liquidation time. The Eurosystem sets haircuts to cover 99 per cent of the price changes within the assumed orderly liquidation time of the respective asset class.

Moreover, the Eurosystem currently applies variation margins. The Eurosystem requires that the market value of the underlying assets used in its reverse transactions cover the provided liquidity over the life of the transaction. Thus if this value, measured on a daily basis, falls below a certain level, counterparties have to supply additional assets or cash. Similarly, if the value of the underlying assets exceeds a certain level, the counterparty may retrieve the excess assets or cash.

For a loss due to liquidity-related risks to occur, the price drop of the asset after the default of the counterparty has to be greater than covered for by haircuts. This makes it quite a rare event, if liquidation time assumptions are sufficient, since haircuts are calculated using a 99 per cent confidence level, meaning that price drops smaller than 2.33 volatilities are covered for. Denoting by X the price movement which is drawn from a normal distribution, a loss due to liquidity risk from a single exposure from a counterparty–issuer pair (i, j) is

$$\mathrm{L_{i,j}} = \mathrm{default}(i)^{*}(1 - \mathrm{default}(j))^{*}\mathrm{exposure}_{i,j}{}^{*}\sigma^{*}\max(-2.33 - X, 0)$$

where default(i) equals one if entity i defaults, and equals zero otherwise.

For the estimation of liquidity-related risk in the Eurosystem's credit operations, some further assumptions have to be made. First of all, a distributional assumption for price movements is necessary. The usual practice is followed here, meaning that a normal distribution for price changes is assumed.

As regards the assumption on volatility, due to technical reasons and since the simulation will not be performed on a bond-by-bond basis, the same volatility will be assumed for all the assets in the collateral pool. For a derivation of this volatility figure, a simple approach was chosen. The volatility estimate was determined by calculating a series of day-to-day volatilities from

[7] See also ECB (2006b).

a monthly sliding window during the last three years, separately for different maturities, by using a government yield curve.[8] In order to be conservative in the volatility estimate, a maximum out of the series of volatilities is taken to derive the volatility figure. Then this daily volatility figure is scaled into a weekly volatility. The result obtained from these calculations is a value of around 1.2 per cent.

Given the fact that the collateral must be valued on a daily basis according to the Eurosystem collateral framework, the basic assumption will be that the value of collateral assigned to it by the Eurosystem reflects its market value at the time of default. Given this assumption, the relevant time horizon for the calculation of price fluctuations is the time it takes to liquidate the instrument. It is assumed that the liquidation time assumptions of the risk control framework (see Table 10.2) hold.

4. Issues related to concentration risks

Portfolio risks are obviously determined to a large extent by concentration and correlations of defaults and price changes. Default correlations are in addition particularly relevant for a repo portfolio in which credit risk is linked to PD.

This section therefore deals with concentration-related risks that are an important source of credit risk in the Eurosystem collateral framework. The dimensions of concentrations are manifold. Figure 10.1 gives an overview of the most important types of concentrations the Eurosystem is exposed to in its collateral operations. Since all these different types of concentrations interact, a fully comprehensive assessment on the current level of concentration and on the maximum level of acceptable concentration in the Eurosystem collateral framework is a complex issue.

For the residual risk estimations, similar assumptions as in the Basel II framework are made. Following the Asymptotic Single-Risk Factor model that underpins the internal ratings based approach in the new Basel capital accord, it is assumed that i) there is only one source of systematic risk and that ii) the portfolios are perfectly fine-grained meaning that idiosyncratic risk has been fully diversified away. Assumption i) implies that the commonality of risk between any two individual credits is uniquely determined by the intensity of their respective sensitivities to the single systematic

[8] An approximation for price volatility can be obtained by multiplying the yield volatility with the instrument's modified duration.

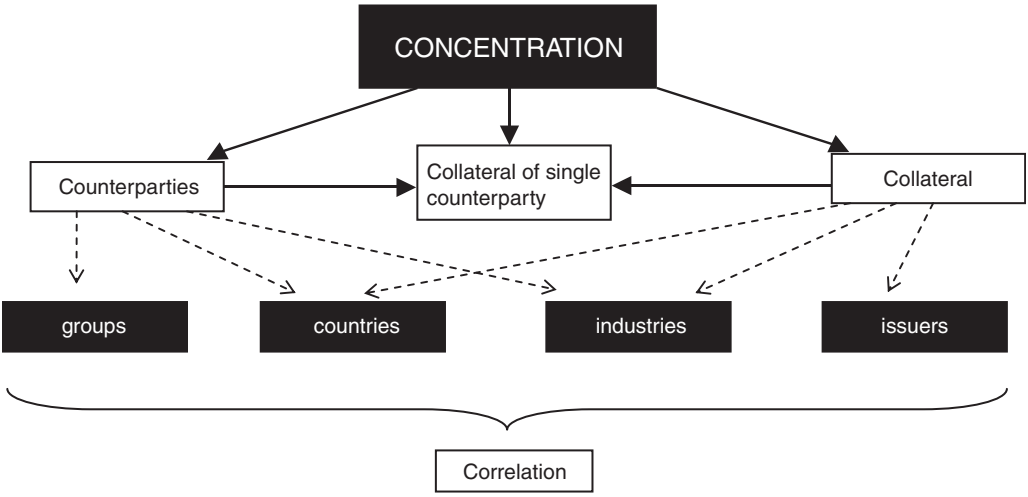

Figure 10.1 The most important types of concentrations in the Eurosystem collateral framework.

factor. That means that it is assumed that the portfolios are well-diversified across sectors and geographical regions, so that the only remaining systematic risk is to the performance of the economy. In practical terms, this is modelled by assuming a unique and constant correlation[9] between and across all the counterparties and collateral issuers. As already mentioned in Section 2.2, the standard assumption chosen for the residual risk estimations is a uniform asset correlation of 24 per cent.

In the following, the most important potential sources of concentration in the Eurosystem collateral framework are analysed more in-depth. Identified concentration risks are then translated into a granularity adjustment for credit risk or might be translated into a corresponding adjustment of the above mentioned correlation assumption.

4.1 Concentration on the level of counterparties

Concentration can arise on the level of counterparties, meaning that collateral may be submitted to the Eurosystem by only a few counterparties. As can be seen from the Lorenz curve[10] in Figure 10.2, there is indeed a high

[9] This approach is also necessitated due to technical restrictions, since the correlation matrix needs to be positive definite.

[10] The Lorenz curve of a probability distribution is a graphical representation of the cumulative distribution function of that probability distribution. In the case of a uniform distribution, the Lorenz curve is a straight line.

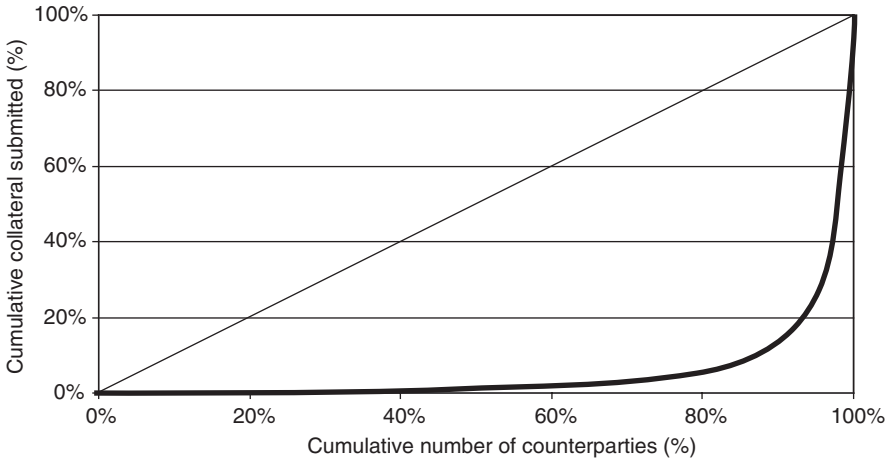

Figure 10.2 Lorenz curve for counterparties with respect to amount of collateral submitted.
Source: own calculations.

degree of concentration on the level of counterparties. The corresponding Gini coefficient[11] is 0.898. Indeed, out of the 1264 counterparties submitting collateral, 25 already account for half of the submitted collateral.

Within the counterparties, there can moreover be *concentration on a group level*, meaning that different counterparties effectively belong to the same banking group. Depending on the concrete design of the group structure, especially as regards support mechanisms and unlimited liability for other group members in case of default of one of the entities belonging to the group, the implications of concentrations on a group level can be quite different. The highest degree of concentration can be observed if there is full joint liability from the whole banking group in case of default of one entity. However, due to technical, data and resource restrictions, a comprehensive assessment of the current level of concentration on a banking group level is not made for the purposes of this analysis.

Another concentration that can arise on the level of counterparties is *concentration by countries* since certain risk factors may be country specific. Given the fact that quite often counterparties from different countries form

[11] The Gini coefficient is a measure of inequality of a distribution, defined as the ratio of the area between the Lorenz curve of the distribution and the Lorenz curve of the uniform distribution (which is a straight line), to the area under the Lorenz curve of the uniform distribution. It is a number between zero and one, where zero corresponds to perfect equality (i.e. all counterparties submitted the same amount of collateral) and one corresponds to perfect inequality (i.e. only one counterparty submits collateral).

one banking group, the determination of the ultimate country risk becomes more and more difficult. To get some idea on the distribution of counterparties by countries, the counterparties can be grouped together according to their country of residence of the ultimate parent. Such an analysis reveals that counterparties are mainly concentrated in Germany whose counterparties have submitted almost 57 per cent of the total amount of collateral submitted to the Eurosystem. Among the twenty-five most important counterparties, seventeen are located in Germany, three in Spain, two each in the Netherlands and Belgium and one in France.

Finally, there is *concentration on the level of industries*. Since the Eurosystem's counterparties belong by definition to the banking sector, there is a maximum degree of concentration by industry.

As regards the risk implications of counterparty concentration, the following can be concluded: overall, there is currently no perfect granularity on the level of counterparties. This type of concentration is, however, an exogenous factor that is driven by structural facts. In this respect it should be noted that counterparty concentration could be even higher if the Eurosystem's monetary policy framework did not aim at ensuring the participation of a broad range of counterparties.

4.2 Concentration on the level of collateral

Concentration on the level of collateral can arise in many respects. First, there can be concentration on the level of issuers. Second, concentration may be observed on the level of different asset categories. Third, concentration may arise by industry to which the collateral issuers belong. Finally, there may be concentrations on the level of collateral as regards country exposure.

Concentration by issuers is illustrated in Figure 10.3. The chart shows the Lorenz curve as regards collateral issuer concentration. The Gini coefficient of 0.823 indicates a slightly lower level of concentration than in the case of counterparties. However, forty-two issuers already account for half of the collateral submitted. Among them, German issuers – mainly banks and (regional) government – form the most important share.

Another level of *concentration is by industries*. Concentration by industry can first be approached by looking at concentrations by asset categories: There is a dominance of unsecured bank bonds that is submitted as collateral (33.2 per cent). This implies a remarkable exposure to the banking industry, in particular in view that the first line of defence, the

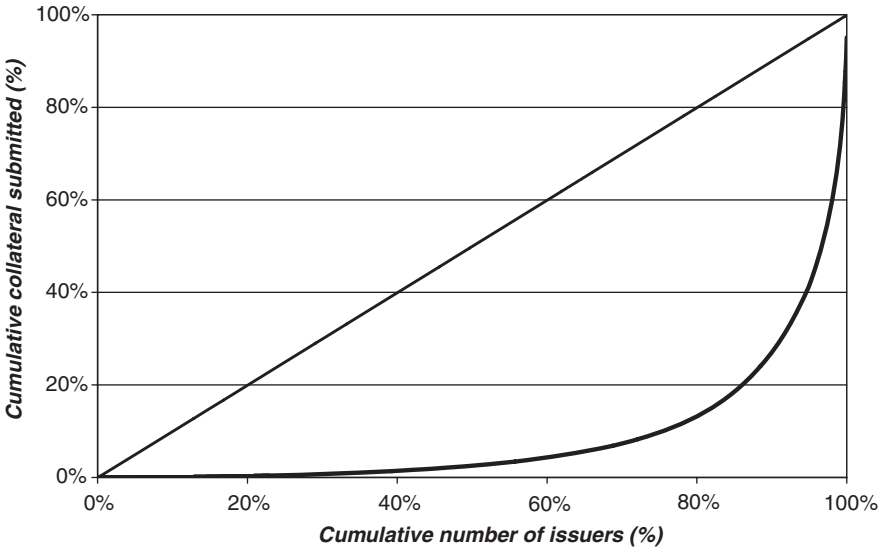

Figure 10.3 Lorenz curve for collateral issuers with respect to amount of collateral submitted.
Source: own calculations.

counterparties, exclusively belong to this industry. Government bonds follow with 26.3 per cent, covered bonds with 18.4 per cent, ABS with 12 per cent and corporate bonds with 7.5 per cent.

Covered bonds which are as well issued by banks also bear an exposure to the banking sector. Depending on the legal framework governing the covered bonds, this exposure may vary. For example, in covered bond legislations where the cover asset pool is bankruptcy-segregated from the holding entity, the insolvency will not necessarily trigger the acceleration of the covered bonds and therefore the exposure to the banking industry is quite low. By contrast, if there is no bankruptcy segregation of cover assets, that is, the cover pool is not separated from other bank assets but the covered bondholders have a superior preferential right to these assets, an insolvency of the issuing bank also triggers the acceleration of the covered bond. In such a case, the exposure to the banking industry is higher. Another important link of covered bonds to the banking industry relates to the composition of the cover pool: the level of voluntary over-collateralization and the quality of the assets in a cover pool which are both an important quality feature for a covered bond depend on the current situation of a bank. In case a bank runs into financial difficulties, it might not be willing to put more than the required quantity and quality of assets in its cover pool.

The main industry exposure of covered bonds as well as of ABSs, however, depends on the assets forming the cover pool. For covered bonds, this might be either mortgage loans or government bonds. According to statistics from the European Covered Bond Council,[12] around half of the covered bonds outstanding on a European level are mortgage covered bonds, and half are public sector covered bonds. These figures might serve as an indication for the ultimate industry exposure the Eurosystem is exposed to with respect to the covered bonds that are submitted as collateral.

With regard to ABSs that constitute 12 per cent of the total amount of submitted collateral, a clear attribution to specific industries is more difficult. It would be necessary to decompose the asset cover pools of each submitted ABS in order to know the exact industry exposure. Around 80 per cent of the ABSs submitted to the Eurosystem can be clearly attributed to the real estate sector.

Finally, the last significant asset category is corporate bonds. When having a closer look at the industrial sectors forming part of this asset category, it turns out that for 71 per cent of the total amount submitted in the form of corporate bonds, the issuers are part of the financial industry (like investment companies, insurance companies, etc.). Other sectors are: communications (6 per cent), utilities (6 per cent), industrial (5 per cent), governmental agencies (4 per cent) and consumer (4 per cent).

Putting all these pieces of information together, a rough indication on the ultimate sector exposure of the Eurosystem in its collateral operations can be provided. There are three dominating sectors: banks/financials (42 per cent), governments (36 per cent) and real estate (18 per cent).

As regards the risk implications of collateral concentration, the following can be concluded. Overall, the degree of concentration on the level of collateral can currently be considered as high. Although the concentration by collateral issuers is in principle slightly lower than in the case of counterparties, this fact cannot be satisfactory per se. In particular the high sector concentration of Eurosystem collateral is remarkable. While the Eurosystem collateral framework is designed as such in order to allow a wide range of collateral from various industries, in reality three sectors dominate the sector exposure: banks, government and real estate. The high exposure to the banking sector is especially noteworthy because on the counterparty side the structurally given sector exposure is also to the

[12] According to these statistics, as of end 2005, there were EUR 885.6 billion mortgage covered bonds outstanding and EUR 865.5 billion public sector covered bonds outstanding (see www.hypo.org).

banking industry. Therefore, there is overall a high exposure to the banking industry.

To take account of these findings in the residual risk estimations, a higher level of correlation between and across counterparties and issuers could be assumed in the residual risk estimations.

4.3 Concentrations in collateral from a single counterparty

Besides looking at concentrations within the group of counterparties and the group of issuers separately, possible concentrations at the level of single counterparties with respect to the usage of assets have also risk implications. As already mentioned, one base assumption underlying the residual risk estimations is that the portfolio of assets that each counterparty submits to the Eurosystem is perfectly fine-grained and therefore idiosyncratic risk has been fully diversified away. In case this assumption might not hold true, a risk adjustment is necessary.

The analysis in this section focuses on the concentration by collateral issuers for each single counterparty. Concentrations by a single counterparty as regards collateral issuer can be illustrated by the Herfindahl–Hirschmann Index (HHI).[13] It takes the value one if the submitted collateral of one counterparty is concentrated on only one collateral issuer and zero (in the limit) if the submitted collateral is equally distributed to a very high number of collateral issuers. The HHI can be calculated for each single counterparty and takes the form

$$HHI = \frac{\sum_{i=1}^{n} Collateral_i^2}{\left(\sum_{i=1}^{n} Collateral_i\right)^2}$$

with i being the different collateral issuers from a counterparty.

To illustrate the calculation of HHI, assume bank X has submitted EUR 100 million from issuer A, EUR 200 million from issuer B, and EUR 250 million from issuer C. The HHI for bank X is then

$$HHI = \frac{100^2 + 200^2 + 250^2}{(100 + 200 + 250)^2} = 0.3125$$

[13] While the Gini coefficient is a measure of the deviation of a distribution of exposure amounts from an even distribution, the HHI measures the extent to which a small number of collateral issuers account for a large proportion of exposure. HHI is related to exposure concentration and therefore the appropriate concentration measure in this context.

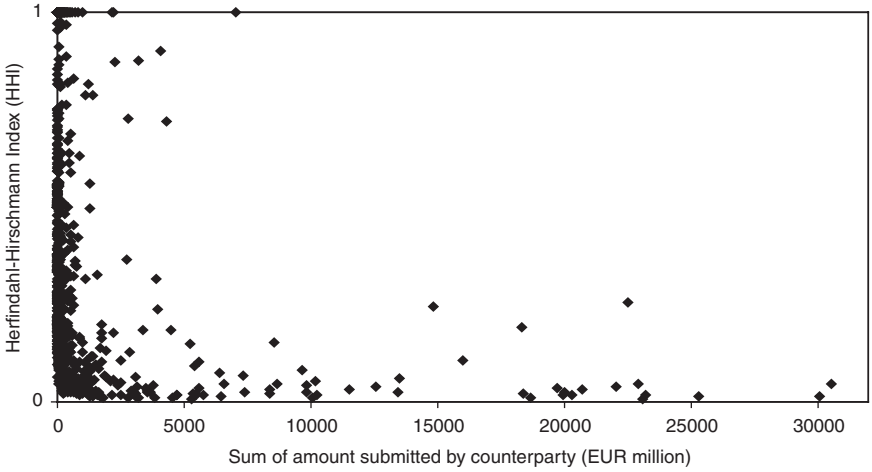

Figure 10.4 Herfindahl–Hirschmann Indices (HHI) of individual counterparties with respect to their collateral submitted.
Source: own calculations.

This index is calculated for all counterparties that submit assets to the Eurosystem. The results are presented in Figure 10.4 in relation to the sum of amount submitted by each counterparty. The average HHI of all counterparties – weighted by their respective sum of amount submitted – is around 0.119. To take account of this concentration in collateral from single counterparties in the risk estimations, a granularity adjustment can be made.[14]

For the purposes of this analysis, a granularity adjustment is approximated following the simplified approach as described in Wilkens *et al.* 2001 for all the counterparties submitting assets to the Eurosystem. According to this approach, the Credit Value-at-Risk ($CVaR_n$) of a portfolio can be decomposed into two components: the $CVaR_\infty$ resulting from a perfectly diversified portfolio and a factor (β*HHI) that accounts for granularity, whereby β is a constant depending on PD and loss given default (LGD), taking the form

$$\beta = (0.4 + 1.2^{*}LGD)^{*}(0.76 + 1.1^{*}PD/F)$$

F is a measure of the systematic risk sensitivity. It takes the form

$$F = N(a_1^{*}G(PD) + a_0) - PD$$

[14] For more details on the calculation of a granularity adjustment, see Gordy 2003; Gordy and Lütkebohmert 2007; BCBS 2001a; BCBS 2001b; Wilkens *et al.* 2001.

where a_0 and a_1 are constants that depend only on the exposure type. For corporate, bank, and sovereign exposures, the values of these coefficients were determined within the IRB granularity adjustment calculations as: a_0=1.288 and a_1=1.118. These values will also be used in this context. G (PD) denotes the inverse cumulative distribution function for PD. Given a PD of 10 basis points, F takes the value of 0.014. Given F and making an assumption on the average LGD, β can be calculated. Assuming for example an average recovery rate of 40 per cent (and hence an LGD of 60 per cent), β takes a constant value of 0.94. Then, the granularity adjustment for each counterparty can be easily calculated if its HHI is known.

For the calculation of a granularity adjustment a constant PD of 10 basis points and a constant recovery rate of 40 per cent (or respectively, an LGD of 60 per cent) are assumed for all the assets submitted by counterparties. Since the granularity adjustment is – following the simplified approach as described above – a linear function of the HHI, an average granularity adjustment can be easily calculated by multiplying the average HHI with the β obtained using a PD of 10 basis points and a LGD of 60 per cent. This results in an average granularity adjustment of around 11 per cent. Technically, in the residual risk estimations the granularity adjustment will be taken into account in the credit risk component of the ES calculations.

As regards the risk implications of concentrations in collateral from a single counterparty, the following can be concluded: Overall, there is a high variety among counterparties as regards their collateral concentration. While some counterparties submit a highly diversified collateral pool to the Eurosystem, there is a sizeable amount of counterparties with collateral pools that are very little diversified.

To include these findings in the residual risk estimations, a granularity adjustment for credit risk could be taken into account.

5. Risk measures: Credit Value-at-Risk and Expected Shortfall

Like Value-at-Risk for market risk, Credit Value-at-Risk (Credit VaR) is defined as a certain quantile of the (portfolio) credit loss distribution. For example, a 99 per cent Credit VaR is defined to be the 99th percentile of the loss distribution. Since Credit VaR is defined with respect to a loss distribution (negative losses mean profits), losses are located in the right tail of the distribution, i.e. in the upper quantiles. Expected Shortfall (ES), also known as Conditional VaR or Expected Tail Loss, at a given confidence level

α per cent is defined as the expected value of losses exceeding the α per cent VaR, or equivalently the expected outcome in the worst $(1-\alpha)$ per cent of cases.

To be more precise, let $x \in R^d$ denote a random variable with a positive density $p(x)$. For each decision vector ξ, chosen from a certain subset ψ of R^n, let $h(\xi,x)$ denote the portfolio loss random variable, having a distribution in R, induced by that of x. For a fixed ξ the cumulative distribution function for the portfolio loss variable is given by

$$F(\xi,\theta) = \int_{h(\xi,x)\leq\theta} p(x)dx = P\{x \leq \theta\}$$

and its inverse is defined as

$$F^{-1}(\xi,\omega) = \min\{\theta : F(\xi,\theta) \geq \omega\}$$

The VaR_α and ES_α values for the loss random variable associated with ξ and a specified confidence level α are given by

$$VaR_\alpha(\xi) = F^{-1}(\xi,\alpha)$$

and

$$ES_\alpha(\xi) = (1-\alpha)^{-1} \int_{h(\xi,x)\geq VaR_\alpha(\xi)} h(\xi,x)p(x)dx$$

As mentioned above, VaR_α is the α-quantile of the portfolio loss distribution and ES_α gives the expected value of losses exceeding VaR_α.

As is well known, unlike VaR, ES is a coherent risk measure.[15] One of the main problems with VaR is that, in general, it is not sub-additive which implies that it is possible to construct two portfolios A and B such that VaR(A+B) $>$ VaR(A) + VaR(B). In other words, the VaR of the combined portfolio exceeds the sum of the individual VaRs, thus discouraging diversification. Another feature of VaR which, compared to ES, makes it unattractive from a computational point of view, is its lack of convexity. For these reasons, ES is used as the preferred risk measure here.

Generally, the approaches used for the computation of risk measures can be classified either as fully- or semi-parametric approaches. In a *fully*-parametric

[15] Artzner *et al.* (1999) call a risk measure *coherent* if it is transition invariant, positively homogeneous, sub-additive and monotonic with relation to stochastic dominance of order one.

approach the portfolio loss distribution is assumed to follow some parametric distribution, e.g. the normal- or t-distribution, based on which the relevant risk measures (VaR and ES) can be easily estimated. For example, in the case the portfolio loss random variable $h(\xi, x)$ would follow a normal distribution, which is widely applied as the basis for market VaR calculations (Gupton *et al.* 1997), the 99 per cent VaR would be calculated as

$$VaR_{99\%}(h) = \mu_h + N^{-1}(0.99)\sigma_h \approx \mu_h + 2.33\sigma_h,$$

where μ_h and σ_h denote the mean and volatility of the losses.

In a *semi*-parametric approach the portfolio loss distribution is not explicitly known. However, what is known instead is the (multivariate) probability distribution of the random variable x, driving the portfolio losses, and the mapping to the portfolio losses $h(\xi, x)$. In many cases, as is also done here, the estimation of the portfolio risk measures has to be done numerically by utilizing Monte Carlo (MC) simulation techniques.

The estimation of portfolio α-VaR by plain MC simulation can be achieved by first simulating a set of portfolio losses, organizing the losses in increasing order and, finally, finding the value under which $\alpha \times 100\%$ of the losses lie. ES can then be calculated by taking the average of the losses exceeding this value.

However, this sorting-based procedure fails in case the generated sample points are not equally probable, as happens e.g. when a variance reduction technique called importance sampling is used to improve the accuracy of the risk measure estimates. Fortunately, there exists an alternative way to compute VaR and ES simultaneously, that is also applicable in the presence of arbitrary variance reduction techniques.

Rockafellar and Uryasev (2000) have shown that, ES_α (with confidence level α) can be obtained as a solution of a convex optimization problem

$$ES_\alpha(\xi) = \min_{\nu \in R} \nu + (1-\alpha)^{-1} \int_{x \in R^d} [h(\xi, x) - \nu]^+ p(x) dx, \tag{10.1}$$

where $[z]^+ = \max\{z, 0\}$, and the value of ν which minimizes equation (10.1) equals VaR_α. An MC-based estimate for ES_α and VaR_α is obtained by generating a sample of realizations for the portfolio loss variable and by solving

$$\overline{ES}_\alpha(\xi) = \min_{\nu \in R} \nu + (1-\alpha)^{-1} \frac{1}{N} \sum_{i=1}^{N} [h(\xi, x_i) - \nu]^+.$$

This problem can easily be solved either by formulating the above problem as a linear program, as in Rockafellar and Uryasev (2000, 2002), which requires introducing N auxiliary variables and inequality constraints to the model, or by directly solving the one dimensional non-smooth minimization problem e.g. with a sub-gradient algorithm; (see Bertsekas 1999 or Nesterov 2004 for more information on sub-gradient methods).

These expressions are readily applicable in optimization applications where the optimization is also performed over the portfolio holdings ξ and the objective is to find an investment portfolio which minimizes the portfolio ES_α (see Rockafellar and Uryasev 2000, 2002; Andersson *et al.* 2001).

Another complication related to the specifics of credit risk estimation is that credit events are extremely rare so that without a highly diversified portfolio none will occur outside the e.g. 1 per cent tail and thus 99 per cent VaR will be zero. ES on the other hand also accounts for the magnitude of the tail events and is thus able to capture the difference between credit exposures concentrated far into the tail.

As an example, consider a portfolio consisting of n different obligors with 4 basis point PD, which corresponds to a rating of AA–. Assuming for simplicity that the different debtors are uncorrelated, it can be calculated with elementary probability rules how many obligors the portfolio should contain in order to obtain a non-zero VaR:

$$\begin{aligned} P(\textit{at least one obligor defaults}) &= 1 - P(\textit{none of the obligors default}) \\ &= 1 - 0.9996^n > 0.01 \Leftrightarrow n > 25 \end{aligned}$$

Including the effect of correlation would mean that the number of obligors should be even higher for VaR to be positive.

6. An efficient Monte Carlo approach for credit risk estimation

The simplest and the best-known method for numerical approximation of high-dimensional integrals is the Monte Carlo method (MC), i.e. random sampling. However, in the literature of numerical integration, there exist many techniques that can be used to improve the performance of MC sampling schemes in high-dimensional integration. These techniques can be generally classified as *variance reduction techniques* since they all aim at reducing the variance of the MC estimates. Most widely applied variance reduction techniques in financial engineering include importance sampling

(IS), (randomized) quasi-Monte Carlo (QMC) methods, antithetic variates and control variates; see Glasserman 2004 for a thorough introduction to these techniques. Often the use of these techniques substantially improves the accuracy of the resulting MC estimates, thus effectively reducing the computational burden required by the simulations; see e.g. Jäckel (2002), Glasserman (2004) and the references therein. In the context of credit risk (rare event) simulations the application of importance sampling usually offers substantial variance reductions, but combining IS with e.g. QMC sampling techniques can improve the computational efficiency even further.

Sections 6.1 and 6.2 provide a brief overview and motivation for the variance reduction techniques applied in this study: importance sampling and randomized QMC methods. Section 6.3 demonstrates the effectiveness of different variance reduction techniques in the estimation of credit risk measures, in order to find out which combination of variance reduction techniques would be well suited for the residual risk estimation in the Eurosystem's credit operations.

6.1 Importance sampling

Importance sampling is especially suited for rare event simulations. Roughly speaking, the objective of importance sampling is to make rare events less rare by concentrating the sampling effort in the region of the sampling space which matters most to the value of the integrand, i.e. the tail of the distribution in the credit risk context.

Recall that the usual MC estimator of the expectation of a (loss) function h

$$\mu = \int_{x \in R^d} h(\xi, x) p(x) dx = E[h(x)]$$

where $x \in R^d$ is a d-dimensional random variable with a positive density $p(x)$ is

$$\hat{\mu} = \frac{1}{N} \sum_{i=1}^{N} h(x_i)$$

where N is the number of simulated sample points.

When calculating the expectation of a function which gets a non-zero value only under the occurrence of a rare event, say a default with a probability e.g. 0.04 per cent, it is not efficient to sample points from the

distribution of x, since the majority of the generated sample points will provide no information about the behaviour of the integrand. Instead, it would seem intuitive to generate the random samples so that they are more concentrated in the region which matters most to the value of the integrand. This intuition is reflected in importance sampling by replacing the original sampling distribution with a distribution which increases the likelihood that 'important' observations are drawn.

Let g be any other probability density on R^d satisfying $p(x) > 0 \Rightarrow g(x) > 0, \forall x \in R^d$.

Then μ can be written as

$$\mu = E[h(x)] = \int_{x \in R^d} h(\xi, x) \frac{p(x)}{g(x)} g(x) dx = \tilde{E}\left[h(x) \frac{p(x)}{g(x)}\right]$$

where $\tilde{E}$ indicates that the expectation is taken with respect to the probability measure g, and the MC estimator of μ is given by

$$\overline{\mu} = \frac{1}{N} \sum_{i=1}^{N} h(x_i) = \frac{1}{N} \sum_{i=1}^{N} h(\tilde{x}_i) \frac{p(\tilde{x}_i)}{g(\tilde{x}_i)}$$

with $\tilde{x}_i, i = 1, \ldots, N$ independent draws from g and the weight $p(\tilde{x}_i)/g(\tilde{x}_i)$ is the ratio of the original and the new importance sampling density.

The IS estimator of ES_α and VaR_α can now be obtained by simply solving the problem

$$\overline{ES}_\alpha(\xi) = \min_{\nu \in R} \nu + (1-\alpha)^{-1} \frac{1}{N} \sum_{i=1}^{N} \frac{p(\tilde{x}_i)}{g(\tilde{x}_i)} [h(\xi, \tilde{x}_i) - \nu]^+ \tag{10.2}$$

where $\tilde{x}_i, i = 1, \ldots, N$ are independent draws from the density g.

Successful application of importance sampling requires that:

1) g is chosen so that the variance of the IS estimator is less than the variance of the original MC estimate;
2) g is easy to sample from;
3) $p(x)/g(x)$ is easy to evaluate.

Generally, fulfilling these requirements is not at all trivial. However, in some cases e.g. when $p(x)$ is a multivariate normal density – that will be used in the credit risk estimations in Section 7 – these necessities are easily met. For a more detailed discussion on IS with normally distributed risk factors and its applications in finance; see Glasserman (2004) and the references therein.

The density $p(x)$ of a d-dimensional multivariate normal random variable x, with mean vector θ and covariance matrix Σ, is given by

$$p(x) = \frac{1}{(2\pi)^{d/2} Det(\Sigma)^{1/2}} \exp\left(-\frac{1}{2}(x-\theta)^T \Sigma^{-1}(x-\theta) \right)$$

With normally distributed risk factors, the application of IS is relatively straightforward. Impressive results can be obtained by choosing the IS density, g, as a multivariate normal distribution with mean vector $\hat{\theta}$ and covariance matrix Σ, i.e. by simply changing the mean of the original distribution; see Glasserman *et al.* (1999).

This choice of g clearly satisfies requirement 2) above, but it also satisfies requirement 3) since the density ratio is simply

$$\frac{p(x)}{g(x)} = \frac{c \exp\left(-\frac{1}{2}(x-\theta)^T \Sigma^{-1}(x-\theta) \right)}{c \exp\left(-\frac{1}{2}(x-\hat{\theta})^T \Sigma^{-1}(x-\hat{\theta}) \right)} = \exp\left((\frac{1}{2}(\hat{\theta}+\theta) - x)^T \Sigma^{-1}(\hat{\theta}-\theta) \right) \quad (10.3)$$

where $c = \frac{1}{(2\pi)^{d/2} Det(\Sigma)^{1/2}}$. As demonstrated in Section 6.3, an appropriate choice of $\hat{\theta}$ effectively reduces the variance of the IS estimator in comparison to the plain MC estimate, thus also satisfying the most important requirement 1).

6.2 Quasi-Monte Carlo methods

QMC methods can be seen as a deterministic counterpart to the MC method. They are deterministic methods designed to produce point sets that cover the d-dimensional unit hypercube as uniformly as possible, see Niederreiter (1992). By suitable transformations, QMC methods can be used to approximate many other probability distributions as well. They are just as easy to use as MC, but they often result in faster convergence of the approximations, thus reducing the computational burden of simulation algorithms. For a more thorough treatment of the topic the reader is referred to Niederreiter 1992 and Glasserman 2004.

It is well known that if the function $h(x)$ is square integrable then the standard error of the MC sample average approximation $\hat{\mu}$ is of order $1/\sqrt{N}$. This means that cutting the approximation error in half requires increasing the number of points by a factor of four. In QMC the convergence rate is $\ln(N)^{d-1}/N$, which is asymptotically of order $1/N$, which is

much better compared to MC. However, this asymptotic convergence rate is practically never useful since even for a very moderate dimension $d = 4$, $N \geq 2.4 \times 10^7$ for QMC to be theoretically superior to MC. Fortunately in many applications, and especially in the field of financial engineering, QMC methods produce superior results over MC and the reason for this lies in the fact that, even though the absolute dimension of the considered problems could be very high, the *effective dimension*, i.e. the number of dimensions that account for the most of the variation in the value of the integrand, is often quite low. For a formal definition of the term effective dimension see e.g. Caflisch *et al.* (1997) and L'Ecuyer (2004). Therefore, obtaining good approximations for these important dimensions, e.g. using QMC, can significantly improve the accuracy of the resulting estimates. These effects will be demonstrated in the next section.

The uniformity properties of QMC point sets deteriorate as a function of the dimension. In high-dimensional integration the fact that the QMC point sets are most uniformly distributed in low dimensions can be further utilized:

1) by approximating the first few (1–10) dimensions with QMC and the remaining dimensions with MC; and
2) by transforming the function h so that its expected value and variance remain unchanged in the MC setting, but its effective dimension (in some sense) is reduced so that the first few dimensions account for the most variability of h.

Detailed descriptions for implementing 2) in case the risk factors are normally distributed, as is done here, are described in L'Ecuyer (2004), and a procedure based on principal component analysis (PCA) (Acworth *et al.* 1997) is also outlined and applied in the following.

The fact that QMC point sets are completely deterministic makes error estimation very difficult, compared to MC. Fortunately, this problem can be rectified by using randomized QMC (RQMC) methods. To enable practical error estimation for QMC a number of randomization techniques have been proposed in the literature; see L'Ecuyer and Lemieux (2002) for an excellent survey. An easy way of randomizing any QMC point set, suggested by Cranley and Patterson (1976), is to shift it randomly, modulo 1, with respect to all of the coordinates. After the randomization each individual sample point is uniformly distributed over the sample space, but the point set as a whole still preserves its regular structure. Randomizing QMC point sets allows one to view them as variance reduction techniques which often produce significant variance reductions with respect to MC in empirical applications, see e.g. L'Ecuyer (2004).

The best-suited combination of the described variance reduction techniques for Credit VaR calculations has to be further specified based on empirical findings.

6.3 Empirical results on variance reduction

This section studies the empirical performance of the different variance reduction techniques by estimating portfolio level ES figures caused by credit events, i.e. defaults, using the following simplified assumptions: The portfolio contains $d = 100$ issuers distributed equally within the four different rating categories AAA, AA, A and BBB. The issuer PDs are all equal within the rating classes and the applied rating class specific default probabilities are presented in Table 10.1. For simplicity, the asset price correlation between every issuer is assumed to be either 0.24 or 0.5. The value of the portfolio is arbitrarily chosen to be 1000 and it is invested evenly across all the issuers. The recovery ratio is assumed to be 40 per cent of the notional amount invested in each issuer. The Credit VaR will be estimated over an annual horizon at a confidence level of 99 per cent.

The general simulation algorithm can be described as follows:

1. Draw a point set of uniformly distributed random numbers

$$U_N = \{u_1, \ldots, u_N\} \subset [0,1)^d$$

2. Decompose the issuer asset correlation matrix as $\Sigma = CC^T$, for some matrix C.

For i = 1 to N:

3. Transform u_i, component by component, to a normally distributed random variable x_i through inversion $x_i = \Phi^{-1}(u_i^j)$, where Φ^{-1} denotes the inverse of the cumulative standard normal distribution and $i = 1, \ldots, N$, $j = 1, \ldots, d$,
4. Set $\tilde{x}_i = Cx_i + \hat{\theta}$, where $\hat{\theta}$ is the mean vector of the shifted density g. Now $\tilde{x} \sim \Phi(\hat{\theta}, \Sigma)$
5. Identify defaulted issuers $Y_j = \mathbf{1}\{\tilde{x}_{j,i} > z_j\}$, $j = 1, \ldots, d$
6. Calculate the portfolio loss $h(\xi, \tilde{x}_i) = c_1 Y_1 + \cdots + c_d Y_d$

End

Find the estimate $\overline{ES}_a(\xi) = \min_{\nu \in R} \nu + (1-\alpha)^{-1} \frac{1}{N} \sum_{i=1}^{N} e^{0.5(\hat{\theta} - \tilde{x}_i)^T \Sigma^{-1} \hat{\theta}} [h(\xi, \tilde{x}_i) - \nu]^+$

In step 1) the point set can be simply an MC sample or a sample generated through an arbitrary RQMC method. In step 2) the most common choice for C is the *Cholesky factorization* which takes C to be a lower triangular matrix.

Another possibility is to select C based on a standard *principal component analysis* (PCA) which concentrates the variance, as much as possible, to the first coordinates of x, with the aim of reducing the effective dimension of the problem. This choice yields $C = QD^{1/2}$, where D is a diagonal matrix containing the eigenvalues of Σ in decreasing order, and Q is an orthogonal matrix whose columns are the corresponding unit-length eigenvectors. Even though this technique completely ignores the function h whose expectation we are trying to estimate, it has proven empirically to perform well in combination with (R)QMC techniques; see Acworth *et al.* (1997), Moskowitz and Caflisch (1996) and L'Ecuyer (2004). In step 5) $\mathbf{1}\{.\}$ is an indicator function, in 6) c_i denotes loss given default for issuer i in monetary terms, and the final expression for the ES is obtained simply by combining (10.2) and (10.3) with $\theta = 0$.

The simulation experiments are undertaken with a sample size of N=5000 and the simulation trials are repeated 100 times to enable the computation of error estimates for RQMC methods. The numerical results presented below were obtained using a randomized Sobol sequence,[16] but comparable results were also achieved with e.g. Korobov lattice rules. The accuracy of the different variance reduction techniques are compared by reporting variance reduction factors with respect to plain MC sampling for all the considered methods. This factor is computed as the ratio of the estimator variance with the plain MC method and the variance achieved with an alternative simulation method.

Figure 10.5 illustrates the variance reduction factors achieved with the importance sampling approach described in Section 6.1, where the mean $\theta = 0$ of the d-dimensional standard normally distributed variable x will be shifted to $\hat{\theta}$. The results show that the simple importance sampling scheme can reduce the estimator variance by a factor of 40 and 200 for an asset correlation of 0.24 and 0.50, respectively. Figure 10.5 also illustrates, reassuringly, that the variance reduction factors are fairly robust with respect to the size of the mean shift, but shifting the mean too aggressively can also substantially increase the estimator variance. The highest variance reductions are obtained with the mean shifts equalling 1.3 and 1.8 in case of 0.24 and 0.5 asset correlation assumptions, respectively.

As indicated above, additional variance reduction may be achieved by combining importance sampling with RQMC and dimension reduction techniques. Tables 10.3 and 10.4 report the results of such experiments

[16] See: Sobol (1967).

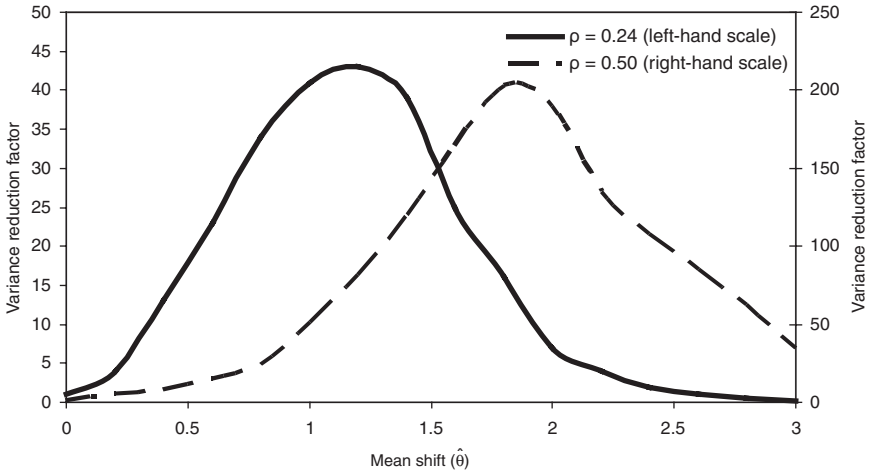

Figure 10.5 Variance reduction factors, for varying values of $\hat{\theta}$ and asset correlations.

under the asset correlation assumption of 0.24 and 0.5, respectively. In addition to MC, three different combinations of variance reduction techniques are considered, namely Monte Carlo with importance sampling (MC+IS), a combination of MC and randomized Sobol sequence with IS (SOB+IS) where the first five dimensions are approximated with Sobol point sets and the remaining ones with MC,[17] and finally SOB+IS combined with the PCA dimension reduction technique to pack the variance as much as possible to the first dimensions which are hopefully well approximated with the Sobol points (SOB+IS+PCA).

The results in Table 10.3 and Table 10.4 show that all the applied techniques produce significant variance reduction factors (VRF) with respect to MC and the VRFs grow substantially as the confidence level increases from 99 per cent to 99.9 per cent. In all experiments SOB+IS+PCA produces the highest VRFs and the effectiveness of the PCA decomposition increases with the asset correlation which reduces the effective dimension of the problem as the asset prices tend to fluctuate more closely together.

The conducted experiments indicate that, among the considered variance reduction techniques, the implementation based on SOB+IS+PCA produces the highest VRFs and therefore this is the simulation approach chosen to derive residual risk estimates for the Eurosystem's credit operations.

[17] Empirical tests indicated that extending the application of Sobol point sets to higher dimensions than five, generally had a detrimental effect on the accuracy of the results. Tests also showed that applying antithetic variates instead of plain MC does not improve the results further. Therefore, plain MC is used for the other dimensions.

Table 10.3 Comparison of various variance reduction techniques with 0.24 asset correlation

	MC	MC + IS	SOB + IS	SOB + IS + PCA
$ES_{99\%}$	16.629	16.612	16.545	16.575
$\sigma^2_{99\%}$	1.809	0.045	0.038	0.035
$\mathrm{VRF}_{99\%}$	1	40	47	51
$ES_{99.9\%}$	33.960	34.300	34.157	34.198
$\sigma^2_{99.9\%}$	41.323	0.239	0.204	0.120
$VRF_{99.9\%}$	1	173	202	343

Table 10.4 Comparison of various variance reduction techniques with 0.5 asset correlation

	MC	MC + IS	SOB + IS	SOB + IS + PCA
$ES_{99\%}$	29.049	29.368	29.418	29.409
$\sigma^2_{99\%}$	21.539	0.106	0.075	0.051
$\mathrm{VRF}_{99\%}$	1	202	286	422
$ES_{99.9\%}$	83.544	85.915	85.935	85.911
$\sigma^2_{99.9\%}$	586.523	0.816	0.381	0.184
$VRF_{99.9\%}$	1	719	1540	3186

7. Residual risk estimation for the Eurosystem's credit operations

This section presents the results of the residual risk estimations for the Eurosystem's credit operations. The most important data source used for these risk estimations is a snapshot on disaggregated data on submitted collateral that was taken in November 2006. This data contains information on the amount of specific assets submitted by each single counterparty as collateral to the Eurosystem. In total, Eurosystem counterparties submitted collateral of around EUR 928 billion to the Eurosystem.

For technical reasons, the dimension of the problem needs to be reduced without impacting the risk calculations. The total collateral amount is spread over more than 18,000 different counterparty–issuer pairs. To reduce the dimension of the problem, only those pairs are considered where the submitted collateral amount is at least EUR 100 million. As a consequence, the number of issuers is reduced to 445 and the number of counterparties is reduced to 247. With this approach, only 64 per cent of the total collateral submitted is taken into account. Therefore, after the risk calculations, the resulting risks need to be scaled up accordingly.

The other assumptions used for the risk estimations were discussed in Sections 2 and 3. In the following, the most important ones are briefly recalled. The annual PDs of the counterparties and issuers are derived from the credit ratings on a second-best basis. These annual PDs are scaled down linearly according to the time it takes to liquidate the least liquid instrument that has been submitted from the issuer. The same liquidation times are used as those applied for the derivation of haircut levels (see Table 10.2).

With regard to the recovery rate in case of an issuer default, a uniform recovery rate of 40 per cent is assumed for all the assets. For the default correlation between and across counterparties and issuers only one uniform level of correlation of 24 per cent is assumed. To take account of granularity in the counterparties' collateral pools, a granularity adjustment of 11 per cent for credit risks is made.[18]

The necessary assumptions for the calculation of liquidity-related risk are the following: as regards the distributional assumption for price movements, a normal distribution for price changes is assumed. Concerning the assumption on price volatility, the same weekly volatility of 1.2 per cent is assumed for all the assets in the collateral pool.

Another important assumption for the risk calculations is that it is assumed that there is no over-collateralization. This means that the amount of submitted collateral equals the amount lent to the bank. Since there is normally some voluntary over-collateralization, this presents a conservative assumption.

Section 7.1 summarizes the results of the residual risk estimations when using (conservative) assumptions under normal conditions. Section 7.2 illustrates some possible developments in risk under 'stress' conditions. Section 7.3 presents an application of the model to show the development in risks over time.

7.1 Expected shortfall in a base case scenario

In the base case scenario, residual risks are at a very low level. As can be seen from Table 10.5, ES at a 99 per cent confidence level is only around EUR 18.8 million for the total amount of assets submitted to the Eurosystem. This corresponds to 0.2 basis points of total lending. As regards the breakdown into different asset categories, by far the biggest share of EUR 13.2 million is allotted to bonds issued by banks. For this asset category, the

[18] Technically, this is done by scaling up the resulting credit risk by a factor of 1.11.

Tabel 10.5 Breakdown of residual risks in the base case scenario

Asset categories	ES (in EUR mn)	ES in relation to total lending (in basis points)
Bank bonds	13.2	0.32
Government bonds	2.1	0.07
ABS	1.3	0.13
Corporate bonds	1.0	0.26
Other	1.2	0.18
All categories	18.8	0.2

Source: own calculations.

ES in relation to the total amount submitted of this asset category is 0.32 basis points.

With respect to the distribution into credit risk and liquidity-related risks, around EUR 16 million are due to credit risk while only EUR 2.8 million result from liquidity-related risks.

These results suggest that the risks stemming from Eurosystem collateral operations are very low. This is on one hand certainly due to the risk mitigation of the risk control framework, but on the other hand, the risks are also at such a low level because the average credit quality of issuers and counterparties can be considered as good.

7.2 Stability of risk calculations in terms of assumptions

This section presents the risk implications of isolated changes of different key input parameters. The charts in Figures 10.6 to 10.8 illustrate these effects for the following cases: i) a change in the liquidation time assumptions; ii) a change in the assumptions on the credit quality of issuers and counterparties; and iii) a change in the assumptions on issuer-counterparty correlations, respectively. All the three charts show the development of ES in basis points of total lending in relation to a change in the respective parameter. All the other relevant input parameters for the risk estimations are chosen according to the base case scenario (see Section 7.1).

A change in the *liquidation time assumptions* has both effects on credit and on liquidity-related risks. If it takes longer to liquidate the instrument, the time that the issuer might default is longer and therefore the credit risk increases. At the same time, a higher liquidation time assumption leads to a higher price volatility and therefore affects the liquidity-related risks.

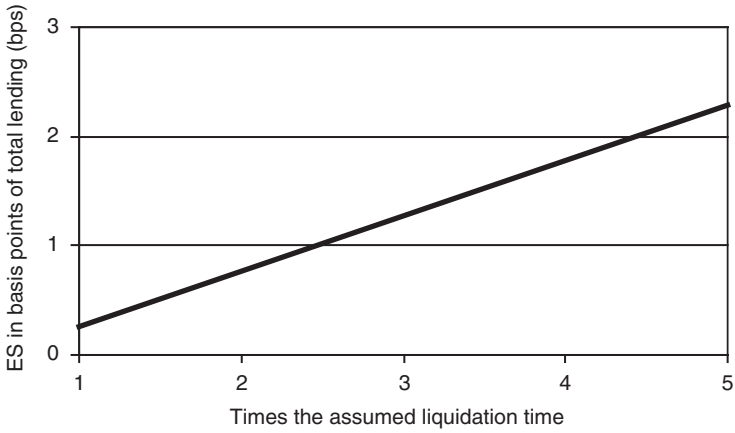

Figure 10.6 The effect on Expected Shortfall of changed liquidation time assumptions. *Source*: own calculations.

In sum, the overall effect of a higher liquidation time is a roughly linear increase in risks. This relationship is also illustrated in Figure 10.6. The values on the x-axis of this chart have to be read as follows: '1' means that it takes exactly the time as defined in the risk control framework to orderly liquidate these assets. '2' means that it takes twice the time as defined in the risk control framework to orderly liquidate them, and so on.

A change in the assumptions on the *credit quality of issuers and counterparties*, measured in changed assumptions on the PDs for both the issuers and the counterparties, has also effects both on credit and liquidity-related risks. Liquidity-related risks increase because of the higher likelihood of a counterparty default. The increase in credit risk is quite obvious. It is both motivated by the higher PDs for counterparties and for issuers. As illustrated in Figure 10.7, a simultaneous increase in the assumption on the PD for the counterparties and issuers results in a roughly quadratic risk increase.

Finally, the risk effects of a change in the assumptions on *issuer-counterparty correlations* are illustrated in Figure 10.8. It shows that risks grow roughly exponentially with increasing correlation. This relationship is especially of interest as regards the impact of high sector concentration in the Eurosystem's collateral that was discussed in Section 4. It shows the possible effects on ES if the sector concentration between counterparties and issuers is high. Moreover, the average issuer–counterparty correlation would have to be adjusted upwards if counterparties submitted a sizeable amount of assets from issuers with which they have close links. Since the increase in ES would be significant for a higher average correlation level, it is appropriate

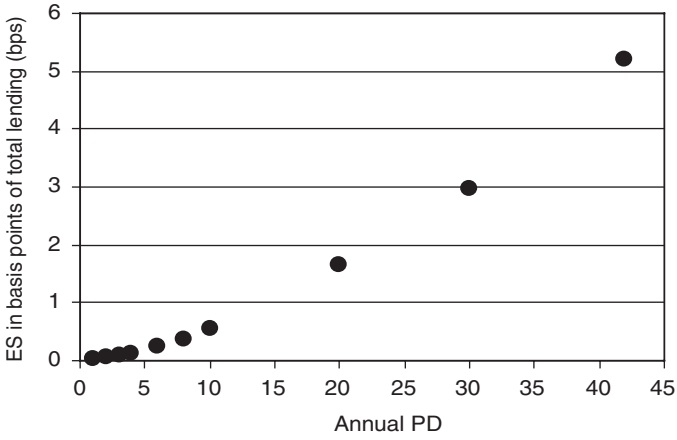

Figure 10.7 The effect on Expected Shortfall of changed credit quality assumptions. *Source*: ECB's own calculations.

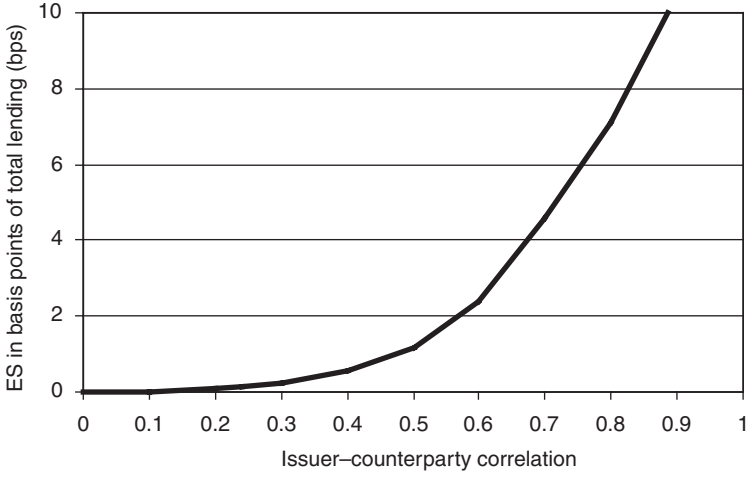

Figure 10.8 The effect on Expected Shortfall of changed assumptions on issuer-counterparty correlations.

that the Eurosystem does in principle not foresee the existence of close links in its collateral operations.

It should be kept in mind that normally one of these three input parameters does not change in an isolated way but that for example a drying-up of liquidity conditions could as well be accompanied by a concurrent deterioration in credit quality of counterparties and issuers. Therefore, the residual risks for Eurosystem credit operations could increase quite dramatically under such circumstances. Such developments can be simulated in stress scenarios.

Tabel 10.6 Composition of submitted collateral over time and composition of residual financial risks over time

	Total submitted collateral (in EUR billion)						Residual financial risks (in EUR million)					
	2001	2002	2003	2004	2005	2006	2001	2002	2003	2004	2005	2006
Bank bonds	337	343	347	379	418	467	10.7	10.9	11.1	12.1	13.3	14.9
Government bonds	269	255	186	311	299	262	1.8	1.7	1.3	2.1	2.0	1.8
ABS				45	83	109				0.6	1.1	1.5
Corporate bonds	21	28	33	28	46	61	0.5	0.7	0.8	0.7	1.2	1.6
Other	61	50	164	55	54	53	1.1	0.9	2.9	1.0	1.0	0.9

7.3 Risk development over time

Finally, the model may be used in order to quantify the development in residual financial risks over time. Table 10.6 illustrates the development and composition of collateral over time. It can be seen that between 2001 and 2006 the total amount of submitted collateral has increased by around 30 per cent. While the amount of government bonds has remained stable over time, especially the more risky asset categories like bank bonds, corporate bonds and ABS are now submitted to a larger extent to the Eurosystem than five years ago.

For the computation of residual risks over time, some simplifying assumptions need to be made since for the previous years no information on a disaggregated level for the submitted collateral was available for the purposes of this study. Therefore, it is assumed that the average credit quality of the counterparties and collateral issuers is the same as the one observed in 2006, and the assumptions for the other important input parameters are the same as the one assumed for the estimation of residual risks for 2006 in the base case scenario (see Section 7.1).

The resulting residual financial risks are also presented in Table 10.6.[19] It can be seen that overall risks grow by 45 per cent. The increase in risks is on one hand due to the higher absolute amount of submitted collateral and on the other hand it is because of the higher share of more risky collateral in the pool. While government bonds for example had a share of around 13 per cent in overall financial risks in 2001, their share has now decreased to less than 9 per cent.

[19] The figures reported for 2006 differ slightly from the figures presented in Section 7.1 since for the computation of residual risks the annual average of submitted collateral was taken, while for the risk estimations in Section 7 the data is based on a one time data snapshot taken in November 2006.

8. Conclusions

This chapter has presented an approach to estimate tail risk measures for a portfolio of collateralized lending operations. The general method was applied to quantitatively assess the residual financial risks for the Eurosystem's collateralized lending operations. The risk measure chosen was ES at the 99 per cent quantile of the loss distribution over an annual horizon. In order to avoid making distributional assumptions on the shape of the credit loss distribution, ES was estimated on a Eurosystem-wide basis by using sophisticated Monte Carlo simulation techniques.

Overall, risk taking from policy operations appears very low. Risk estimations in a base case scenario revealed that ES in relation to the total amount of collateral submitted amounts to only around 0.2 basis points. This corresponds to an absolute exposure of EUR 18.8 million. However, when incorporating trends in collateral use with stressed assumptions, risks are driven up considerably. Especially, a rise in the correlation between issuers and counterparties or a deterioration of average credit quality leads to significant increases in risks.

In view of the size of the Eurosystem's monetary policy operations portfolio, a regular quantitative assessment of the residual risks is necessary in order to check if the collateral framework ensures that the risk taken in refinancing operations is in line with a central bank's risk tolerance. Finally, the quantification of these risks is also an important step towards a more comprehensive and integrated central bank risk management.

11 Central bank financial crisis management from a risk management perspective

Ulrich Bindseil

1. Introduction[1]

Providing emergency liquidity assistance (ELA) or being the lender of last resort (LOLR) are considered to be amongst the most important tasks of central banks, and literature on the topic is correspondingly abundant (see e.g. Freixas *et al.* 1999, for an overview). To avoid confusion relating to the specific definitions given and uses made of these terms in the literature and in the central banking community, this chapter uses the broad concept of 'central bank financial crisis management' (FCM), which encompasses ELA and LOLR. Apart from some important general clarifications of direct usefulness for the practitioner (from Bagehot 1873 to e.g. Goodhart 1999), the literature on FCM takes mainly an academic perspective of microeconomic theory (e.g. Freixas and Rochet 1997; Repullo 2000; Freixas *et al.* 2003; Caballero and Krishnamurthy 2007).[2] While this microeconomic modelling of the functioning of FCM is certainly relevant, doing practical FCM at the right moment and in the right way requires more than that. In particular, it requires three disciplines of applied central banking, namely central bank liquidity management, central bank risk management and prudential supervision. The role of risk management has been stressed recently by W. Buiter and A. Sibert in their internet blog posted on 12 August 2007, just days after the break-out of the financial market turmoil:

[1] I wish to thank Denis Blenck. Fernando Gonzalez, Jose Manuel Gónzalez-Páramo, Paul de Grauwe, Elke Heinle, Han van der Hoorn, Fernando Monar, Benjamin Sahel, Jens Tapking, and Flemming Würtz for useful comments. Remaining mistakes remain mine of course.

[2] See Goodhart and Illing (2002) for a comprehensive panorama of views on financial crises, contagion and the lender-of-last-resort role of central banks.

A credit crunch and liquidity squeeze is ... the time for central banks to get their hands dirty and take socially necessary risks which are not part and parcel of the art of central banking during normal times when markets are orderly. Making monetary policy under conditions of orderly markets is really not that hard. Any group of people with IQs in three digits (individually) and familiar with (almost) any intermediate macroeconomics textbook could do the job. Dealing with a liquidity crisis and credit crunch is hard. Inevitably, it exposes the central bank to significant financial and reputational risk. The central banks will be asked to take credit risk (of unknown) magnitude onto their balance sheets and they will have to make explicit judgments about the creditworthiness of various counterparties. But without taking these risks the central banks will be financially and reputationally safe, but poor servants of the public interest.

This chapter attempts to summarize and structure some key messages from the academic literature, to the extent they seem important for practice. It appears plausible that central bank financial operations in times of crisis imply, or are even inherently associated with, particular risk taking, and that considerations underlying FCM decisions must therefore follow also a risk management philosophy and related technical considerations. As it will become clear, risk management considerations will not only be relevant for the practice of FCM, but also shed a new light on existing academic debate. Noticing that the recent literature does not pay much attention to risk management aspects does not mean that this issue has always been neglected. At the contrary, the founding fathers of the concept of FCM, Thornton (1802 – see quotation in Goodhart 1999, 340), Harman (1832, quoted e.g. in King 1936, 36 – '*We lent ... by every possible means consistent with the safety of the Bank*'), and Bagehot (1873) were all clear that liquidity assistance should only be granted subject to, as one would say today, adequate risk control measures, such as to protect the Bank of England against possible losses. For instance Bagehot (1873) explained:

These advances should be made on *all good banking securities*, and as largely as the public ask for them ... *No advances indeed need be made by which the Bank will ultimately lose.* The amount of bad business in commercial countries is an infinitesimally small fraction of the whole business. That in a panic the bank, or banks, holding the ultimate reserve should refuse bad bills or bad securities will not make the panic really worse.

Besides providing further insights on long-debated FCM issues, studying risk management aspects of FCM is also a crucial contribution to an efficient FCM framework, since once a potential need for FCM occurs, time for analysis will be scarce, and facts will be complex and opaque. In such a

situation, it can only help if all risk management policies and procedures for FCM have been thought through, been documented internally and understood, even if one would take the view that the policies should neither be mechanistic, nor fully known to the outside world to prevent moral hazard.

The rest of this chapter proceeds as follows. Section 2 provides a typology of FCM cases, to clarify the subject of this paper, and since a great amount of confusion often exists in public debate on what FCM exactly is. Section 3 summarizes, mainly from a risk management perspective, a number of key conclusions of the FCM literature. Section 4 argues that a first crucial central bank contribution to financial stability lies in the normal operational framework of the central bank. Section 5 develops the 'inertia' principle of central bank risk management under crisis situation, which provides the bridge between central bank risk management under normal circumstances and FCM actions. Section 6 discusses, again largely from a risk management perspective, FCM providing *equal* access to all central bank counterparties to some exceptional form of liquidity provision. Section 7 discusses ELA/LOLR to single banks under conditions granted only to this bank. Section 8 summarizes and draws some key conclusions.

2. Typology of financial crisis management measures

Often, the literature and public debate on FCM is astonishingly imprecise on what type of FCM operations it has in mind, although this makes a big difference in every respect. Hence, conclusions seem often naïve or wrong – again in the summer of 2007. It is for instance difficult to find a clear commonly agreed distinction in the literature between ELA and LOLR. Here, it is therefore simply assumed that the two are equivalent, and only the expression 'ELA' is used. The following typology encompasses all FCM measures of central banks. It essentially distinguishes between three types of measures whereby the first type is further subdivided. Consider each type of FCM measures in more detail.

A Equal access FCM measures

This type of FCM measures is equally addressed to all banks with which a central bank normally operates. It may also be called 'ex ante' FCM or 'preventive' FCM, as it is undertaken *before* single banks need specific liquidity help due to actual illiquidity.

A-I *Aggregate excess liquidity injection through normal open market operations (OMOs).* Demand for excess reserves is normally rather limited in modern financial systems with reliable payment systems (see Bindseil *et al.* 2006). However, under financial stress, banks tend to demand excess liquidity, and if the central bank would not offer it, at least interbank short-term rates would go up. Therefore, at various occasions (Bank of England 1825, Y2K transition, days after 11 September 2001, second half of 2007), central banks have massively injected reserves to satisfy exceptional demand for excess reserves and to thereby calm down the market and contribute to avoid potential illiquidity of some banks.

A-II *Narrowing the spread of the borrowing facility vis-à-vis the target rate.* In case that liquidity and/or infrastructure problems force banks to use extensively the borrowing facility, the central bank may want to alleviate associated costs by lowering the penalty rate applied on the borrowing facility. The ECB did so for instance for the first two weeks of the euro (in January 1999), and the Fed in August 2007.

A-III *Widening of collateral set.* While ELA to individual banks typically involves accepting unusual collateral (since otherwise the central bank borrowing facility could do the job), it can also be conceived that the collateral set is widened on an *aggregate* basis, i.e. maintaining equal access of all central bank counterparties. For instance, in case of a liquidity crisis and lack of collateral, the central bank could increase available collateral by accepting a class of paper normally not accepted (e.g. Equity), or it could lower rating requirements for an existing asset class (require a BBB rating instead of an A rating). In case of two separate collateral sets for open market operations (narrow set) and standing facilities (wider set), a measure could also be to accept exceptionally the latter set for open market operations.

A-IV *Other special liquidity supplying operations to address a liquidity problem in the market.* For instance, after 11 September 2001, the ECB provided dollar liquidity to euro area banks through a special swap operation (see press release of ECB dated 13 September 2001) and in December 2007 it again provided US dollars, but this time against euro-denominated collateral (see press release of 12 December 2007). Typically, international banks operate in several key markets (e.g. USD, EUR), and have collateral in each market, but there are normally no bridges between the collateral sets, increasing the vulnerability to liquidity shocks.

B Individual access FCM measures

Once a *single* bank can no longer fulfill its payment duties, or anticipates that this state becomes unavoidable in the near future, unless it obtains help, it may turn to the central bank to obtain the necessary help. This is normally associated with the terms 'ELA' or 'LOLR', although these terms are also used sometimes in a looser way, encompassing equal access FCM measures. ELA as understood here may also be called an *ex post* FCM measures, because it is provided after a liquidity problem has concretely materialized. Giving up the principle of equal treatment is of course something considered undesirable, and will have to be counterbalanced by making the bank(s) concerned pay some high price, also for incentive / moral hazard reasons. Emergency liquidity assistance to single banks often cannot be separated from potential solvency assistance, since (i) liquidity problems tend to correlate with solvency; (ii) the reasons to support banks through ELA are not fundamentally different from those to also save insolvent banks. As an old financial stability saying goes: 'A bank alive is worth more for society than a dead bank, even if its net worth is negative.'[3]

C Organize emergency/solvency assistance to be provided by other financial institutions.

The most important motivation for the central bank's (or another public agent's) ELA or solvency assistance to individual banks is the negative externality of a bank failure (see the next section). Therefore, the failure of the banking system to address this on its own may be interpreted as a Coasean social cost problem (Coase 1960), in which transaction costs (due e.g. to asymmetric information, high number of agents, etc.) preclude economic actors to agree on the socially optimal outcome. The 'state' (in this case the central bank) can address such failures either by acting itself, or by contributing to reduce the transaction costs of reaching a private agreement, such that negative externalities are avoided. Indeed, in the case of individual ELA or solvency assistance, often central bank action consists in bringing a group of financial institutions together which would bear, together, a large part of the costs of the failure, and by making them agree to support the bank, which is eventually also in their interest.

[3] Solvency assistance by the Government, including e.g. nationalization by the Government, may also be considered to fall under this type of FCM measures. However, this chapter does not elaborate on these cases in detail.

In the rest of this paper, the above-proposed typology will be followed. Deviating from other attempts to typologize FCM measures (e.g. Daniel *et al.* 2004, 14 for the Bank of Canada), the typology proposed here does *not* consider the following two measures as ELA measures on their own. First, *lowering the short-term monetary policy target interest rate* seems to be only indirectly linked to liquidity problems. Indeed, it does not address them, but may only be associated with them because liquidity problems may worsen the macroeconomic outlook and thus create deflationary pressures (and thereby justify rate cuts), or the reason for liquidity problems may be a macroeconomic shock that is at the same time the justification of a lowering of interest rates from the monetary policy perspective. Second, *communicating* in itself is also not considered a direct financial crisis management measure. For instance Daniel *et al.* (2004, 14) explains that 'in times of heightened financial stress, the Bank can also reinforce its actions through public statements that indicate that the Bank stands ready to ensure the availability of sufficient liquidity in the financial system'. This was done for instance by the Fed after the stock market crash of 12 October 1987, or by both the Fed and the ECB as an immediate reaction on 11 September 2001. However, verbal communications are not measures on their own, but only relevant as far as associated with possible implied central bank actions. In so far, each type of action described above has its corresponding communication dimension. Table 11.1 summarizes again the proposed typology, including examples from the period July to December 2007.

3. Review of some key results of the literature

This section reviews some of the main conclusions of the FCM literature. As far as relevant, it takes a risk management perspective, although for some of the issues, there is no such specific perspective. Still, it was deemed useful to summarize briefly how these key issues are understood here.

3.1 Key lessons retained from nineteenth-century experience

The origins of modern ELA/FCM theories are usually traced back to Bagehot (1873), and, as Goodhart (e.g. 1999) remarks, to Thornton (1802). The main conclusions of these authors are summarized (again, e.g. by Goodhart 1999) as (i) lend freely in times of crisis; (ii) do so at a high interest rate; (iii) but only against good collateral/securities. Amongst the

Table 11.1 FCM typology and illustration from August–December 2007

(A) Equal access FCM measures ('ex ante FCM')	(A-I) Inject aggregate excess liquidity through OMOs. *Example: ECB injects EUR 95 billion on 9 August 2007 through a spontaneous fixed rate tender with full allotment at the overnight target rate.*
	(A-II) Reduce penalty associated with standing facilities. *Example: Fed lowers on 17 August 2007 the discount rate by 50 basis points, such as to half the spread vis-à-vis the Fed founds target rate.*
	(A-III) Widening collateral set. *Examples: (1) Bank of Canada announces on 15 August 2007 that it will accept the broader collateral set for its borrowing facility also for open market operations. (2) On 6 September 2007, the Federal Reserve Bank of Australia announces that it will accept ABCPs (asset backed commercial paper) as collateral. (3) On 19 September the BoE announces to accept MBS paper for special open market operations. (4) On 12 December 2007, the Fed announces that it will conduct for the first time in history open market operations (reverse repos) against the broad set of collateral eligible for discount window operations; At the same time, Fed accepts a wide set of counterparties for the first time in an open market operation; (5) On the same date, the Bank of Canada and the Bank of England widen again their collateral sets for reverse repo open market operations.*
	(A-IV) Non-standard operations, including cross-currency. *On 12 December 2007, the Swiss National Bank and the ECB announce that they would provide USD funds for 28 days against collateral denominated in euro (and Swiss Franks).*
(B) Individual access FCM (ELA). *On 14 September 2007, the Bank of England provides ELA to Northern Rock PLC.*	
(C) Organize emergency/solvency assistance to be provided by other financial institutions. *Implemented in Germany for IKB on 31 July 2007, and for Sachsen LB on 17 August 2007.*	

classical quotes on the topic, one is in fact due to a central banker, and not to Bagehot: an eventual wilful massive injection of liquidity in a financial panic situation by the Bank of England in 1825 was summarized in the words of Bank member Jeremiah Harman in the Lords' Committee in 1832 (quoted from King 1936, 36, but also to be found in Bagehot 1873):

> We lent ... by every possible means, and in modes that we never had adopted before; we took in stock of security, we purchased Exchequer bills, we made advances on Exchequer bills, we not only discounted outright, but we made advances on deposits of bills to an immense amount; in short, by every possible means consistent with the safety of the Bank; ... seeing the dreadful state in which the public were, we rendered every assistance in our power.

Three details are noteworthy in this statement, which maybe have been noted too little in the recent literature. *First*, Harman distinguishes explicitly between collateralized lending to banks, and outright purchases of securities, which have rather different risk properties, and different implications with regard to the dictum 'lend at high prices', since this seems to apply potentially only to collateralized lending ('advances'). *Second*, the liquidity injection is not only against good banking securities, but 'in modes never adopted before', and '*by every means consistent with the safety of the bank*'. In other words, the only constraint was a central bank risk management constraint, but not an a priori constraint on the types of securities to be accepted. The quotation seems to suggest that finding these unusual modes to inject liquidity *with limited financial risk* for the central bank was considered the key challenge in these operations. *Third*, Harman does not mention as a crucial issue the 'high rate of interest', and indeed, as supported by some recent commentators (e.g. Goodhart 1999, 341), a high level of interest rates charged should not be the worry in such circumstances, not even from an incentive point of view.

It seems noteworthy that earlier statements on the principles of financial crisis management are mainly about *aggregate* liquidity injection into the financial system under circumstances of a collective financial market liquidity crisis (i.e. case A in the typology introduced in Section 2), and not about ELA in the narrow sense, i.e. support to an individual institution which has run into trouble due to its specific lack of profitability or position taking (type B of liquidity provision). Most authors today writing on ELA or LOLR do not note this difference, i.e. they start with Thornton or Bagehot when introducing the topic, but then focus on individual access FCM measures.

Today, the set of eligible assets for central bank borrowing facilities tends to be rather wide, and access is not constrained in any other way than by collateral availability. So one could say that for the type of FCM Thornton and Bagehot had in mind, central banks have gone a long way to incorporate them into a well-specified and transparent framework. Moreover, it is undisputed amongst central bankers today that liquidity absorption due to autonomous factor shocks (see e.g. Bindseil 2004, 60) should be fully neutralized through open market operations.

3.2 The nature of liquidity problems of banks

A key feature of banks is that their assets are largely illiquid (in particular under stressed financial market conditions), while their liabilities comprise

to a considerable extent short-term deposits by non-banks (and to a lesser extent by other banks) and short-term paper (see e.g. Freixas *et al.* 1999, or Freixas and Rochet 1997). This feature makes them susceptible to a run: if there is a belief that the bank has a solvency problem, customers may start rushing to withdraw their deposits. Since the first-come-first-served principle applies, rumours that a bank has problems are sufficient to trigger a run and to become self-fulfilling. The same will hold for unsecured interbank lending: a bank which is considered to have an increased likelihood of defaulting for whatever reason will see other banks cutting down rapidly credit lines to this bank. The bank run problem has been modelled first by Diamond and Dybvig (1983), with a very extensive subsequent literature (see e.g. Caballero and Krishnamurthy 2007, for a recent related paper). The breakdown of the interbank market may be considered a classical adverse selection problem (see e.g. Flannery 1996).

As long as the bank has high-quality collateral which is accepted easily in interbank markets (e.g. Government fixed-income securities), it should always be in a position to obtain liquidity. Such collateral is typically also the kind of asset that could be sold outright relatively easily and without losses at short notice, even if markets are stressed. Only if such assets are exhausted, may the liquidity situation of the bank may turn critical. In principle, it may be able to sell or pledge more credit risky assets at short notice, but only at a substantial discount, to reflect that those who are supposed to take them over quickly cannot analyse their value at short notice (implying a 'conservative' valuation). A substantial discount in asset fire sales may also put the solvency of the bank further at risk, aggravating the bank's crisis. The 'correlation' between liquidity problems and solvency problems is thus based on a two-sided causality, which may develop its own dynamic. First, a bank with solvency problems of which depositors or other banks become suspicious will also quickly face liquidity problems since depositors or short-term paper holder will not want to prolong their investments due to the perceived high credit risk. Second, a bank with liquidity problems which needs to do asset fire sales to address these problems will see its net value deteriorate due to the losses associated with the discounts to be made to achieve the fire sales.

Daniel *et al.* (2004, 7) highlight that widespread liquidity problems of banks may also arise out of pure extra caution in view of some external developments: when lending institutions become concerned *that their own sources of liquidity may be less reliable than usual*, banks may reduce the volume of funds that they lend in the interbank market, setting up a

situation of self-fulfilling expectations of liquidity hoarding, even if there is no credit quality issue at the origin. Only indirectly, through the causality mentioned above, will credit problems come into play. The eventual victims of the liquidity crisis may thus not be the ones that triggered it by cutting lines in the interbank market.

For central bank FCM, the following may be concluded, taking the risk management perspective. Liquidity problems of banks will be related in one or two causal ways to credit problems. Therefore, if the central bank wants to help, it will also have to get involved in credit risk taking, whereby it can and should do its best to limit actual additional risk taking through careful risk management. Subject to its policy objectives, it should minimize additional risk taking through adequate haircuts, and sufficiently precise (i) credit quality assessments and (ii) valuation of collateral.

3.3 Motivations for central banks FCM, and in particular ELA

In the following, a number of potential reasons for justifying FCM measures by the central bank are provided. Of course, in some cases, more than one reason is relevant, but each of them is in principle sufficient. Maybe these motivations are not so relevant for the FCM measures of type A-I and A-II, and more for the really substantial FCM measures of types A-III and B, which are about accepting non-standard types of collateral. A-I and A-II could be seen to be mainly about steering short-term interest rates, and be somehow less substantial in terms of the central bank 'getting its hands dirty', so they may need less of additional motivation – i.e. they could be simply motivated by continuing to control short-term interest rates. Anyway, their case, including justifications going beyond short-term interest rate control, will be developed further in Sections 6.1 and 6.2.

3.3.1 Negative externalities of illiquidity (and bankruptcy)

The central bank may be ready to engage in FCM measures because of the potential negative externalities for society of letting a bank(s) go down, relating to 'systemic risk' and knock-on effects associated with bank failures (see e.g. Freixas *et al.* 1999, 154–7). These negative externalities would affect a lot of players in the market, so whoever of these players would risk their money to save a bank from illiquidity, would appropriate only a small part of the avoided negative externalities, but bear the full risk and expected cost. In theory, affected market players could coordinate to internalize the full social cost, but as Coase (1960) explained, this does not need to happen due

to transaction cost. In so far, a situation would not be unlikely in which no single bank is willing to help, but the central bank, comparing *social* cost and benefit of FCM measures, would conclude that it should help. The externality argument also applies in the case where it is certain that a bank also has solvency problems, since for society, 'A bank alive will be worth more than a dead bank, even if net worth of the bank is negative.' The negative externality of a real bankruptcy relates in particular to the two following issues: (i) Difficulties and large transaction costs to continue managing the banks claims in a way to minimize the further losses to its asset value. In any case, asset values will suffer further from the discontinuity in their management. If a bank is rescued or taken over, and only the senior management and other main culprits for the failure are replaced, continuity of operations can be ensured to a higher extent. (ii) Contagion risks: as the failed bank stops to honour its liabilities (until it is sorted out how much its lenders can actually receive after liquidation of all assets, which takes time) and if assets lose more of their value, other banks, corporate claimants and private investors may become illiquid or insolvent, implying further costs to society and the danger of widespread knock-on effects.

From a central bank risk management perspective, it is to be highlighted that a social cost–benefit analysis requires to quantify the costs of risk taking (or expected losses) of the central bank arising in most FCM cases. Without estimation of risk, collateral valuation, and an analysis of the role of risk control measures (e.g. haircuts), the economic (or 'social') cost–benefit analysis of FCM measures would remain fundamentally incomplete, in particular for A-III and B, when the collateral base is enlarged.

3.3.2 Central bank is only economic agent not threatened by illiquidity

Central banks have been endowed with the monopoly and freedom to issue the so-called legal tender: central bank money. Therefore, central banks are never threatened by illiquidity, and it seems natural that in case of a liquidity crisis, in which all agents rush towards securing their liquidity, the central bank remains more willing than others to hold (as collateral or outright) assets which are not particularly liquid. This should be the case even if the central bank would not have the public mandate to act such as to avoid negative externalities as described in the previous section. The price mechanism, leading to depressed prices of illiquid assets in a liquidity crisis, should itself, everything else equal, be sufficient to have the central bank

move more into illiquid assets. Of course, it could be argued that the central bank does not have sufficient expertise to assess the value of illiquid assets in a crisis. However, as far as collateral is concerned, this issue is mitigated by another central bank specificity as explained in section 3.3.5.

3.3.3 Superior knowledge of the central bank

One may imagine that due to supervisory functions, or close relation to the banking supervisor, a central bank has more knowledge than the market, and knows that a bank (or banking system) with liquidity problems is solvent, such that emergency liquidity assistance can be done with little or no risk, in view of the collateral the bank can offer. Other banks in contrast, on the basis of their inferior knowledge, see a high probability that the bank also has solvency problems, and therefore do not want to lend any further to the bank, or at least not against the collateral that the bank can provide (see e.g. Freixas *et al.* 1999, 153–4). Berger *et al.* (1998) test this possibility empirically and conclude that shortly after supervisors have inspected a bank, supervisory assessments of the bank are more accurate than those of the market, but that the opposite is true if the supervisory information is not recent. This argument may possibly also be restated as follows: once a bank has problems and gets in touch with the central bank and supervisor to ask for help, these public institutions can always update their supervisory information quickly. This is also the case since the bank will be willing to fully open its books to the central bank and supervisor, trusting that it will help, and that it will not misuse the information. In contrast, it would not be willing to open its books and share all information with a competitor, who it can expect to seek only its own advantage. In other words, not only its status as powerful regulator, but also the one of a fair broker interested in the social good allows the central bank to obtain more information than competitors of the bank in trouble could. The central bank risk management perspective on this motivation for FCM measures is clear, as the motivation is based on superior risk management relevant knowledge of the central bank, which in this case may be seen to overlap to a significant extent with prudential supervisory expertise.

3.3.4 Superior ability of the central bank to secure claims

Finally, a third independent sufficient reason for FCM measures could be that the central bank has for legal or other reasons more leverage on a

bank in liquidity and/or solvency problems; i.e. after provision of liquidity against some non-standard collateral, it can ensure, better than banks could, that it will be paid back. The central bank is generally in a powerful position vis-à-vis banks as it has influence on supervisory decisions, including the eventual bank's closure. Also, the central bank may have access to legal tools to make a collateral pledge more effective than the market ever could. Financial risk management typically relies on the assumption of legal reliability of claims, and leaves doubts on this to operational risk management experts and lawyers. Still, once the financial risk manager knows what can be relied upon and what is non-secure, it will consider this as crucial input in its financial risk analysis and in composing optimal risk mitigation measures. If it is true that the central bank can better secure its claims than others, it will also be able to better manage related financial risks.

3.3.5 Haircuts as powerful risk mitigation tool if credit risk is asymmetric

Haircuts are a powerful tool to mitigate liquidation risk of collateral in the case of the default of the cash taker (i.e. collateral provider) in a repo operation – *however only if the cash taker is more credit risky than the cash lender*. Indeed, in case of a haircut, the cash taker is exposed to the risk of default of the cash lender, since in case of such default, she is uncertain to get her collateral back. For instance if she received EUR 100, but had to provide collateral for EUR 110 due to haircuts, she has an unsecured exposure of EUR 10 to the cash lender. This is why haircuts between banks of similar credit quality tend to be low, while banks impose potentially high haircuts if they lend cash to e.g. hedge funds. This also explains why banks would never question haircuts imposed by the central bank, or at least would never consider a high haircut to be a reason not to borrow from the central bank for credit risk considerations. Therefore, the central bank will be able to use (possibly quite high) haircuts as an effective tool to mitigate risk, and will be able to lend to credit risky counterparties without taking disproportional risks. These credit risky counterparties will be happy to accept the haircuts because they do not have to fear a default of the central bank. They would be far less happy to accept such haircuts when imposed by other banks on them, as these banks themselves have non-negligible default risk in a crisis. Therefore, adverse selection and rationing phenomena are likely to lead to a breakdown of the interbank repo market as far as less liquid collateral and more risky banks are involved (e.g. Ewerhart and Tapking, 2008).

3.4 ELA provided by other banks, coordinated by the supervisor or the central bank

Quite often, individual bank FCM (ELA) and solvency assistance have not been provided by the central bank, but by banks, however, on the basis of a collective deal engineered by the supervisor and/or the central bank (case C of the typology of liquidity functions proposed in Section 2). This was the case for the majority of recent cases of ELA, namely LTCM, BAWAG, IKB and Sachsen LB. Why can't banks get together themselves, but need a broker to come to such agreements which are in their own interest, and why are public authorities suitable to be the broker for such deals? For instance LTCM was prevented from going bankrupt by fourteen banks agreeing to supply funds to allow LTCB to continue operating and to achieve a controlled unwinding and settlement of its operations. The Fed brought together the banks, but did not contribute money itself. The banks were made better off by this agreement, and indeed ex post even did not lose any of their money. Still, no single bank had incentives to rescue LTCM on its own, since the private expected benefits of that did not exceed the private expected costs. The main reason for the public authorities to take the initiative to achieve a Coasean agreement amongst banks to rescue another bank is that it should be trusted by all players, having no self-interest, but ideally only the social good in mind. If being trusted, all players will with a higher chance wilfully reveal relevant information on facts and on their true preferences, such as to overcome at least partially the well-known problems of bargaining under asymmetric information (see e.g. Myerson and Satterthwaite 1983 as a model of these mechanism design issues). The public authorities will broker a fair agreement on the basis of the information it has collected, and not reveal private information more than necessary. A further reason may be the power of the supervisor and central bank to persuade potential free riders to participate. Even if a bank which is asked to contribute to a liquidity or solvency rescue feels that it can fully trust the supervisor and central bank as brokering agents, it may prefer to find excuses not to participate to save costs and avoid risks. However, the public authorities may be in a position to exert strong 'moral suasion', as any bank may fear that making the central bank its 'enemy' will backfire in the future.

Supervisory and risk management expertise seems essential for the supervisor and central bank to be respected as catalysts, since the banks will be invited to provide money and thereby to take risks. Also the central bank

needs to be able to understand this financial risk taking, also to be sure that it is reasonable and fair to invite banks to take it.

3.5 Moral hazard

FCM measures by central banks may be regarded as an insurance against adverse outcomes, and, as any insurance, it would weaken incentives to prevent bad outcomes to occur. Incentives issues are not generally seen to prevent insurance from being useful for society, although in some cases, it may be a reason to not insure some activity, which, in a world without any moral hazard, would be ensured. Also, incentive issues are generally something taken into account when designing insurance contracts (e.g. by foreseeing that the insured entity takes a part of the possible losses). Thus the negative impact of moral hazard on the welfare improvements from insurances is not just a given fact, but can be minimized through intelligent institutional design. Still, different schools on applying this to FCM can be found in the literature. For instance de Grauwe (Financial Times, 12 August 2007, 'ECB bail-out sows seeds of crisis') suggests that the liquidity-injecting open market operations by the ECB done in early August 2007 created a moral hazard dilemma. Another school represented by e.g. Humphrey (1986) or Goodfriend and Lacker (1999) arguing that only aggregate equal access emergency liquidity injections (though open market operations) are legitimate from an incentive point of view, while any individual access ELA would be detrimental overall to welfare because of its moral hazard inducing effects. Other authors like Flannery (1996), Freixas (1999) Goodhart (1999) and most central bankers would in contrast argue that also ELA may be justified (namely as explained in Section 3.3). The fact that an insurance always has some distorting effects is not a proof that there should be no insurance, in particular if bank failures have huge negative externalities. This does not mean that there should always be a bail-out, but in some cases there should. Also it does not mean that the bail-out should be complete in the sense of avoiding losses to all stakeholders – quite the contrary (see below).

Another widespread view is that moral hazard and incentive issues imply that crisis-related liquidity injections 'should only be given to overcome liquidity problems as a result of inefficient market allocation of liquidity'[4], and in particular not in the case of true solvency problems. However, as

[4] See also Goodhart (1999, 352–3). For instance Goodfriend and Lacker (1999) seem to take this conservative view, and also Sveriges Riksbank (2003, 64).

argued above, the welfare considerations on solvency problems are not so different from those regarding liquidity problems. In both cases, a rescue scheme may be useful ex post, and also the incentives issues are similar after all. So in fact, for both, one needs to accept that looking for the right design of the 'insurance contract' is essential. In both cases, it seems important to not take away incentives to monitor banks to those who are in a position to do so. Moreover, it is difficult to distinguish insolvency and illiquidity ex ante.

Goodhard (1999, 352), for whom the claim that 'moral hazard is everywhere and at all times a major consideration' is one of the four myths that have been established around ELA, distinguishes four groups of stakeholders with regard to the incentives they need to be confronted with. The most important decision makers and thus those most directly responsible for liquidity or solvency problems are the **senior management** of the bank. Therefore, often a bank requesting support from the central bank has to sacrifice one or several of its responsible Board members. This was for instance also the case in the most recent bank rescue operation, namely of IKB in Germany in August 2007, in which both the CEO and the CFO had to quit within two weeks. Goodhart (1999, 353) considers it a failure that the executives of LTCM were never removed after the rescue of this hedge fund, exactly for these incentive reasons. He notes that unfortunately 'the current executives have a certain monopoly of inside information, and at times of crisis that information may have particular value'. Second, **equity holders**, which have the highest incentives to ask the management to run a high-risk strategy with negative externalities, should obviously suffer. Third, also **bond holders**, and probably interbank market lenders should suffer if equity holders cannot absorb losses, as these groups should still be made responsible for monitoring their exposures. Fourth, it is according to Goodhart (1999, 354) today considered as socially wasteful to require ordinary **small depositors** to monitor their bank, and that some (though preferably not 100 per cent) deposit insurance for those would be justified.

To become more precise, consider quickly the moral-hazard relevance of all types of FCM measures:

- *A-I: OMO aggregate liquidity injection.* This is supposed to be for the benefit of all banks and the system, which has fallen or risks falling into an inferior equilibrium, as a whole. Therefore, referring to moral hazard would imply a collective responsibility of the system, which is difficult to see. If banks have collectively embarked in reckless activity, then the banking and financial market supervision, and/or the legislator have not done their job either, and it may be difficult to hold responsible the

competitive banking system which was led towards this by the environment set or tolerated by the legislator. In any case, in times of money market tensions, interbank interest rates move to levels above the target rate set by the central banks, if the profile of the liquidity supply by the central bank is unchanged. Therefore, it is natural for a central bank to re-adjust the profile of liquidity supply to demand in order to stabilize short-term interbank rates around their target. => ***Moral hazard considerations hardly relevant***.

- *A-II: lowering the penalty level associated with the borrowing facility*. Same arguments. => ***Moral hazard considerations hardly relevant***.
- *A-III: widening the collateral set*. Admittedly, this type of measure could have moral hazard implications, particularly (1) if the general widening of the collateral set in fact targets a small number or, say, even a single bank under liquidity stress, which is rich in the specific type of additional collateral; and (2) if the central bank offers facilities, for instance a standing borrowing facility, to effectively refinance such collateral. This type of action invites moral hazard as it may indeed be decisive to establish whether the single bank fails or not, while at the same time sparing to the banks' management and shareholders the substantial costs associated with resorting to real emergency liquidity assistance. Therefore, a widening of the collateral set accepted for monetary policy purposes should probably only be considered if this measure would substantially help a significant number of banks and if the set of assets were very narrow. In this case the lack of collateral obviously seems to be more of a systemic issue, and the central bank should consider taking action.[5] => ***Moral hazard is an issue, in particular if problems are limited to a few banks***.
- *B: individual ELA*. Moral hazard is in principle an issue, but it can be addressed by ensuring that shareholders and managers are sanctioned. Being supported individually by the central bank always comes at a substantial cost to the bank's management, which can expect that the central bank / supervisor will ask for the responsible persons to quit, and which will control closely the actions of the bank for some while. Also the reputation damage of a liquidity problem is substantial (as soon as the support becomes public). In so far, it seems implausible to construct a

[5] The Institute of International Finance (2007) proposes the following guiding principle with regard to moral hazard: As a principle, central banks should be more willing to intervene to support the market and its participants and be more lenient as to the type of collateral they are willing to accept, if the crisis originates outside the financial industry.

scenario in which the bank deliberately goes for much more liquidity risk than would be optimal from the point of view of society. This does not mean that the calculus of the bank is untouched by the perspective to be bailed out. But probably the distortion remains weaker than the one that might be caused by solvency aid – as far as the two are clearly distinct. Also in the case of solvency aid, the authorities should ensure, to the extent possible ex ante and ex post, that in particular shareholders and senior managers suffer from losses. => ***Moral hazard is an issue, but can be addressed to a significant extent***.

- *C: Public authorities as catalysts for peer institution's help*. Again, the public authorities and the helping institutions can and should ensure that shareholders and senior management are sanctioned. => ***Moral hazard is an issue, but can be addressed to a significant extent***.

In sum, it is wrong to speak *generally* about moral hazard associated to FCM measures, since it makes a big difference what type of FCM measure is taken. An individual ELA (and solvency aid) framework can be designed with a perspective to preserve to the extent possible the right incentives, the optimum having to be determined jointly with the prudential supervision rules. In the optimum, some distortions will still occur, as they occur in almost any insurance or agency contract. Recognizing the existence of these distortions is not a general reason for concluding that such contracts should not exist at all. In the case of individual ELA, the concrete issue of incentives may be summarized as stated by Andrew Crocket (cited after Freixas *et al.* 1999, 161): 'if it is clear that management will always lose their jobs, and shareholders their capital, in the event of failure, moral hazard should be alleviated'. For equal access widening of collateral, moral hazard issues are potentially tricky, and would deserve to be studied further. Risk management expertise of the central bank is relevant in all this because for all measures except A-I and A-II, asset valuation, credit quality assessment and haircut setting are all key to determine to what extent the different measures are pure liquidity assistance, and when they are more likely to turn out to also consist of solvency assistance. In the latter case, moral hazard issues are always more intense that in the former.

3.6 Constructive ambiguity

'Constructive ambiguity', which is a term due to Corrigan, is considered one important means to limit moral hazard. The idea is that by not establishing any official rules on FCM measures, banks will not count on it

and their incentives to be prudent will not be weakened. Still, ex post, the central bank may help. Already Bagehot (1873, chapter 7) touches on the topic, and taking the perspective of the market, criticizes the ambiguity surrounding the Bank of England's FCM policies (whereby it needs to be admitted that this refers to equal access FCM measures, and not to what today's debates mainly have in mind, which is individual ELA):

> Theory suggests, and experience proves, that in a panic the holders of the ultimate Bank reserve (whether one bank or many) should lend to all that bring good securities quickly, freely, and readily. By that policy they allay a panic; by every other policy they intensify it. *The public have a right to know whether the Bank of England the holders of our ultimate bank reserve acknowledge this duty, and are ready to perform it. But this is now very uncertain.*

Apparently central bankers remained unimpressed by this claimed 'right' of the public, and still 135 years after Bagehot, the financial industry continued expressing the wish for more explicitness by central banks (Institute of International Finance 2007, 42):

> Central banks should provide greater clarity on their roles as lenders of last resort in both firm-specific and market-related crises ... Central banks should be more transparent about the process to be followed during extraordinary events, for example, the types of additional collateral that could be pledged, haircuts that could be applied, limits by asset type (if any), and the delivery form of such assets.

Freixas (1999) proposes an explicit model of the role of constructive ambiguity, in which he shows that **mixed strategies**, in which the central bank sometimes bails out and sometimes does not, can be optimal when taking into account the implied incentive effects. In mixed strategies, a player in a strategic game randomizes over different options applying some optimal probabilities, and it can be shown that such a strategy may maximize the expected utility of the player (in this case the central bank, the utility of which would be social welfare; see e.g. Myerson (1991, 156) for a description of the concept of mixed equilibrium, or any other game theory textbook). In so far, constructive ambiguity could be considered as reflecting the optimality of mixed strategies. As it may however be difficult to make randomization an official strategy (as this would raise legal problems), hiding behind 'constructive ambiguity' may appear optimal. Another interpretation of constructive ambiguity could be that it is **a doctrine to avoid legal problems**: if there would be a clear ELA policy, the central bank would probably be forced to act accordingly, and not become exposed to

legal proceedings. In addition, even if it follows in principle its policies, cases may be so complex that legal proceedings may always be opened against the central bank in order to blame it for losses that eventually occurred.

In particular three central banks worldwide have made attempts to be transparent on their ELA policies, thereby rejecting in principle the doctrine of constructive ambiguity, namely Sveriges Riksbank (2003, 58), Hong Kong Monetary Authority (1999), and Bank of Canada (Daniel *et al.* 2004).[6] Of course, all of these three central banks do not promise ELA under any conditions, they only specify necessary conditions for ELA, such that they could in principle still randomize over their actual actions. The specifications they provide are to a substantial part focused on risk management.

From a **risk manager's perspective**, one would generally tend to be suspicious of the concept of constructive ambiguity. It is the sort of concept which triggers alarm bells in the risk manager's view of the world, since it could be seen to mean non-transparency, discretion and deliberate absence of agreed procedures (see also e.g. Freixas *et al.* 1999, 160). In the risk manager's perspective in a wider sense, the following doubts could be expressed vis-à-vis constructive ambiguity, if it is supposed to mean that no philosophy, principles and rules have been studied *even internally* in the central bank. First, it could be interpreted to reflect a lack of thinking by central banks, i.e. an inability to formulate clear contingent rules on when and how to conduct FCM measures. Second, it will imply longer lead times of actual conduct of FCM measures, and more likely wrong decisions. Third, constructive ambiguity would concentrate power with a few senior management decision makers, who will not be bound to policies (as such policies would not exist, at least not in an as clear way as if they were written down) and will have more limited accountability.

If constructive ambiguity is supposed to mean that a philosophy, principles and rules *should exist, but are kept secret*, then the following points could still be made. First, it could still seem to be the opposite of transparency, a value universally recognized nowadays for central banks, in

[6] For instance Sveriges Riksbank (2003, 58) explains: 'Some central banks appear unwilling to even discuss the possibility of possible LOLR operations for fear that this could have a negative effect on financial institutions' behaviour, that is to say, that moral hazard could lead to a deterioration in risk management and to a greater risk taking in the banking system. The Riksbank on the other hand, sees openness as a means of reducing moral hazard . . . A well reasoned stance on the issue of ELA reduces the risk of granting assistance un-necessarily . . . [and is] a defence against strong pressure that the Riksbank shall act as a lender of last resort in less appropriate situations.'

particular if large amounts of public money are at stake.[7] Second, it could be argued that this approach would reiterate an old fallacy from macroeconomics, namely the idea that one can do things ex post, but as long as they are not transparently described ex ante, they will not affect behaviour ex ante. This fallacy was overcome in macroeconomics by the theory of rational expectations. From the perspective of rational expectations theory, it could be argued that having a policy but trying to be non-transparent on it eventually does not mean that policies will not be taken into account ex ante, but will be taken into account in a more 'noisy' way since market players will estimate the 'reaction function' of the central bank under more uncertainty. Noise in itself is however unlikely to be useful. Finally, constructive ambiguity could still imply delays in implementing measures, since even if the central bank is internally well prepared and has pre-designed its reaction function as far as possible, banks requesting FCM measures would be likely to be much better prepared if they knew in advance the relevant rules.[8]

Generally, it could be argued that constructive ambiguity is the opposite from what the regulators expect from banks, namely to have well-documented risk taking policies in particular in crisis situations, and to do stress testing and to develop associated procedures. Risk management becomes most important in stress situations, and it is counter-intuitive to say that exactly for these situations, no prior thinking should take place. Prior thinking does not mean believing in the possibility to anticipate every detail of the next financial crisis, but only by the belief that one will be in a much better position to react, compared with the case of no preparation at all.

3.7 At what rate to provide special lending in a crisis situation?

This issue has been debated for a long time, as it is part of the Bagehot legacy. Interestingly, both Thornton (1802) and Harman in 1832 were less

[7] Full transparency in the middle of a crisis and associated rescue operations may also be harmful, and information on banks accessed by the central bank may be confidential. Ex post, a high level of transparency appears desirable as a key element of accountability of public authorities operation with public resources.

[8] This is also the opinion expressed by the industry in Institute of International Finance (2007, 42): 'there is a fear that greater transparency on the part of central banks would lead to moral hazard. It is the Special Committee's belief, however, that the benefits of increased clarity on how central banks would respond to different types of crises outweigh this risk. In times of crisis involving multiple jurisdictions and regulators, there will always be challenges in the coordination of information collection, sharing, and decision making. To the extent possible, the more protocol that is established prior to such an event, the better prepared both firms and supervisors will be to address a crisis.'

explicit on this. Debates on it are linked to the incentive / moral hazard issue discussed above. The key passage in Bagehot (1873, 197 – see also Goodhart 1999) is:

> The end is to stay the panic; and the advances should, if possible, stay the panic. And for this purpose there are two rules: First. That these loans should only be made *at a very high rate of interest*. This will operate as a heavy fine on unreasonable timidity, and will prevent the greatest number of applications by persons who do not require it. The rate should be raised early in the panic, so that the fine may be paid early; that no one may borrow out of idle precaution without paying well for it; that the Banking reserve may be protected as far as possible.

First, it is important to recall one more time that Bagehot referred to the case of equal access FCM, not to what is mostly debated today, namely individual bank ELA. Second, it may be noted that today, central banks offer borrowing facilities, typically at +100 basis points relative to the target rate, i.e. at some moderate penalty level, but that even this penalty level is apparently considered too high, since central banks e.g. in August 2007 injected equal access emergency liquidity through open market operations almost at the target level, instead of letting banks bear a 100 basis points penalty. So this would not at all have been in the sense of Bagehot (1873). Without saying that it was necessary to shield banks in August 2007 from paying a 100 basis point penalty for overnight credit, it is difficult also to believe that a 100 basis point penalty would have been very relevant in terms of providing incentives. For aggregate FCM measures, the topic simply does not appear overly relevant, at least not in terms of providing or not the right incentives for banks. A general liquidity crunch in the money market is anyway a collective phenomenon, which may have been triggered only by the irresponsible behaviour of a few participants, or even by completely exogenous events. Therefore, collective punishment (anyway only by small amounts) does not make too much sense.

The same seems to hold true for single access ELA: single access ELA implies a lot of problems for a bank and its stakeholders, and this is how it should be (as argued above). Also, in expected terms, ELA often means subsidization of banks, since ELA tends to correlate with solvency problems. The rate at which an ELA loan is made to a bank is in this context only a relatively subordinated issue, which will not decide on future incentives. For the sake of transparency of financial flows, it would probably make sense to set the ELA rate either at a market rate for the respective maturity (in particular if one is confident that there will be enough 'punishment' of the

relevant stakeholders anyway), or at say +100 basis points, which is a kind of penalty level, but none which would compensate for the risks.[9]

4. Financial stability role of central bank operational framework

The setting-up of a central bank's risk management framework for collateralized lending operations is discussed in-depth in Chapters 7 and 8. When designing a framework and setting thereby the eventual amount of eligible collateral and the implied borrowing potential of the banking system with the central bank, the relevance of this for financial stability needs to be recognized. Financial markets are inherently unstable as the prospect of a systemic crisis may be self-fulfilling (as has been modelled in various ways ever since Diamond and Dybvig 1983). Bank runs (both by non-banks and by other banks) are inherently dynamic, and in principle the state of the money market can switch from smooth and liquid to tense and dry up from one moment to the next, as observed at least since Bagehot (1873) and as experienced most recently in August 2007. Potential reasons (more substantial than sunspots) which could trigger a liquidity crisis are always out there in the market, whereby the intensity of these potential triggers for a crisis can be thought as a stochastic process across time. The central bank's normal credit operations framework is key in deciding what the critical level of stress will be before an actual liquidity crisis breaks out. For instance, it is plausible that the critical stress level will depend on the following five dimensions.

(1) **Availability of collateral for central bank credit operations**. It will be stabilizing that: (i) the set of collateral eligible for central bank credit operations is wide; (ii) amounts of collateral available to the banks are large, in comparison to average needs with regards to central bank credit operations, and needs at high confidence levels; (iii) collateral buffers are well-dispersed over the banking system; (iv) risk control measures imposed by the central bank such as limits and haircuts are not overly constraining (e.g. avoidance of limits). The collateral and risk control framework may, or may not be differentiated across the three different types of central bank credit operations (open market

[9] The Hong Kong Monetary Authority (1999, 79) puts emphasis on the idea of a penalty rate: 'The interest rate charged on LOLR support would be at a rate which is sufficient to maintain incentives for good management but not at a level which would defeat the purpose of the facility, i.e. to prevent illiquidity from precipitating insolvency.'

operations, borrowing facility, intra-day). If they are somehow differentiated, then also the impact on money market stability needs to be differentiated across these three dimensions – even if 'the more the better' will somehow hold across all of these dimensions.

(2) **Existence of an end-of-day borrowing facility at a moderate penalty rate and which does not stigmatize**. Most leading central banks impose a penalty level of 100 basis points on their borrowing facility. This is not really an important cost factor for a bank if we talk about covering exceptional short-term liquidity needs even of a considerable size. For instance a EUR 1 billion loan overnight at a 100 basis point penalty means penalty costs of EUR 28 thousand. If a bank takes EUR 10 billion over ten days, it would thus mean a cost of EUR 2.8 million. For a medium- or large-sized banks, this is literally peanuts compared with the potential damage caused by a run on the bank (or a cutting of credit lines from other banks causing a liquidity squeeze). What will really be relevant will be the availability of collateral and the certainty perceived that the information on a large access to the borrowing facility will be kept secret (unless the bank decides voluntarily on an outing).[10]

(3) **A high number of financial institutions has direct access to the central bank borrowing facility**. Wide access to open market operations may also be relevant in so far as it is considered that open market operations have their own merits in terms of contributing to financial stability. A wide range of counterparties is relevant since one core characteristics of a money market crunch is that due to general mistrust and scarcity, no one lends, not even at a high price. Therefore, the central bank becomes the lender of last resort for all financial institutions which lack liquidity, and it does not help a financial institution to know that others could borrow from the central bank and pass on the liquidity to itself, as exactly this will not happen.

(4) **The intra-day payment system is well designed to avoid deadlocks** – for which also limits in intra-day overdrafting and possible collateral requirements are important. There is an extensive literature on this topic see e.g. CPSS (2000).

[10] An important issue in this context is how close the borrowing facility is to an emergency facility. In the US, the discount window had been understood before 2003 as being something in between a monetary policy instrument and an automated emergency liquidity facility. In contrast, the Eurosystem's facility had been designed from the outset more as a monetary policy tool, as suggested by (i) the identity of the collateral set with the one for open market operations; (ii) the absence of any quantitative limitation; (iii) the absence of any follow-up investigations by the central bank.

(5) **Reserve requirements and averaging**. At least initially, i.e. in the US until the early 1930s, reserve requirements were considered as a means to ensure financial stability (see e.g. Bindseil 2004 chapter 4 and the literature quoted there). Although this is no longer seen to be an essential motivation of reserve requirements, the existence of reserve requirements must be relevant in some sense for financial stability. It may be argued that obtaining central bank reserves to fulfil reserve requirements drains available collateral, which is negative for financial stability. On the other side, averaging allows to buffer away moderate end-of-day liquidity shocks without recourse to a borrowing facility, which may be positive if recourse to the borrowing facility is stigmatized. The averaging functionality also allowed the Eurosystem to inject massively excess reserves into the system in the second half of 2007, as it knew it could absorb these excess amounts again in the second half of the reserve maintenance period.

In sum, a first crucial contribution to financial stability that the central bank can provide lies in its normal-times credit operations and collateral framework. When designing this framework, the central bank should always also have financial stability in mind, and not only risk mitigation of the central bank in normal times. This will bias the framework in the directions indicated above, and should be based on some cost–benefit analysis, although estimating benefits is very difficult, since it is not observable how many liquidity crises are avoided by building in one or other stability-enhancing features into the framework.

5. The inertia principle of central bank risk management in crisis situations

In the previous section, it was argued that the operational framework for central bank credit operation was a first major contribution a central bank can make and should make to financial stability. In particular, a wide range of eligible collateral (to be made risk-equivalent through risk control measures) is crucial in this respect. In the subsequent section, equal access FCM measures will be discussed. In between, something very fundamental needs to be introduced, which is called here the 'inertia principle' of central bank risk management. The inertia principle says that the central bank's risk management should not react to a financial crisis in the same way as banks' risk managers should, namely by restricting business such as to limit the

additional extent of risk taking. Instead, the central bank should maintain its risk control framework at least inert, and accept that its risk taking will therefore rise considerably in a crisis situation. While central bank risk management is normally conservative and reflects the idea that, probably, the central bank is a moderately competitive risk manager compared to private financial institutions,[11] it becomes an above-average risk taker in crisis situations – first of all by showing *inertia* in its risk management framework. There is thus some *fundamental transformation* occurring because the central bank continues operating in a financial crisis as if nothing had changed – even if all risk measures (PDs of collateral issuers and counterparties, correlations, expected loss, CreditVaR, MarketVaR, etc.) have gone up dramatically, and all banks are cutting credit lines and are increasing margins in the interbank market. The inertia principle can be traced back to Bagehot (1873) who formulates it as follows (emphasis added):

> If it is known that the Bank of England is freely *advancing on what in ordinary times is reckoned a good security on what is then commonly pledged and easily convertible*, the alarm of the solvent merchants and bankers will be stayed. But if securities, really good and usually convertible, are refused by the Bank, the alarm will not abate, the other loans made will fail in obtaining their end, and the panic will become worse and worse.

Bagehot thus does not say: 'only provide advances on what is a good security also in the crisis situation', so he does not invite the central bank to join the flight to quality, but he says that advances can be provided on what was good collateral 'in ordinary times'. It may also be noted that Bagehot does not try to make a distinction between: (i) securities of which the *intrinsic* quality has not deteriorated relative to normal times, but of which only the qualities in terms of market properties (liquidity, sale price that can be achieved, availability of market prices, etc.) have worsened; and (ii) securities of which the intrinsic quality is likely to have deteriorated due to the real nature of the crisis (i.e. increased expected loss from holding the security, regardless of need to mark-to-market or sell the security). Not distinguishing these two is a very crucial issue. On one side, it appears wise as mostly, these two types of securities are not clearly distinguishable in a

[11] Also, the central bank should focus on its core business (monetary policy to achieve price stability), which is a sufficiently complicated job; second, it is unlikely to be a competitive player (with 'tax payer's money') in sophisticated risk taking. Third, it may encounter conflicts of interest when engaging in such business.

crisis situation, i.e. a liquidity crisis does typically arise if market players are generally suspicious and do not know yet where the actual losses will materialize. On the other side, not even trying to make the distinction means that the central bank's stabilization function does not only stem from its willingness to bridge the liquidity gap (which it should do as it is the only agent in the economy which can genuinely create liquidity), but to really take some expected losses.

The inertia ends when the central bank starts widening its collateral set, or when it relaxes risk control measures. Indeed, the Harman description of the 1825 events, where the Bank of England widened the set of assets it accepted ('*We lent ... by every possible means, and in modes that we never had adopted before*'), suggests that inertia sets a *minimum* constraint in terms of liberality of the central bank risk management in crisis situations, but that if the seriousness of the crisis passes some threshold, equal access FCM measures becomes necessary. Anyway, the striking feature of the inertia principle is that the increasing social returns to additional risk taking by a central bank in a crisis situation appear to always outweigh the increasing costs of the central bank taking more risks (although it is not a specialist in risk taking), such that there is for quite a range of events no point in tightening or loosening the risk mitigation measures of the central bank when moving within the spectrum from full financial system stability to various types and intensities of tensions. That this general inertia is optimal seems somehow surprising, since the two factors determining the trade-off are very unlikely to always support the same optimum. A number of arguments in favour of inertia per se may however be brought forward. First, only inertia ensures that banks can really plan well for the case of a crisis. The possibility that the central bank would make more constraining risk control measures or would reduce collateral eligibility in crisis situation would make planning by banks much more difficult. As the optimal changes of the credit risk mitigation measures would be likely to be dependent on various details of the ongoing crisis, it would also become almost impossible to anticipate these contingent central bank reactions in advance. Second, the central bank is unlikely to be able to re-assess the complex trade-off between optimal financial risk management (avoiding financial losses to the central bank and eventually to the taxpayer in view of its limited risk management competence) and optimal contribution to financial stability anyway at short notice, since both sides are difficult to quantify even in normal static conditions. Third, ex ante equivalence of

repeated access to the borrowing facility to a refinancing through open market operations (besides the penalty rate and possible stigmatization) also requires full trust of the bank into central bank inertia with regard to all access conditions to the borrowing facility. As it will be argued further below, convincing the banks that the borrowing facility is as good as open market operations (but for the penalty) is of help for financial stability as it may avoid to conduct large-scale open market operations which eventually can lead to unstable short-term interest rates, and an uneven reserve fulfillment path.

In implementing the inertia principle, some tricky issues arise. For instance how should rating requirements be handled, if important issuers start getting downgraded to below the rating threshold for collateral eligibility? Does the inertia principle refer to the actual set of collateral, or to the eligibility requirements? This of course depends on how rating agencies map short-term shocks and uncertainties into downgrades, and therefore probably no universal answer can be provided. If a central bank would conclude that substantial downgrades to below the rating requirement imposed by the central bank lead to a critical aggravation of the crisis due to the implied shrinking of the set of assets eligible for central bank credit operations, then it may lower the rating threshold to maintain these assets eligible. This may then be either interpreted as a form of application of the inertia principle, or as equal access ELA through a widening of the collateral set (as discussed further in Section 6.3). Another interesting question is whether the inertia principle also refers to central bank investment decisions (e.g. foreign reserve management, and domestic investment portfolios). In this area, it could be argued that the central bank is not a policy body, but an investor, which should feel obliged to prudently manage public money. On the other side one could feel that consistency issues would arise between policy and investment operations if they were treated too differently. Also, the central bank faces reputation risk if on one side it tries to persuade the market to relax, while on the other it does the same as other banks (e.g. cut credit lines) and thereby contributes to a liquidity crisis. Finally, inertia does not mean absence of knowledge about central bank risk taking in a crisis situation. On the contrary, the readiness to take more risks should go hand-in-hand with *sophisticated measurement of these risks*, such as to put decision makers in the position to judge on a continuous basis on the costs (in terms of risks) of inertia, or to possibly react if they feel that too much of additional risk arises.

6. Equal access FCM measures

In an *equal access* FCM measure, access conditions to liquidity are modified in the same way for all relevant counterparties of the central bank. In Section 2 of this chapter, four sub-cases of equal access FCM measures were introduced, each of which are addressed in more detail below.

6.1 Emergency liquidity injections through open market operations

Open market operations (OMOs) in relation to market turmoil were conducted recently by major central banks both after 11 September and in August–December 2007. An increase in interbank overnight rates beyond the target level was observed by the central bank and seemed to reveal a demand for excess reserves. Demand for excess reserves is normally rather limited in modern financial systems with reliable payment systems. However, under financial stress, banks tend to demand excess liquidity, and if the central bank does not respond, at least interbank short-term rates would go to above the central bank's target level.

What needs to be distinguished from this modern form of injection of excess reserves is the injection of additional central bank money through open market operations to address *changes in autonomous factors*, and in particular the increase of banknotes, or the decrease of foreign exchange reserves of the central bank in case of a fixed exchange rate regime. For instance the classical nineteenth-century liquidity crisis was often triggered by an increased demand for cash or for gold, which are classified as autonomous liquidity absorbing factors in the taxonomy of monetary policy implementation (see e.g. Bindseil 2004, chapter 2). Today, it is considered obvious that autonomous factor changes need to be reflected in corresponding changes of the supply of reserves through open market operations, and in any case, today's typical liquidity crisis no longer relates to increases in liquidity absorbing autonomous factors. By not caring about this difference, many academic or journalistic commentators (and even central banks), applying directly nineteenth-century insights to today's liquidity crisis, are potentially led to wrong conclusions. To get these differences right, commentators need first to understand the precise logic of monetary policy implementation.

So, again, today's question is: should a temporary demand for excess reserves, which is revealed in an increase of interbank overnight rates, be

satisfied through additional funds injected via open market operations? First, it needs to be understood why exactly banks would want to hold excess reserves in such situations. A liquidity crisis means essentially uncertainty about whether: (i) expected payments will come in from other banks; (ii) other banks will cut credit lines, i.e. will no longer be available to provide liquidity in case of need; (iii) customers may withdraw more deposits than expected; (iv) securities that would under normal circumstances be liquid can be liquidated without substantial discounts at short notice. All these effects may be summarized in model terms as meaning an increase in the variance of liquidity shocks, as it can be described in the standard liquidity management model originating from Poole (1968), as stated by Woodford (2003) and Bindseil and Nyborg (2008), in the case of a symmetric corridor set by standing facilities. Let S_j be the reserves bank j chooses to hold (through dealing in the interbank market) at the beginning of the day. The bank is subsequently subject to a shock in its holdings of ε_j, taking its end of day holdings to r_j. The shocks are independently distributed across banks with $E[\varepsilon_j|S_j]=0$, $Var[\varepsilon_j|S_j]=\sigma_j^2$. For each j, $\varepsilon_j|\sigma_j$ has a cumulative density function F, with a mean of zero, variance of 1, and $F(0)=0.5$. Let i, i_B and i_D denote the market rate, the rate of the borrowing facility, and the rate of the deposit facility, respectively, and let R be the aggregate reserves of banks with the central bank set at the beginning of the day. It can then be shown that (see e.g. Bindseil and Nyborg 2008):

$$i = i_D + F\left(\frac{-R}{\sum_j \sigma_j}\right)(i_B - i_D) \tag{11.1}$$

Thus by choosing R, for example through open market operations at the beginning of the day, the central bank can achieve any market interest rate within the corridor set by the two standing facilities. If $R = 0$, the market rate would be in the middle of the corridor (since $F(0) = ½$). Interestingly, this model would suggest that an increase in the variance of liquidity shocks would have no effect on interbank overnight rates. So why can one observe an increase in overnight rates whenever a money market crisis arises? The true corridor must in fact be different from what it seems to be. The following reasons for this can be found. First, an end-of-day excess of funds does not oblige a bank necessarily to go to the deposit facility, in particular not if the bank has still a considerable amount of reserve requirements to fulfill. If the excess of funds can be used for fulfilling reserve requirements, then there is no immediate cost to it. Second, if a bank runs out of collateral

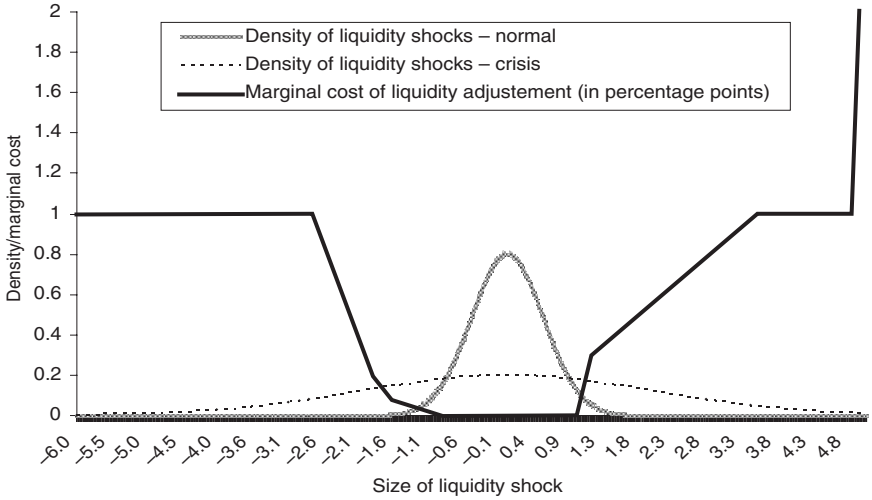

Figure 11.1 Liquidity shocks and associated marginal costs to a specific bank.

due to the liquidity absorbing shock, and cannot refinance in the interbank market, then it needs individual ELA, which is certainly a rather catastrophic and costly event. It may well be that some banks seeking funds in the unsecured interbank market already know that they are short of collateral, so the willingness to pay is very high. Third, even if a bank has enough collateral to refinance at the borrowing facility, the stigmatization problem arises. Will the central bank ask questions? Will other banks find out that the bank made this recourse and be will thus even more suspicious and will cut further their credit lines to the bank? The number of persons who will know that you took the recourse will always be considerable (both in the bank and in the central bank). The two large recourses of Barclay's to the Bank of England's borrowing facility in August 2007 in fact all became public – Barclay's made them public, probably anticipating that it would be worse if the market would find out itself.

Under normal market conditions, the last two points are far less relevant, which explains why a central bank like the ECB can normally consider that it offers a symmetric corridor system. The more intense a crisis, the less symmetric the effective corridor will be, and thus the higher the equilibrium rate in the overnight interbank market will be. Consider Figure 11.1, which illustrates the idea for a single bank. The bank is subject to daily liquidity shocks, i.e. unexpected in- or outflows of reserves which need to be addressed through money market operations or recourse to central

bank facilities. Every bank will have its own 'marginal cost of liquidity adjustment' curve, depending on parameters such as the credit lines other banks have granted to it, the credit lines it has granted to other banks, the size and equipment of its money market desk, its reserve requirements, and last but not least the availability of central bank eligible collateral. Small liquidity shocks can be buffered out at almost no cost through reserve requirements (with averaging), whereby this buffering effect is asymmetric because of the prohibition to not run a deficit at day end. Beyond using the buffering function associated with reserve requirements, the bank can use the interbank market, however, taking into account the bid–ask spread and increasing marginal costs of interbank trades due to limitations imposed by credit lines and market depth. At the end, the bank needs to make use of standing facilities, which in the example of Figure 11.1 are available at a cost of +/−100 basis points. Finally, banks can run out of central bank collateral when making use of the borrowing facility, and then the marginal costs of the liquidity shock suddenly grow very quickly or almost vertically. In a next step, marginal cost of liquidity adjustments curve needs to be matched against the density function of liquidity shocks. Figure 11.1 assumes under normal conditions a variance of liquidity shocks of EUR 0.5 billion, and of EUR 2 billion during a crisis. Assuming that the collateral basis of the counterparty considered is EUR 5 billion, then the probability of running out of collateral is around 10 E-24 under normal circumstances, but 45 basis points in a crisis, which makes a dramatic difference. It is important to note that for every bank, each of the three curves in Figure 11.1 will be different, and that it is not sufficient for a central bank to consider some 'aggregate' curves or representative banks.

Another reason for why interbank rates soar in case of a liquidity crisis is increased credit risk: as long as this does not lead to a total market breakdown, it would at least lead to higher unsecured interbank rates to reflect the increased risk premium.

The central bank will dislike the increase of short-term interbank rates first for *monetary policy reasons*. The target rate reflects the stance of monetary policy, and it is the task of monetary policy implementation to achieve it. Financial turmoil is, if anything, bad news on economic prospects, and therefore should, if anything, be translated into a loosening, and not a tightening of the monetary policy stance. Anyway, there is no need to adapt the stance of monetary policy within the day to macroeconomic news – it is almost always sufficient to wait until the next regular meeting of the policy decision-making body. If really needed, an ad hoc meeting of the

decision-making body could be called. Moreover, there could be *financial stability related reasons* why the central bank may conduct additional open market operations in a liquidity crisis are as follows. First, liquidity injections through OMOs could show the pro-activeness of the central bank: even if open market operations were not crucial in themselves, they may be a signal that the central bank is ready to act through more substantial means. Second, liquidity injections may incite those banks who have less problems (e.g. because they have comfortable liquidity or collateral buffers) to lend to those who have problems, because e.g. they may run out of central bank eligible collateral. For instance, one could imagine that good banks want to be sure to not have to take recourse to the borrowing facility at day end, but once this is sure, they would be ready to lend their funds in the overnight interbank market, and at least their propensity to do so will increase if they obtained additional cheap liquidity through an OMO injection. Third, banks may generally like to avoid recourse to the borrowing facility due to stigmatization. If open market operations allow avoiding that some banks have to go to the borrowing facility, and that this allows avoiding identification of those banks who have exceptional central bank refinancing needs, then this relative suppression of information could stabilize the system, since it avoids that runs on single banks take place. Fourth, if the collateral set in a special open market operation would be wider than in normal open market operations or in the borrowing facility, then this would be a way to inject additional liquidity in a controlled way against such additional collateral (as the Fed did in December 2007). By limiting the volume of the operation, the central bank could also limit the volume of such collateral it receives. Finally, banks may have strong preferences on the maturity structure of their liabilities vis-à-vis the central bank, i.e. they may prefer long-term OMOs relative to short-term OMOs, and even more relative to recourse to an overnight borrowing facility. This could reflect the fact they do not have full trust in central bank inertia (see Section 5), or that they are subject to some liquidity regulation which sets constraints on the maturity of liabilities (this was particularly relevant for German banks in the second half of 2007), or that they are just keen to get longer-term liquidity at a relatively low rate.

Injecting additional central bank reserves into the banking system through OMOs in a financial crisis could also have some *disadvantages*. First, by being 'activist', the central bank may send a signal that it knows unpleasant things about the true state of the market, which the market itself does not know yet. Second, in a reserve maintenance period with averaging,

injecting excess reserves on some days probably means, at least to some extent, that subsequent reserve deficits need to rise before the end of the reserve maintenance period. Third, lengthening the maturity of outstanding central bank operations could have the drawback of 'blocking' more liquidity with some banks over the longer term, which, in case of non-functioning interbank markets, would mean that it is no longer available to be channeled in the short term through short-term OMOs to the banks which have particular liquidity needs.

It may also be worth noting that in the Institute of International Finance (2007) study on liquidity risk management, not one single explicit recommendation to central banks seems to relate to injection of aggregate liquidity through OMOs as an FCM measure. Instead, all explicit recommendations (41–2) relate to collateral issues. This could suggest that indeed, special OMOs in a financial crisis context may have mainly the purpose of controlling short-term interest rates. The policies to be formulated ex ante in FCM-OMOs are also of interest from a risk management perspective, since these operations will increase risk taking linearly, and therefore need to be well justified. The eventual justification of these operations is likely to be related to liquidity risk considerations of banks. This should be analysed in more depth.

6.2 Narrowing the spread of the borrowing facility vis-à-vis target rate

In case that liquidity and/or infrastructure problems force banks to use extensively the borrowing facility, the central bank may want to alleviate associated costs by lowering the penalty rate applied to the borrowing facility. The ECB did so for instance for the first two weeks of the euro (in January 1999), and the Fed did so in August 2007. Again, this may appear at a first look more as a psychological measure, as it should not be decisive whether banks take overnight loans from the central bank at e.g. +100 or +25 basis points. The following advantages of narrowing the penalty spread associated with a borrowing facility could still be considered. First, it could be argued that any sign of central bank pro-activeness is useful in a crisis situation. Second, decreasing costs to banks, even if only marginally, cannot harm in a crisis situation. Third, this measure could avoid some of the major disadvantages of an excess reserve injection through OMOs, as in particular the destabilizing of the reserve fulfillment path. Also, lowering the borrowing facility rate may appear less alarmist and may less be misinterpreted as revealing that the central bank knows something bad that the

market does not know yet. Finally, it could be seen as an invitation of the banks to use the facility, and to reiterate that there should be no stigmatization associated with recourse. This is what the Fed may have tried to achieve in August–September 2007.

Possible disadvantages of a narrowing of the spread could be: First, as not decisive – why do it at all if it still may alarm banks, and may be misunderstood as a monetary policy move? Second, by reducing the penalty spread relative to the target rate, the central bank weakens incentives to reactivate the interbank money market. If a spread of e.g. 100 basis points is deemed optimal under smooth interbank market conditions in terms of providing disincentives against its use, then e.g. 50 basis points is clearly too little under conditions of a dysfunctional interbank market.

Maybe there is some low spread level where many banks would make use of the facility such as to overcome the stigmatization effect, so suddenly reducing the perceived full costs dramatically. For example, if the spread would be lowered to 5 basis points, use would probably become common, and stigmatization would vanish. Central banks probably want that: (i) stigmatization of recourse is avoided, which requires that there are quite some banks that take recourse for pragmatic reasons; (ii) the interbank market, however, has still room to breath, i.e. that banks continue lending to good banks in the interbank market. Ideally, the lowering of the spread could lead to a situation in which the gain of confidence effect would be such that interbank-market volumes at the end increase again due to this measure. Comparing the Eurosystem with the Fed suggests the following: as the Fed has anyway an asymmetric corridor (because it has no deposit facility) and as US banks have low reserve requirements, surplus banks have far stronger incentives to try to get rid of their excess funds in the interbank market, and this is not affected by a lowering of the spread between the target rate and the discount rate. Therefore, a lowering of the discount rate has more chances to have predominantly positive effects on the interbank market than it would have in the case of a symmetric narrowing in case of the Eurosystem.

6.3 Widening of collateral set

While ELA to individual banks typically involves a widening of the collateral set (since otherwise the borrowing facility could do the job), it can also be conceived that such a widening is done on an aggregate basis, i.e. maintaining equal access of all central bank counterparties. The central bank

increases available collateral by accepting a class of paper normally not accepted (e.g. Equity, commercial paper, bank bonds, credit claims, real estate, foreign currency assets etc.), or it could lower rating requirements for an existing asset class (require a BBB rating instead of an A rating). In case of two separate collateral sets for open market operations (narrow set) and standing facilities (wider set), a measure could also be to accept exceptionally the latter set for open market operations, such as was done by the Bank of Canada and the Fed in the second half of 2007. Even if applying equally to all banks, the central bank could choose a way to widen the collateral set which effectively targets a certain group of banks (or even a single bank), who are known to have the relevant collateral (but this could be inappropriate due to moral hazard).

The literature ever since the nineteenth century has taken different positions with regard to the widening of eligible collateral. As seen above, one would interpret Harman's 1832 description of Bank of England action in 1825 ('*We lent . . . by every possible means, and in modes that we never had adopted before*') as the position that being innovative and extending the collateral set is crucial and useful in a crisis (he is not saying: 'we lent volumes as never before', but is really referring to innovative collateralization or asset types purchased outright). Bagehot in contrast could be interpreted as arguing to lend only against assets which are normally eligible ('advancing *on what in ordinary times is reckoned a good security* on what is then commonly pledged and easily convertible').

Still today, different positions are found. For instance Rajan (Financial Times, 7 September 2007, 'Central banks face a liquidity trap') seems to have Bagehot in mind and argues that the central bank should not accept in particular illiquid securities. In contrast, Buiter and Sibert (blog: 12 August 2007, 'The central bank as market maker of last resort 1') take side with Harman's view of 1832 and argue that it is crucial that the central bank accepts illiquid security in a crisis situation, even if this is a particular challenge and implies additional risks (see also introductory quote to this paper from this blog). They argue provocatively that the central banks should do nothing less than becoming the 'market maker of last resort'. While some parts of the Buiter–Sibert propositions may appear doubtful and exaggerated (in particular the suggestion to do both repos and reverse repos at once, and to do outright operations), the substance, namely that the list of eligible collateral could be specifically widened in a crisis period to illiquid assets, makes potential sense under some circumstances, and is in principle in line with Harman's view of 1832.

In contrast to the two previously discussed equal access FCM measures (liquidity injection through OMOs and narrowing the corridor set by standing facilities), there is no doubt that an aggregate widening of the collateral set is highly effective to defuse liquidity problems of banks, in particular if the additional collateral volumes are substantial and if they are with those banks which could otherwise be most vulnerable to the liquidity squeeze. That a widening of the collateral set may be the most substantial equal access FCM measure is also suggested by the fact that all the explicit conclusions addressed to central banks in the Institute of International Finance (2007) report on liquidity risk management relate to collateral. In comparison with individual bank ELA, equal access widening of the eligible collateral set seem to have mainly two advantages: first, it may be more efficient since it allows helping a high number of banks and requires little if any demanding administrative and legal procedures; second, by avoiding making public failed counterparty names as it would be the case for individual ELA, it may avoid further deterioration of the market confidence and disruption of business relations.

However, some disadvantages of a general widening of the collateral set also need to be considered when taking such a decision. First, one may argue that ideally, the central bank *always* accepts a wide range of collateral, and this is, together with a stability-oriented monetary policy, the best general contribution that the central bank can make to financial stability. If the central bank did not accept a certain asset class in normal times, then this was probably for good reasons, such as for example (i) legal ambiguity; (ii) high handling or risk assessment costs; (iii) low liquidity compared to other types of financial assets; (iv) difficulties to value the assets. All these drawbacks also hold in a crisis scenario. Some of them, like handling costs, become less relevant in case of a crisis. In contrast, the weaknesses in terms of risk management (credit assessment and valuation difficulties) are likely to intensify in a crisis situation. Therefore the result of the cost–benefit analysis remains undetermined a priori. Moreover, the central bank will have little experience with these assets in a crisis situation, which increases operational and financial risks further. Individual (ex post) ELA has lower such risks due to the more limited scale of the use of additional types of collateral and due to the possibility to set up a tailor-made arrangement. Finally, as this measure is an effective one in helping banks that are in trouble and that have the respective collateral, while at the same time allowing these banks to avoid the humiliation and substantial consequences

for management and equity holders of having to request ELA from the central bank, it may be considered to be particularly subject to moral hazard issues.

Deciding on a widening of the set of eligible collateral (or a relaxation of risk control measures such as limits), will depend on: (i) the benefits of doing so in terms of financial stability; (ii) operational and legal considerations, including lead times; (iii) risk management considerations – i.e. how much additional risk would the central bank be taking, and how can it contain this risk through appropriate risk controls; (iv) moral hazard considerations. It is useful to have thought in depth through all of these aspects well in advance, as this increases the likelihood of taking the right decisions under the time pressures of a crisis.

6.4 Other equal access FCM measures

For instance, after 11 September 2001, the ECB provided US dollar liquidity to euro area banks through special swap operations (see press release of ECB dated 13 September 2001). In December 2007, the ECB provided USD liquidity against euro-denominated collateral (see press release of the ECB of 12 December 2007). Generally, it appears that access to foreign currency and the use of foreign collateral is an important topic under financial stress which may deserve special FCM measures in crisis times. International banks typically do business in different currencies, and are thus subject to liquidity shocks in these different currencies. At the same time, they may have collateral in different currencies and settled in different regions of the world. Ideally, collateral could be used as one large pool for liquidity shocks arising in whatever currency. In normal practice, central banks however tend to limit collateral eligibility to such assets located and denominated in their own currency/jurisdiction[12]. In a crisis situation, readiness of central banks to relax such constraints may increase, whereby a currency mismatch typically needs to be addressed through some extra haircut on collateral.

Cross-border use of collateral has for instance been analysed by Manning and Willison (2006). Also the Basel Committee on Payment and Settlement

[12] CPSS (2006, 3) suggests some reasons: 'Issues relating to jurisdictional conflict, regulation, taxation and exchange controls also arise in crossborder securities transactions. Although these issues may be very complex, they could be crucial in evaluating the costs and risks of accepting foreign collateral.'

Systems (CPSS) has published a report in 2006 on 'Cross border collateral arrangements'. Page 3 of this report explained that:

> Some central banks note that the emergency use of cross-border collateral has the potential to promote financial stability during a crisis ... [S]uch cross-border collateral arrangements could allow banks to access collateral assets in a market that may not have been directly affected by the emergency. Further, if foreign assets are only accepted in case of emergency and there is a low probability that an emergency-only facility will be triggered, banks may have a lower incentive to economise on precautionary collateral holdings and will, therefore, have a larger pool of collateral on which to draw in the event of an emergency arising.

Also the Institute of International Finance (2007) report on liquidity risk management invites central banks to take steps to permit cross-border collateralization.

As any widening of the set of eligible collateral, the acceptance of international collateral is potentially highly effective to address liquidity tensions. It may imply only limited additional risks to the central bank if operational, legal and financial risk management aspects have been studied carefully ex ante.

6.5 Conclusions: the role of equal access FCM measures

A modern central bank credit framework is typically characterized by including a borrowing facility which allows obtaining central bank funds at a limited penalty (up to 100 basis points) against a very wide range of eligible assets. The relevance of any equal access FCM measure has to be assessed against the existence of such a facility, implying in particular that the effectiveness of the injection of excess reserves through OMO as an FCM measure seems to relate to a number of not totally obvious effects, like (i) the stigmatization of recourse to the borrowing facility; (ii) uncertainty of banks about the central bank sticking to the inertia principle (i.e. restricting access to the borrowing facility); (iii) maturity preferences of banks vis-à-vis central bank refinancing, for instance because of liquidity regulations; (iv) an increased willingness of banks which hold excess reserves to lend in a liquidity crisis to other banks which have no collateral. All of these effects may be relevant and thus may justify FCM-OMOs. In any case, additional liquidity injections through OMOs in a crisis are sufficiently justified as monetary policy measures, namely to steer interbank market rates again closer to the target policy rate, which is a sufficient justification.

An equal access widening of the collateral set is quite of a different nature: it may be a highly effective measure, but it also could invoke substantial additional risk taking of the central bank and can cause moral hazard. A careful design of risk mitigation measures is obviously warranted. The decision between equal access and individual bank (ELA) widening of the collateral set has to be taken on the basis of a number of parameters characterizing the crisis, whereby the equal access variant could be the more attractive: (i) the higher the number of banks that need help and that would be helped through the widening of the collateral set; (ii) the more easily one finds additional collateral types that can be made eligible with a non-excessive increase in operational, credit and market risks, taking duly into account that the central bank is not necessarily an institution with high expertise and capacities in handling these risks; (iii) the more one is worried that individual bank ELA should be avoided because it further would intensify the sentiment of disaster and financial market crisis; (iv) the less one needs to be concerned with moral hazard (e.g. because a really extreme external event has triggered the crisis).

For the central bank, it seems essential to be well prepared to take the relevant decisions in a relatively short period of time. In particular for the equal access FCM measures, open market operations and the lowering of the borrowing facility rate, there are no important moral hazard issues at stake, and therefore there is no reason to not prepare detailed internal policies, as there is little harm if these are leaked. That one attempts to formulate such policies does not mean that actual execution would be mechanical. Almost by definition, every financial market crisis will be different, and therefore the optimal central bank actions cannot be anticipated in every detail. But accepting this uncertainty is far from accepting the conclusion that one should not try hard to think through the different cases and agree on a general policy. This may allow saving decisive time in a crisis situation, and thereby to avoid mistakes. Also on the possibilities to widen the collateral set, some internal policies should be elaborated in advance by the central bank. Many issues, such as legal and operational ones, take time to be analysed, and it will certainly help if the central bank knows in advance for all possibly relevant asset classes what the respective issues would be if considering making them eligible in some crisis situation. Also, risk control measures (credit quality requirements, valuation, limits and haircuts) should have been considered in advance. It is useful to analyse in advance the relevant crisis scenarios, and under which circumstances which types of assets could be useful. As this will necessarily be speculative and will not

imply any commitment, moral hazard arguments should again not be a sufficient argument for internal preparation. Preparation also gives more time to think about how to counteract moral hazard, and maybe what rescue operations to avoid under almost any circumstances due to moral hazard.

7. FCM measures addressed to individual banks (ELA)

ELA to individual banks may also be called 'ex post' FCM, since it is done once serious liquidity problems have materialized. Some (e.g. Goodhart 1999) suggest using the 'lender of last resort' (LOLR) expression only in this case of liquidity assistance, which sounds reasonable as it comes *last*, after possible ex ante FCM measures. Individual bank FCM is typically made public sooner or later, and then risks to further deteriorate the market sentiment, even if its intention is exactly the opposite, namely to reassure the system that the central bank helps. A decision to provide single-bank ELA will have to consider in particular the following parameters, assuming the simplest possible setting:

- $B =$ the social benefits of saving a bank from becoming illiquid. This will depend on the size of the bank, and on its type of business. It could also depend on moral hazard aspects: i.e. if the moral hazard drawbacks of a rescue are large, then the net social benefits will be correspondingly lower.
- $L =$ size of the liquidity gap of the bank.
- $C =$ value of the collateral that the bank can post to cover the ELA.
- $A =$ net asset value of the bank ($\approx$ net discounted profits).

In principle, all four of the variables can be considered to be random variables, whereby prudent supervision experts may contribute to reduce the subjective randomness with regard to A, central bank risk managers with regard to C, and financial stability experts with regards to B. Lawyers' support is obviously needed in all of this. Assuming for a moment that these variables would be deterministic, one could make for instance the following statements:

- If $C > L$, then there is no risk implied from providing ELA, and therefore no need to be sure about $B > 0$ and/or $A > 0$.[13]

[13] According to Hawtrey (1932), the central bank can avoid making a decision as to the solvency of a bank if it lends only on collateral (referred to in Freixas *et al.* 1999).

- If $A<0$, then the bank should probably be shut down in some orderly way. The 'orderly' would probably mean achieving to the extent possible B.
- If $C < L$, then $A > 0$ is important, if the central bank does not want to make losses.
- If $C < L$ and $L - C > B$ and $A = 0$, then do not do ELA.
- If $A < 0$, i.e. the bank is in principle insolvent, the state may still want to help if B is very large (i.e. $B > -A$).

Assessing the sizes (or, more precisely, the probability distributions) of the four variables will be crucial. This makes the joint analysis of prudent supervisors, financial stability experts and central bank risk managers even more relevant. In a stochastic environment, the central bank will risk making mistakes, such as in particular (i) to provide ELA although it should not have done so (maybe because it overestimated social benefits, or the value of collateral, or the value of the net assets of the central banks) and (ii) to not provide ELA although it should have (e.g. because it underestimates the devastation caused by the failure, etc. – see also Sveriges Riksbank 2003, 64). The likelihood of making mistakes will depend on the ability of the different experts to do their job to reduce the uncertainty associated with the different random variables, and to cooperate effectively on this.[14]

A number of considerations may be briefly recalled here:

- **Collateral set**: The collateral will consist in non-standard collateral, so probably less liquid, less easy to value, less easy to settle, etc. than the normal central bank collateral. Central bank risk managers will not only be required to assess this collateral and associated risk control measures ex ante, but also to monitor the value of collateral across time.
- **Moral hazard** would be addressed mainly by ensuring that equity holders and senior management suffer. This issue should be thought through ex ante. Setting a high lending rate may also be useful under some circumstances.
- **Good communication** is critical to ensure that the net psychological effect of the announcement of an individual ELA is positive.

As mentioned, the Sveriges Riksbank, the Bank of Canada (BoC), and the Hong Kong Monetary Authority (1999) have chosen an approach to specify ex ante their policy framework for individual bank ELA. With regard to the

[14] ELA to individual banks is only marginally a central bank liquidity management issue, since the liquidity impact to a single bank relating to ELA will simply be absorbed by reducing the volume of a regular open market operation. Central bank liquidity management issues will thus probably never be decisive to decide on whether or not to provide individual ELA.

precondition for ELA, the HKMA (Hong Kong Monetary Authority) specifies inter alia: (i) sufficient margin of solvency – at least 6 per cent capital adequacy ratio after making adjustments for any additional provisions; (ii) adequate collateralization of ELA support – adequate instruments are precisely defined as consisting of purchase of the institution's placements with other banks, acceptable for HKMA, reverse repos against investment grade securities, and credit collateralized by residential mortgage portfolios; (iii) institution has sought other reasonable sources of founding; (iv) no prima facie evidence that the management is not fit and proper; (v) institution must be prepared to take remedial action to deal with its liquidity problem. The Hong Kong Monetary Authority (1999, 79) also sets a limit for ELA: 'Notwithstanding the availability of suitable collateral, the HKMA will set a limit on the maximum amount of LOLR support provided to an individual institution via repos or the credit facility ... The limit would normally be set between 100 per cent to 200 per cent of the capital base ... depending on the margin of solvency the institution can maintain ... subject to a cap of HK$ 10 billion.'

Daniel *et al.* (2004, 8–9) also specifies explicitly a number of conditions for and specifications of individual ELA in Canada:

(i) maximum maturity of six months and one six-months renewal;
(ii) rate of interest – bank rate or higher (so far Bank of Canada applied its bank rate, i.e. no penalty rate);
(iii) collateral:

> In practice, it would be expected that the borrowing institution would use its holdings of marketable securities to obtain liquidity from the private sector before approaching BoC for ELA. If appropriate, the Bank could provide ELA loans on the pledge or hypothecation of assets that are not subject to as precise a valuation as are readily marketable securities. For example, the Bank may provide loans against the security of the Canadian-dollar non-mortgage loan portfolio of the institution, which can make up a significant portion of the institution's assets. Because the composition of a loan portfolio changes over time and the valuation of individual loans is subject to fluctuation, the bank would likely take as security a floating charge against the institution's loan portfolio (under the law, mortgages are considered to be a conveyance of 'real property', which the bank cannot take as collateral) ... the bank endeavours to minimize its exposure to loss in the event of default by the borrowing financial institution. Thus, it is important for the bank to have a valid first-priority security interest in any collateral pledged to support ELA. (Daniel 2004, 10)

(iv) eligibility of banks – only to banks which are judged to be solvent. ELA does not create new capital;
(v) ELA agreement creates a one-day, revolving facility in which the BoC has discretion to decline to make any further one-day loans (e.g. if it is judged that the institution is insolvent, or available collateral has a higher risk of being inadequate).

8. Conclusions

The summer 2007 liquidity crisis has revealed that views on adequate central bank FCM are heterogeneous. Central banks took rather different approaches, and split views were expressed by central bank officials on what was right or wrong. This could appear astonishing, taking into account that: (i) FCM and associated liquidity support is supposed to come second in the list of central bank core functions, directly after monetary policy; (ii) FCM is a topic on which economists have made some rather clear statements already more than 200 years ago; and (iii) there is an extensive theoretical microeconomic literature on the usefulness and functioning of (some) FCM measures.

How can this be explained? Most importantly, many commentators did not care that many of the FCM issues frequently quoted (e.g. moral hazard) are applicable to some of its variants, but not to others. Also the academic literature has contributed to this, by often not starting from a clear typology of FCM measures. The relevance of some comments in the summer of 2007 also suffered from a lack of understanding of the mechanics of the central bank balance sheet and how it determines the interaction between central bank credit operations and the 'liquidity' available to banks.

This chapter aimed at being pragmatic by, first of all, proposing a typology of FCM measures, such that the subject of analysis becomes clearer. The central bank risk manager perspective is relevant, since FCM is about providing unusual amounts and/or unusually secured central bank credit in circumstances of increased credit risk, valuation difficulties and liquidity risk. While the central bank is normally a pale and risk averse public investor, a financial crisis makes it mutate into an institution which wants to shoulder considerable risks. The central bank risk manager is crucial to ensure that such courage is complemented by prudence, that if help is provided, it is done in a way that is not more risky than necessary, and that

an estimate of financial risk taking associated with FCM measures is provided as a key element of the cost–benefit analysis that should underlie any FCM decision. Amongst the conclusions drawn in this paper, the following eight may be highlighted.

First, in terms of putting order into the different types of FCM measures, it is crucial to distinguish between equal access (ex ante) and individual (ex post) measures, as well as between whether they include or not the widening of the set of eligible collateral. Without making these distinctions, it is very difficult to make generally correct statements about FCM measures. Relating to this, it often leads to wrong conclusion to refer to statements by Bagehot (1873) when commenting on some FCM operation if this operation is rather different from what Bagehot had in mind.

Second, a first crucial contribution of the central bank to financial stability is the design of its *normal* operational framework, whereby a wide collateral set and not unnecessarily restrictive risk control measures are key. Also various other dimensions such as the distinction of collateral sets between open market operations and the borrowing facility, the size of the liquidity deficit of banks to be covered through reverse operations, the existence or not of reserve requirements and averaging, are all interesting but under-researched dimensions of the operational framework which will be relevant for the built-in stability of the interbank money market.

Third, the fundamental principle of *inertia* as the central bank's key contribution to financial stability in crisis situation was developed, stating essentially that the central bank should never restrict its collateral and risk management framework in a crisis situation, even if this implies that it ends up with much more financial risk than normally, and is the only agent in the market which does not react to changed circumstances. Central bank commitment to inertia (i.e. to at least not restrict credit) will be a basis for banks to plan how to sort out a crisis and to survive. Uncertainty about the possibility of restrictive moves by the central bank is exactly what the banking system will not easily digest as additional stress in a crisis. Inertia certainly does not mean that the central bank should be blind on its financial risk taking in a crisis situation – quite the contrary – financial risks as also implied by inertia should be measured and reported in a sophisticated way.

Fourth, the usefulness for financial stability of aggregate liquidity injections through open market operations remains relatively little explored, apart from a number of relatively soft arguments such as the psychological effect of showing central bank pro-activeness. At the same time, their

effectiveness in lowering short-term interbank interest rates is undisputed, so they may be understood in any case as monetary policy instrument in a crisis situation. Their possible effectiveness for financial stability also may emerge from a stigmatization of the central bank borrowing facility, which the summer 2007 episode confirms to be an issue. Reducing this stigmatization should therefore be an objective of central banks.

Fifth, an equal access widening of the set of collateral for central bank credit operations is likely to be the most effective ex ante FCM measure to contribute to financial stability. At the same time, it is the most challenging in terms of additional central bank risk taking and possible moral hazard issue. Its usefulness will obviously depend on how wide the set of eligible collateral is under normal circumstances.

Sixth, individual bank (or ex post) FCM measures (ELA) are also about widening eligible collateral, and are thereby also challenging in terms of central bank risk management. Compared to equal access widening of the collateral set, ELA appears to have some advantages, such as allowing the central bank to focus on one (or hopefully not more than a very few) institution(s), and to allow addressing more effectively moral hazard since the central bank and regulators could make sure that shareholders and senior management would be held responsible. On the other side, it is clear that ELA often does not defuse market tensions, and means a large administrative, legal and operational burden of the central bank.

Seventh, from a practitioner, and even more a risk manager perspective, the concept of constructive ambiguity appears to have some drawbacks. It may lead to weaker preparation, less accountability and transparency, more noise and uncertainty for the market, and maybe at the end even less punishment of those who would deserve to be punished since they took irresponsible risk at the expense of the financial system's stability. At least three advanced central banks have demonstrated that an effort for transparency and the establishment of rules can be made. This of course should not mean that there is a mechanistic commitment of the central bank to help, i.e. there is no 'ELA facility'.

Eighth, relating to the previous point, being prepared to implement FCM is crucial in reaching the right decisions when under time pressure, in particular from the risk management perspective. Key questions to be thought through as much as possible in 'times of peace' are for instance: Which asset types are candidates for widening the set of eligible collateral? Why is some type of collateral not accepted under normal circumstances, but should be under certain crisis events? What set-up issues will occur

(e.g. how to make the collateral eligible in the short run without legal and operational risks?). What haircuts will be appropriate? How exactly can one define eligibility criteria to have a clear frontier against hybrid asset types? Under what circumstances would limits be useful? Should additional collateral be only for the borrowing facility, or also for open market operations? How can one measure central bank risk taking in crisis situations, such as to ensure awareness of what price the central bank pays in terms of risk taking for maintaining inertia? A long list of similar questions can be noted down for other areas of ELA, such as the role of additional open market operations.

If central banks work on such practical FCM topics in a transparent way, one should expect that if once again, some day in the future, a liquidity crisis like the one in the summer of 2007 begins, there could be less misunderstandings and less debate about the right way for central banks to act.

Part III

Organizational issues and operational risk

12 Organizational issues in the risk management function of central banks

Evangelos Tabakis

1. Introduction

Risk management as a separate function in a central bank, with resources specifically dedicated to it, is a rather new development in the world of central banking. This may be considered surprising since central banks are, effectively, risk managers for the financial system as a whole. In their core functions of designing and implementing monetary policy and safeguarding financial stability, they manage the risk of inflation[1] and the systemic risks inherent in financial crises. Strangely however, until recently, they had not paid as much attention to the management of their own balance sheet risks that emanate from their market operations. A change is obvious in the last fifteen to twenty years. In part as a consequence of the general acceptance of the principle of central bank independence, central banks have been rethinking their governance structure. After all, financial independence of the central banks is an important element in supporting institutional independence from fiscal authorities and understanding, managing and accurately reporting on financial risks is necessary to control financial results. At the same time central banks as investors are facing increased challenges. Some have accumulated considerable sizes of foreign reserves and need to invest them in a diversified manner. This in turn makes them exposed to more complicated markets and more sophisticated instruments. Finally, risk management expertise is increasingly in demand in order to understand the complexity of risk transfer mechanisms in financial markets and detect potential risks of systemic nature.

In view of these developments, issues relating to the organization of the risk management function in the central bank are actively discussed in the

[1] For an interesting analysis of the parallels between the risk management function in a financial institution and the management of inflation risks in particular by the central bank see Kilian and Manganelli (2003).

central bank community. A number of questions such as the position of the risk management function in the organization of the central bank, the amount of resources dedicated to it, the types and frequency of reporting, the form of cooperation with other business areas and in particular those that take risks, the synergies between financial and operational risk management, the challenges of recruiting and training staff have yet to find their optimal answers. This chapter does not provide them either. However, it does attempt to provide a way of thinking about these questions which is consistent with financial theory, regulatory directives and best practices in financial institutions while, at the same time, considering the idiosyncrasies of the central bank.

2. Relevance of the risk management function in a central bank

What is the added value of risk management for a central bank? Addressing this question would provide guidance as to how to best organize the risk management function and how to distribute available resources. The academic and policy-related literature indicates two ways to approach this question.

First, one could look at the central bank as a financial institution. After all, central banks are active in financial markets, albeit not necessarily in the same way as private financial institutions, have counterparties to which they lend or from which they borrow money, engage in securities and commodities (e.g. gold) transactions and, therefore, face financial risks. Since its establishment in 1974 and perhaps more importantly since the introduction of the first version of the Basel Capital Accord in 1988, the Basel Committee on Banking Supervision (BCBS) has been the driving force for the advancement in measuring and managing financial risks. The goal of the Committee has been the standardization of capital adequacy frameworks for financial institutions throughout the international banking system with the aim to establish a level playing field. As capital and other financial buffers of financial institutions should be proportional to the financial risks that these institutions face, the guidance provided by the Basel Committee in the New Basel Accord in 2004–6 has set the standards that financial institutions need to follow in the measurement and management of market, credit and operational risks. Implementing such standards has become increasingly complicated and has led financial institutions to increase substantially their investment in risk management technology and know-how.

However, the fact that central banks have the privilege to issue legal tender as well as the observation that in some cases central banks have been operating successfully on negative capital, may cast doubts as to whether the capital adequacy argument, and the resulting importance of risk management, is equally relevant for the central bank. These considerations are examined in Bindseil *et al.* (2004a) where it is argued that, while annual financial results may be less important for a central bank, securing adequate (i.e. at least positive) capital buffers in the long run remains an important goal, linked to the maintenance of the financial independence from the government and of the credibility of the central bank. Therefore, ultimately, the *risk management function of the central bank strengthens its independence and credibility.*

Second, the central bank can be seen as a firm. The corporate finance literature has looked into the role of risk management in the firm in general. Smith and Stulz (1985) have argued that managing financial risks of the firm adds value to the firm only if the stockholder cannot manage these risks at the same cost in the financial markets. Stulz (2003) reformulates this result into his 'risk management irrelevance proposition' according to which 'hedging a risk does not increase firm value when the cost of bearing the risk is the same whether the risk is borne within the firm or outside the firm by the capital markets'. This principle is applicable only under the assumption of efficient and therefore frictionless markets. Some central banks are public firms, while for others it could be assumed that they are, ultimately, owned by the taxpayers. In both cases, it is doubtful whether every stock owner or taxpayer could hedge the financial risks to which the central bank is exposed in the financial markets at the same cost. This seems to be even more difficult for risks entailed in very specific operations initiated by central banks such as policy operations. In a very similar way, Crouhy *et al.* (2001) argue that managing business-specific risks (e.g. the risk of fuel prices for an airline) does increase the value of the firm. Interestingly enough, when carrying this argument over to the case of a central bank, a basis is provided to argue that the scope of risk management in central banks needs to go beyond the central bank's investment operations, and *needs in particular to focus on central bank specific, policy-related operations.*

3. Risk management best practices for financial institutions

The financial industry has worked extensively on establishing best practices on organizational issues including, in particular the role of the risk management

function within the financial institution. There is less clarity on the extent to which these guidelines should apply to central banks because of the specificities in the mandate, risk appetite and risk-taking incentives of these institutions (see also Chapter 1 of this book). However, according to the conclusions of the last section, it is certainly useful for central bankers to take into account the general principles and best practices available for financial institutions even if considerations of the specific business of the central bank may require some adaptations. There is no lack of guidance provided for the organization of the risk management function in financial institutions both from regulatory and supervisory agencies and as a result of market initiatives.

In the first category, the BCBS has become the main source of guidance for financial institutions in developing their risk management framework. The publication of the New Basel Capital Accord (Basel II) has set up a detailed framework for the computation of capital charges for market, credit and operational risk. While the purpose of Basel II is not to provide best practices for risk management, it implicitly does so by requiring that banks develop the means to measure their risks and translate them into capital requirements. When following some of the most advanced approaches suggested (e.g. the Internal Ratings-Based (IRB) approach for the measurement of credit risk and the Advanced Measurement Approach (AMA) for operational risk) banks would need to invest considerably in developing their risk management and measuring capabilities. Furthermore the disclosure requirements outlined under Pillar III (market discipline) include specific requests to banks for transparency in their risk management approaches and methodologies (see BCBS 2006b for details).

The relation between the supervisory process and risk management requirements is emphasized also in BCBS (2006c). In the paper, the Committee underlines that

> supervisors must be satisfied that banks and banking groups have in place a comprehensive risk management process (including Board and senior management oversight) to identify, evaluate, monitor and control or mitigate all material risks and to assess their overall capital adequacy in relation to their risk profile. These processes should be commensurate with the size and complexity of the institution.

In addition BCBS has published a number of papers that address particular risk management topics. BCBS (2004) deals specifically with the management and supervision of interest rate risk. It contains sixteen important principles covering all aspects of interest rate risk management ranging

from the technical issues of monitoring and measuring interest rate risk to governance topics (board responsibility and oversight) and internal controls and disclosure requirements. To a great extent this paper complements BCBS (2000b) that focused on principles for the management of credit risk emphasizing that 'exposure to credit risk continues to be the leading source of problems in banks worldwide'.

The BCBS has also looked directly into the corporate governance structure for banking organizations (BCBS 2006a). The paper notes that 'given the important financial intermediation role of banks in an economy, their high degree of sensitivity to potential difficulties arising from ineffective corporate governance and the need to safeguard depositors' funds, corporate governance for banking organizations is of great importance to the international financial system and merits targeted supervisory guidance'. The BCBS already published guidance in 1999 to assist banking supervisors in promoting the adoption of sound corporate governance practices by banking organizations in their countries. This guidance drew from principles of corporate governance that were published earlier that year by the Organisation for Economic Co-operation and Development (see also OECD 2004) for a revised version) with the purpose of assisting governments in their efforts to evaluate and improve their frameworks for corporate governance and to provide guidance for financial market regulators and participants in financial markets.

Finally, the BCBS has already in 1998 provided a framework for internal control systems, (BCBS 1998a) touching also on the important issue of segregation of duties. The principles presented in this paper provide a useful framework for the effective supervision of internal control systems. More generally, the Committee wished to emphasize that sound internal controls are essential to the prudent operation of banks and to promoting stability in the financial system as a whole.

A number of market initiatives for the establishment of sound practices in risk management are also worth mentioning. The 2005 report of the Counterparty Risk Management Policy Group II – building on the 1999 work of Counterparty Risk Management Policy Group I – is directed at initiatives that will further reduce the risks of systemic financial shocks and limit their damage when, rarely but inevitably, such shocks occur. The context of the report is today's highly complex and tightly interconnected global financial system. The report's recommendations and guiding principles focus particular attention on risk management, risk monitoring and enhanced transparency.

Furthermore, the work of COSO (Committee of Sponsoring Organizations of the Treadway Commission) on Enterprise Risk Management (COSO 2004) should be mentioned even if it relates more to the management of operational risks in an organization (see also Chapter 13 of this book). It provides a benchmark for organizations to consider evaluating and improving their enterprise risk management processes. A companion document, 'Applications Techniques', is annexed to the Framework and provides examples of leading practices in enterprise risk management.

Finally, the management of foreign reserves, one of the functions of a central bank that bears clear similarities with asset management in private banks, has attracted special attention by central banks and other institutions (see e.g. IMF 2004; 2005 and the Bank of England Handbook in Central Banking no. 19 (Nugée 2000)). Such publications necessarily include also an analysis of the role of risk management in foreign reserves and are therefore also useful references for the topic of this chapter. There does not seem to be however so far a treatment of the risk management function in the central bank as a whole from which one could deduce general principles for the organization of such function.

4. Six principles in the organization of risk management in central banks

The existence of an abundance of guidance for financial institutions in setting up and maintaining their risk management function has not necessarily made the work of risk managers in central banks easier. Which of these guidelines are applicable or even relevant in the central bank environment remains an issue of debate. This section will concentrate on the six points that, in the experience of the author have been most extensively discussed within the central bank community.

4.1 Independence of the risk management function

The BCBS's 'Framework for internal control systems' (BCBS 1998a) includes the following fundamental principle: 'An effective internal control system requires that there is appropriate segregation of duties and that personnel are not assigned conflicting responsibilities. Areas of potential conflicts of interest should be identified, minimized, and subject to careful, independent monitoring' (Principle 6). The importance of the principle of

segregation has become evident by the case of the Barings collapse in 1995 where unclear or non-existing separations between front, middle and back office allowed one person to take unusually high levels of financial risks. Today, segregation of tasks between risk management in financial institutions and the risk takers of these institutions (responsible for either the loan or the trade book of the bank) is sharp and reaches the top management of the institution where, normally, the Chief Risk Officer (CRO) of the bank has equal footing with the Head of Treasury. This principle is respected even if it results in some duplication of work and hence efficiency losses. So, for example, trading desks and risk managers are routinely required to develop and maintain different models to price complex instruments to allow for a full control of the risk measurement process by the risk managers.

It can be argued that the limited incentives of the central bank investor to take risks imply that there is no significant conflict in responsibilities between the trading desk of the central banks and the risk management function. Hence a clear separation at a level similar to that found in private institutions (that reward management according to financial results)[2] conferring complete independence of risk management from any administrative link to the senior management in risk-taking business areas of the bank, is not necessary.

However, the recent trend of diversification of investments in central banks in particular in the case of accumulation of significant foreign reserves may indicate that this traditional central bank environment of low risk appetite is changing. As the investment universe increases and the type and level of financial risks reaches other orders of magnitude, the need to have a strong risk management function operating independently of the centres of investment decisions in the bank will increase. Furthermore, reputation risks and the need to 'lead by example' are also important central bank considerations: central banks have the obligation to fulfill the same standards that they expect from private financial institutions, either in their role as banking supervisors (where applicable) or simply as institutions with a role in fostering financial stability. In addition, the reputation risks associated with what could be perceived as a weak risk management framework could be considerable for the central bank even if the corresponding true financial risks are low.

[2] For a thorough analysis of motivation and organizing performance in the modern firm see Roberts (2004).

A more complex and even less-elaborated issue is the role of risk management in policy operations. Here it can be argued that the operational area is not actively pursuing exposure to high risks (as this would have no immediate reward) but rather attempts to fulfill the operational objectives at a lower cost (for the central bank or for its counterparties) at the expense of an adequate risk control framework (see also the model in Chapter 7). In this situation as well, conflicting responsibilities arise, and segregation at an adequate level that guarantees independence in reporting to top management is needed.

How far up the central bank hierarchy should the separation of risk management from the *risk-taking* business areas reach? A general answer that may be flexible enough to fit various structures could be: the separation should be clear enough to allow the independent reporting to decision makers while allowing for opportunities to discuss issues and clarify views before such divergent views are put on the decision-making table. An optimal trade-off between the ability to report independently and the possibility to work together with other business areas must be struck.[3] In practice, the choice is often as much a result of tradition and risk culture as it is one of optimization of functionality.

4.2 Separation of the policy area from the investment area of the central bank – the role of risk management (Chinese walls principle)

Central banks are an initial source of insider information on (i) the future evolution of short-term interest rates, and (ii) on other types of central bank policy actions (e.g. foreign exchange interventions) that can affect financial asset prices. Furthermore, central banks may acquire non-public information of relevance for financial asset prices from other sources, relating for instance to their policy role in the area of financial stability, or acquired through international central bank cooperation.

Chinese walls are information barriers implemented within firms to separate and isolate persons who make investment decisions from persons who are privy to undisclosed material information which may influence those decisions. Some central banks have created Chinese walls or other similar mechanisms to avoid that policy insider information is used in an inappropriate way for non-policy functions of the bank, such as for

[3] In its first ten years of experience, the ECB has tried out various structures providing different degrees of independence for the risk management function. Currently, the Risk Management Division has an independent reporting line to the same Executive Board member to whom the Directorate General Market Operations also reports.

investment decisions. However, it appears that different institutions have very different understandings of the scope of this mechanism and the real or reputation risks it is set to mitigate.

A further twist in this context is the role of the risk management function in this schism. Although not taking investment decisions, the risk managers of the central bank could provide indirectly an input in these decisions by, for example, having a role in the strategic asset allocation of the bank's investment portfolios. At the same time, their role in the set-up and monitoring of the risk control framework in policy operations would require that they receive any policy-related information that would assist them in this role.

While this appears to be a delicate issue, there are simple and pragmatic rules that can be applied to overcome the above dilemma. Such rules could for example specify that the input provided by risk managers in the investment process is either provided ex post (e.g. through performance measurement and attribution) or, when provided ex ante, it is based on a transparent methodology and is free of taking private views on future market developments. In practice this means that proposals on asset location formulated (at the strategic level) by risk managers must be the result of a structured procedure involving well-documented models and using only publicly available information. Such procedures should generate an auditable trail that would serve to prove at any given time that the input of risk management in the investment process was not influenced by any insider information.

4.3 Transparency and accountability

Internal and external transparency are prerequisites for the accountability of top-level managers in the management of risks in a central bank. Internal transparency, which is a key issue in the first Counterparty Risk Management Policy Group report, may be achieved by:

- the detailing of the execution of functions in manuals of procedures that are regularly updated; and
- the regular reporting on risks and risk exposures to top-level management.

The core of risk management in any institution is the risk control function. It often comprises several regular tasks to be fulfilled on a high-frequency basis from a number of staff members. This requires detailed documentation of models, tasks and procedures. Although not the most exciting task, preparing and maintaining such documentation is an important, resource-intensive

duty of the risk management function. It allows the completion of the tasks by several staff members and supports knowledge transfer. It provides proof of the correct application of procedures regardless of the person executing the task, minimizing subjectivity and emphasizing rule-based decisions. It guarantees and documents for any audit process that the risk management processes have not been jeopardized by influences from other business areas.

Important decisions on the level of risks taken by central banks must be taken at the top level of the hierarchy. For this, however, top management in any financial institution depend on comprehensive and frequent reporting on all risks and exposures that the bank carries at any moment. An additional difficulty that arises in the central bank environment is that reporting on risks is 'competing' for attention with reporting on core issues in the agendas of decision makers such as information necessary to take monetary policy decisions. That is why risk reporting should be thorough but focused. While all relevant information should be available upon request, regular reports should include the core information needed to have an accurate picture of risks. They should emphasize changes from previous reports and detect trends for the future. They should avoid overburdening the readers with unnecessary numbers and charts and instead enable them to draw clear conclusions for future action. The best reports are in the end those that result in frequent feedback from their readers.

In most central banks, a number of committees have been created that allow that more detailed reporting and discussion of the risks in the central bank is considered by all relevant stakeholders before the core information is forwarded to the top management. Such are, for example, the Asset and Liabilities Committee, that examines how assets and liabilities of the bank develop and impact on its financial situation, the Investment Committee that formulates all major investment decisions and the Risk Committee that prepares the risk management framework for the bank.[4]

External transparency reinforces sound central bank governance. Therefore, ideally, central bank financial reports and other publications should be as transparent as possible regarding the bank's aggregate and lower-level risk exposures.

First, and maybe most importantly, informing the public and other stakeholders about the risks that the central bank incurs when fulfilling its

[4] Some central banks, like the ECB, may have specialized committees by type of risk, for example a Credit Risk Committee or an Operational Risk Committee.

policy tasks, such as holding foreign exchange reserves, seems essential to prevent reputation damages in case large losses actually materialize. If a central bank is not transparent about the risks it takes, critics can argue ex post in case of large losses that the risk taking was irresponsible and that the losses document the incompetence of the central bank. If in contrast the risks had been explained and quantified transparently ex ante, and if it can be documented that it was market movements (or credit events) that explain the losses, then such ex post criticism is much less convincing. Also the publication of risk figures ex ante obliges a central bank to ponder carefully about a thorough justification for taking these risks, and not only once a considerable loss has occurred.

Second, it could be argued that transparency is a value per se for any public institution. In particular central banks are not subject to market pressures and have been entrusted with independence and a particularly valuable franchise (to issue banknotes). It seems therefore natural to ask from them to adhere to the highest standards of transparency, unless there are convincing concrete reasons for exceptions.

Finally, the guidelines provided in the Third Pillar of Basel II, part 4, section II, (see BCBS 2006b), suggest that all internationally active banks should, as a part of their reporting, disclose VaR figures. Even though the central banks are not legally obliged to follow the guidelines set up in Basel II, it would be odd not to follow what is considered best practice for all major international banks. In cases where information is of obvious interest to the public, the central bank should only allow itself to diverge from best practice when strong arguments speak against this.

However, given the power to conduct monetary policy and to support financial stability, central banks are closely watched by financial market participants. Actions undertaken by central banks may easily be interpreted as a signal, even if it was not intended to be one. Thus, central banks must be careful to ensure that their signals are clear, and that actions not intended to convey a signal are not so interpreted. These considerations often prevent central banks from being more transparent about their market operations and the level of risks entailed in them. They may think that disclosing a change in the level of FX risks as result of a currency reallocation may mislead markets to believe that it was based on privileged information that the central bank had or even that it constituted a form of FX intervention. A difference in the duration of some portfolios may similarly send the wrong signal about future movements in interest rates.

Most of these risks can be mitigated by providing for a sufficient lag between the time of disclosure and the time of the actual occurrence of any changes. In most cases, risk-related information can be provided in the annual report of the central bank which is usually published well into the next year. This way, the information content that could potentially affect markets would have already dissipated while the information content for the general public retains its value.

4.4 Adequate resources

Investing in risk management does not come cheap. This is the lesson learned from the private financial institutions that have adopted Basel II guidelines for the management of risk while trying to cope with the ever-more-complex market landscape and innovations in risk transfer mechanisms. In proportion, central banks have followed suit. Not long ago, the need to acquire and maintain resources for risk management in a central bank was not an obvious fact and such tasks were allocated to staff in various business areas (foreign reserves management, monetary policy implementation, and organization).

General organization principles in the firm (see Brickley *et al.* 2007) indicate that the key investment to be made is that in human resources. Independence of the risk management function would be meaningless if this function could not attract and maintain high quality staff. It is therefore important that risk management staff are compensated at par with 'risk takers' in the institution and have equal career prospects. However, maintaining highly qualified staff in the central bank remains a challenge as such professionals will be generally better compensated in the private sector.

Given the quantitative nature of the work, risk management groups tend to attract more quantitatively trained staff, often with degrees in science and engineering. While the same type of skills is also useful in risk management of the central bank, it is also important to maintain in the group a number of economists and financial economists that would provide the link to the core business of the central bank. If, as argued in Section 2, risk management is most important in relation to the central bank specific operations with a policy goal, it becomes all the more important that staff has a good understanding of these goals.

The adequate function of risk management also depends on adequate systems. Chapter 4 discussed in detail the dilemma of 'build or buy' for IT systems. Given the rather constrained type of operations and range of

instruments in central bank operations, full-scale off-the-shelf systems are rarely adequate for the risk management function of a central bank. Fully in-house developed systems may be better suited to the needs of the institution but require a high degree of effort in maintaining them. Central banks are bureaucratic institutions par excellence and selecting, acquiring and implementing new systems takes time and effort. Whether building or buying, risk management staff needs adequate knowledge in IT in order to sustain a minimum amount of independence in running its vital systems.

4.5 Responsibilities of the risk management division

Given a set amount of resources, which are typically restricted in a public institution, priorities have to be set as to which types of functions and tasks in the central bank should fall into the scope of a dedicated risk management function. In most cases the risk management responsibilities in central banks (and in fact also in private financial institutions) have been built around the *middle-office functions*. These include the tasks of developing and implementing a framework of risk controls for all types of financial risks, monitoring the observation of such controls by risk takers and reporting to top management. In addition they include the measurement and attribution of performance which is in turn based on the ability to accurately valuate positions. Such functions are in the heart of the risk manager's responsibility and are the main reason why risk management should enjoy a sufficient degree of independence in its reporting lines.

One way to look at the risk control tasks is to describe them as 'ex post' tasks. They are centred around actions from the side of the risk manager (e.g. measurement of risks, measurement of return and performance, valuation, reporting on limit breaches) that take place *after* actual risks are taken. Of course they include also the design and frequent review of the risk control framework, for example the setting of the market, liquidity and credit risk limits as well as the responsibility to maintain the necessary systems to perform such tasks.

The next step in extending the responsibilities of risk management has been its involvement in market operations 'ex ante', i.e. *before* risks are actually taken. This is often achieved by entrusting risk managers with the preparation of the *strategic level of decisions in the investment process and/or the asset and liability management* of the institution's balance sheet. This of course does not change the fact that all related decisions on such important

issues should be taken by top management which is in a position to consider both the risk and the return dimensions. There are several reasons for this extension of the risk management role. First, considerable synergies exist with the risk control function. Strategic asset allocation and asset and liability management require a set of technical skills similar to those needed for the risk control tasks of performance measurement and attribution and asset valuation. Therefore, entrusting the strategic asset allocation to the same team that performs the risk control tasks would save valuable resources for the institution. Second, performance measurement and attribution are, by definition, based on a comparison of the results achieved by portfolio managers to a benchmark. Such a benchmark should then be selected by a team that enjoys adequate organizational independence from the portfolio managers. For the reasons described earlier, the risk management team, which performs the risk control tasks, is typically such a group and is therefore well placed for setting benchmarks. Finally, in institutions like central banks and other public investors, the role of active portfolio management is typically limited due to the lack of incentives for extensive risk taking and the overall risk averseness of top management. Therefore, the fundamental strategic decisions on asset and liability management mirror the risk–return preferences of the institution and essentially fully determine the return to be achieved. Risk managers are best placed to advise top management on these fundamental investment decisions.

Monetary policy operations, i.e. market operations that do not have an investment character and where securing return is not even a goal of the institution, are a unique feature of central banks. Already in Section 2 it was pointed out that the risk management of such operations is particularly important for the bank. Under normal market conditions, there is no incentive for the central bank to take additional risks on its balance sheet in order to accomplish its policy goals. Therefore central banks tend to minimize financial risks in such operations. Lending to the banking sector, for example, is mostly done through fully collateralized operations. Furthermore, assets accepted as collateral have to satisfy criteria and be submitted to risk control measures that are designed to minimize credit, market and liquidity risk. Designing the collateral policy of the central bank requires therefore the full range of risk management skills used also in the middle-office function of investment operations. Furthermore, useful transfer of knowledge takes place between risk management tasks in policy and investment operations as in both cases knowledge of the relevant markets

and instruments is paramount. Against this background, it is not surprising that many central banks have included the risk management of policy operations in the scope of the central bank's risk management function.

More recently, central banks have also looked more carefully and began dedicating resources to *operational risks*. While there are obvious synergies between the management of financial and operational risks in an institution, best exemplified in the global treatment of both kinds of these risks in Basel II, the organizational merging of the management of operational risk with that of financial risk has not been always the choice of central banks. Central banks face operational risks in the full range of their tasks and not only in the area of market operations where financial risks are prevalent. Furthermore, reputational issues are core concerns of operational risk management in central banks despite the fact that reputational risk is exempted from the Basel II definition of operational risk. Finally, operational risk management benefits the most from a decentralized approach where experts that 'know the business' are best placed to assess the severity of their risks and a central risk management unit is best placed to coordinate and report on such risks (see Chapter 13 for more details).

4.6 Risk management culture

The discussion over risk management in financial institutions (but also non-financial firms) in the last years has progressively moved from processes and tools to risk awareness and institution-wide culture. It has been widely accepted that managing risks is a responsibility affecting all areas of an institution and requires the cooperation of all staff. This is not in contradiction to the fact that dedicated teams of specialists may have well-defined functions within the area of risk management such as a middle-office function or a coordination of operational risk management.

Risk culture in central banks has traditionally been characterized by three aspects. First, reputational consequences of materialized risks are *ceteris paribus* considered more important and attract considerably more attention from top management than financial impact. This is a consequence of the importance that a central bank places on its credibility as a prerequisite in performing its core tasks of monetary policy and preserving financial stability. While financial losses, if considerable, will be a concern for top management, the concern will most probably focus on the reputational impact of such losses in the markets and the general public. Such focus on

reputation profoundly shapes the investment process in the central bank and determines the way it conducts policy operations.

Second, central banks are generally risk averse, at least when operating under normal market conditions. Risk averseness in firms is usually attributed in the literature to agency–principal conflicts and the related compensation schemes for executives and staff (see for example Jensen and Meckling 1976; Jensen and Murphy 1990; and, for a behavioural finance approach, Shefrin 2007). In central banks, risk averseness is exacerbated by special considerations. Until recently, a culture of zero-risk tolerance was part of the tradition of central bank operations. In operational risk management, this could result in suppressing reporting on operational risk incidents and thus underestimating their potential impact. In financial risk management this risk averseness has been incorporated in utility functions that place particular weight on no-loss constraints. A broader familiarity of central bankers with risk management concepts and techniques has changed the zero-risk-tolerance culture to one of institution-wide risk awareness and risk management responsibility. Developments in financial markets and the need for many central banks to manage considerable sizes of public funds have also brought again the risk–return considerations to the fore of the discussions.

Third, central banks have been always aware that while risk management considerations must be known and accounted for when decisions are made, the importance of financial stability may transcend the standard management of financial risks. In practice, this tension between policy actions and risk management is more likely to exist during a significant financial crisis. This is perhaps one of the more unique concerns of central banks, given its policy goals and responsibilities. The policy objective to promote systemic financial stability takes precedence over, e.g. the setting of risk limits, and may lead to accepting financial losses which risk control measures would normally seek to minimize. Nevertheless, the potential of a policy decision to loosen certain risk-diminishing practices in favour of promoting systemic stability does not obviate the necessity of being able to effectively monitor, measure and report the risks to which a central bank is exposed. In such situations it is particularly important that top management has accurate and up-to-date information on the risk position of the central bank, and the risk implications of all contemplated policy decisions. In fact, it is in such situations that risk managers can play an important advisory role in recommending the best ways in which to support the financial system while containing the resulting risks to the central bank.

5. Conclusions

In Chapter 1 of this book it was highlighted that while the central bank can be seen in many respects as just another financial investor, there are also characteristics of that central bank investor that distinguish it from counterparts in the private sector. In this chapter the debate on the similarities and differences between central banks and other financial institutions was used to discuss the impact of the idiosyncrasies of the central bank on governance principles in relation to the risk management function, but also to draw practical conclusions on how to organize such a function.

Despite the various specificities of central banks that stem out of their policy orientation and their privilege to issue legal tender, the core governance principles relating to the function of risk management are not substantially different from the private sector. On the contrary, Section 2 argued that it is particularly in those operations which are specific to central banks, i.e. those that have a policy goal, where a strong risk management framework is necessary. In fact the conclusion could be that the central bank should follow best practices in risk management for financial institutions *as the default rule* and deviate for them only if important and well-documented policy reasons exist for such a deviation.

Finally, it has been argued that what remains an important element of the risk management function of the central bank is the existence and further fostering of an adequate risk management culture in the institution. Such a culture that steers away from both extreme risk averseness, traditionally associated with central banks, and a lack of the necessary risk awareness is imperative for the appropriate functioning of the central bank both under normal circumstances and during a financial crisis.

13 Operational risk management in central banks

Jean-Charles Sevet

1. Introduction

As shown in the previous chapters of this book, *financial* risk management in central banking has come a long way. Managing *non-financial* risks is to a large extent the new frontier.

Central banks face a very wide array of non-financial risks. A few of them, in particular those related to reputation, have actually an importance which is difficult to oversee or overstate. Still, while very significant effort and progress has been made during the past twenty years to provide tangible solutions for the most pressing and visible concerns related to information security, physical security or business continuity, the broader topic of operational risk management (ORM) has stayed in a relative stage of infancy. In recent times, however, unprecedented forces have spurred a new wave of ORM initiatives. In an era where they encourage commercial banks to improve risk management in general and ORM in particular, central banks are more than ever committed to enhance their own competency and demonstrate that they fully practice what they preach. Faced with reinforced scrutiny on the use of public money, they strive to overcome their traditional bias towards risk aversion and further embrace values of effectiveness and efficiency through formal ORM frameworks and explicit risk tolerance policies. Last but not least, in a complex and uncertain new business environment featuring integrated financial markets and infrastructures, digital convergence and development of web-centric applications, and more generally emerging threats of the era of globalization (e.g. international terrorism, criminality or pandemic issues), central banks are resolutely starting to take a fresh and comprehensive look at the key non-financial risks which may compromise their ultimate objectives.

Thanks to the work of the International Operational Risk Working Group (IORWG), in particular, the state of play of ORM in Central Banking

can be apprehended through a substantial and representative information base. The IORWG (www.iorwg.org) was initiated by the Central Bank of Spain in 2005 to promote exchange of ORM best practices in central banking. As from mid 2008, this forum includes thirty-two central banks, reserve banks and/or monetary supervisory authorities from thirty nations in all five continents. Data and insights gained during the two past conferences inter alia helped delineate key trends mentioned in this chapter (see 2006 conference of the IORWG in Madrid and of the 2007 conference in Philadelphia). Overall, existing frameworks in place in central banks refer to a variety of sources of knowledge and experience. The attached reference list compiles the most frequently used ones.

The topic of ORM has been covered in recent years by a few authoritative textbooks – see Cruz (2002) or Marshall (2001) – which provide a set of useful references regarding critical aspects of modern ORM. Yet, the various techniques discussed in these and similar books and articles (e.g. database modelling; stochastic modelling through severity models; extreme value theory, frequency models or operational value at risk; non-linear models and Bayesian techniques; hedging techniques etc.) are not further discussed in the present chapter. More than the impossibility to provide a meaningful summary of quantitative techniques, the main reason for leaving textbook ORM aside here is that central banks have only very marginally followed this path and opted to address their specific needs essentially through the use of qualitative techniques (see rationale in Section 2).

Regarding the latter, the attached list of reference mentions frequently used 'risk management standards' which have been considered and/or adopted by central banks' various departments and were developed by professions as diverse as experts of insurance management (Association of Insurance and Risk Managers 2002), internal audit (Institute of Internal Auditors 2004), information system security (International Organization for Standardization 2002 and 2005; Information Security Forum 2000), physical security (US Department of Homeland Security 2003), project management (Project Management Institute 2004), procurement and outsourcing (Office of Government Commerce 1999), business continuity (British Standard Institutions 2006; Business Continuity Institute 2007), specific business lines (Financial Markets Association 2007), or public sector bodies (International Organization of Supreme Audit Institutions 2004; Standards Australia 2004). In many central banks, COSO (2004), an integrated framework developed in the world of audit and internal control by the Committee of Sponsoring Organizations of the Treadway Commission,

is frequently referred to, in an attempt to glue together various elements related to operational risks and related controls. Due to their significance for financial institutions, ORM standards defined by the Basel Committee on Banking Supervision (BCBS 1998a; 1998b; 2001a; 2002; 2003) in the wider context of Basel II have been considered by most central banks for aspects of immediate relevance (e.g. taxonomy of operational risks events, sound practice of governance). In the US, regulatory requirements of the Sarbanes–Oxley act also served as a catalyst to propagate certain techniques pertaining to risk and control self-assessment.

Over the last three to five years, finally, active benchmarking initiatives within the central banking community have provided the most relevant source of knowledge and experience on ORM matters. At a global level, on top of the aforementioned work of the IORWG, an internal study group organized by the Bank for International Settlement produced a report on this topic ('Risk Management in central banks', Bank for International Settlements 09/2007). In the context of the US Federal Reserve System, the Federal Reserve of Philadelphia assumes a coordination role for all reserve banks. And at the level of the Eurosystem of central banks, a dedicated working group is expected to complete the development of a common ORM framework by the end of 2008.

Reflecting on this experience, this chapter touches upon ten generic aspects of operational risk management in central banks and illustrates them by presenting the respective solutions currently implemented at the European Central Bank. The ECB launched an ORM programme in November 2005 with a view to:

- harmonize and integrate the various risk management frameworks which had been previously developed in a decentralized mode across the various business areas and risk categories of the ECB during its founding years;
- introduce best practice elements of ORM, in consideration of the specific requirements of a not-for-profit organization;
- lay the foundation for a harmonized ORM framework for Eurosystem central banks, which a dedicated working group is expected to finalise by the end of 2008.

The framework developed by the ECB in 2006 and early 2007 has greatly benefited from all previously mentioned sources of knowledge and experience, as well as from an evaluation of fifteen ORM software solutions. Roll-out started in September 2007.

Section 2 of this chapter reflects the wide consensus of the central banking community regarding the fundamental specificity of ORM for the

industry. The remaining sections, while discussing standard concepts and practices, also highlight some specific aspects of the new ECB framework and, doing so, make a plea for revisiting a few elements of conventional ORM in Central Banking.

2. Central bank specific ORM challenges

ORM is a discipline that continuously and systematically identifies, assesses and treats operational events that may impact the key objectives of an institution. Challenges to develop and implement ORM within central banks are both generic and highly specific.

For central bankers, as for any other organization, ORM poses a formidable methodological challenge: financial risk management disciplines consider a small number of fairly homogeneous categories of risk-generating events (e.g. default of credit counterparties, fluctuation of interest or currency rates) and can accordingly, at least theoretically, build statistical models to slice and dice relatively large populations of events and determine critical underlying risk drivers.

By contrast, events related to operational risks are by nature much more complex and heterogeneous. As the types of operational risk events are of a theoretically infinite number, organizations of all size always must cope with inextricable issues of paucity of historical data to validate ORM analyses. Everywhere, ORM practitioners must engage in technically complex and politically sensitive efforts of information pooling and sharing with external partners to complement their own databases. And everywhere, ad hoc tweaks in data sets and risk assessment models are required, in particular to cover very rare 'fat-tail' events where no historical information at all is available.

While central bankers naturally share these generic issues with their colleagues from the private sector, they must additionally, like other not-for-profit institutions, take into account two very specific aspects: in ORM matters, like for most other management disciplines, both their ultimate objective and their key values and incentives are of a fundamentally different nature than those of private sector companies.

2.1 Non-financial objectives

Basics sometime matter: as central banks' ultimate objectives are fundamentally different from private sector companies, so are the specific objectives

of ORM. Private companies' *raison d'être* is to create value to their shareholders. This orientation fully justifies that all operational events, even those linked to intangible aspects like reputation or service quality, should ultimately be captured in a financial 'value at risk'. In simple terms, private sector ORM can and must in essence be based on a quantitative approach. By contrast, central banks' critical goals and assets are of a non-financial nature and relate to the potential non-achievement of specific legal and/or statutory obligations – in the case of the ECB, for instance, the latter are defined in the Maastricht treaty and the ESCB Statute. Because the most severe impacts of their operational risk events cannot be quantified in monetary terms, central banks have a natural and fully legitimate inclination to emphasize qualitative approaches. Admittedly, quantitative risk modelling still can be applied to their few transaction-intensive processes to measure operational losses. Certainly, sound scientific reasoning, starting by validating human judgement by facts and evidence, must always be guaranteed. Yet, at the end of the day, central banking risks are primarily of a qualitative nature and can make only marginal use of more sophisticated quantitative techniques.

2.2 Not-for-profit values and incentives

The second key difference in implementing ORM in central banks as opposed to within private sector companies relates to base values and incentives systems. New requirements like Basel II which encourage private sector banks to invest in ORM frameworks and systems are sometimes perceived as discretionary and exogenous regulatory pressures. Yet, more fundamentally, a very powerful economical rationale is at play: large and/or sophisticated banks do not commit to the costly application for an *Advanced Management Approach* (AMA) accreditation based on sheer considerations of prestige. By managing their operational risks in a more effective manner, commercial bank managers can free up increasingly scarce economic capital, improve the risk-adjusted return on capital (RAROC) of a given business line, which will more or less directly and immediately translate into a monetary reward.

Yet the value and incentive systems of central bankers are of a totally different nature: no market discipline exists, which can aggressively counterbalance natural concerns of risk avoidance by criteria of cost performance. And public officer employment schemes paired with demographic constraints considerably limit opportunities to reward successful managers

or staff and to punish poor performers. Addressing the hidden and yet decisive change management question ('*ORM? – What is in it for me?*'), no simple carrot-and-stick answer is available. More than anywhere else, patience and long-term commitment are of the essence. ORM benefits in central banks are more collective ('*Develop a shared view of our key risks*') than individual ('*Win over the budget on project x*'); more visionary ('*Preserve and enhance our reputation as a well respected institution employing highly trustful and qualified professionals*') than materialist ('*Secure a 25 per cent bonus*'); and also more protective (e.g. '*Rather proactively disclose incidents than be criticized in a negative audit report'*) than offensive ('*Reducing risks in service line x will free up resources for opportunities in service line y*').

3. Definition of operational risk

In the papers of the Basel Committee on Banking Supervision, operational risk is defined as 'the risk of loss resulting from inadequate or failed internal processes, people and systems or from external events'. However, the aforementioned specificities in central banking typically call for a wider definition and scoping of operational risk.

At the ECB, for instance, the latter is defined as 'the risk of negative business, reputational or financial impact for the bank which derives from specific risk events due to or facilitated by root causes pertaining to governance, people, processes, infrastructure, information systems, legal, communication and changes in the external environment'. In comparison to the Basel II definition, this formulation lays the emphasis on ultimate non-financial impacts, stresses the importance of explicitly managing risk-generating events, and extends the scope of functional causes of risks to include governance and legal matters.

Yet what is risk? The aforementioned risk management 'standards' frequently refer to intuitive definitions for objectives of simplicity. As typical examples, the Australia/New Zealand standard AS/NZS 4360:2005 defines risk as 'any threat of an action or event to our industry or activities that has the potential to threaten the achievement of . . . objectives' and the widely used COSO framework understands this notion as 'the possibility that an event will occur and adversely affect the achievement of objectives'. On closer examination, definitions of that kind represent a gross oversimplification of the notion of risk. As notably shown by a series of provocative articles by Samad-Khan (2005; 2006a; 2006b), they generate a few misconceptions which

are still at the core of traditional ORM frameworks. In order to reflect a few critical concepts and parameters used in statistical theory, the notion of risk should be more precisely defined – for instance as 'the *area of uncertainty* surrounding the *expected negative outcome* or impact of a *type of event*, between *normal business conditions* and a *worst-case scenario* assuming a certain *level of confidence*'. By design, risk is a function of the *frequency distribution* of a type of event as well as of the related *impact distribution*.

Of course, such cryptic jargon is inappropriate when communicating to pressured managers or staff. In essence, however, four simple and practical messages must be explained over time.

3.1 Risk as a distribution

The term 'risk' refers to notions of distribution – i.e. it does not mean *the* product of *the* likelihood and *the* impact of *one* single event. In audit and project departments, very useful instruments have been developed to assess the 'probability-weighted cost' of single events or decisions and to manage log books of potential specifically identified incidents. As the term 'risk' is frequently used in common language to qualify these tools (e.g. assessing the 'risks' of investing in system x, reviewing the list of 'risks' of project y), managers have very naturally learnt to think of this notion as of a unique combination of likelihood and monetary (or non-monetary) impact. Using day-to-day examples, it is essential to illustrate that operational risks applying to recurring processes actually reflect *types* of events, which may impact very differently the institution over time.

3.2 Normal business conditions vs. worst-case scenarios

As paucity of historical data always impairs (and generally prohibits) proper modelling of the frequency distribution of risk-generating events for most central banking activities, ORM must consider at least two fundamentally different cases:

- Negative business outcomes *in normal business conditions*, i.e. considering the regular conditions and environment that an institution faces when executing its daily tasks. The frequency of risk-generating events under normal business conditions can be observed, and their financial impact at least can be measured. For good reasons, insurers and credit institutions explicitly factor in and manage such 'average losses' as the 'cost of doing business' requiring a recurrent 'risk premium' for their customers.

- Negative outcomes *under a worst-case scenario*, i.e. stating very unlikely yet plausible assumptions on possible risk events, underlying root causes and ultimate risk impacts for the institution. Due to the absence of relevant internal historical data, some central banks have in the past tended to insufficiently reflect on such worst-case scenarios, leaving them potentially oblivious and vulnerable to some of their most severe risks. Yet, if the 'likelihood' of worst-case scenarios cannot be assessed, their plausibility can be fairly well documented based on experts and managers' judgements: very adverse external events which have historically happened in other countries and/or industries give at least partially relevant hints about potential catastrophic scenarios, and the severity of their impact can be ascertained in a fairly straightforward manner with due respect to the specific control environment of the institution.

3.3 Danger of the classical likelihood/impact matrix

As a consequence of the above, conventional representations of operational risks must be adapted. Nowadays, most central banks continue to use a 'likelihood–impact' matrix to visualize their operational risks. Yet, using such a matrix to report on heterogeneous operational risk events only produces 'apples and pears' comparisons: As an example, a 'worst-case' event like 11 September, assessed as 'very unlikely' and 'very severe' will by design appear in the yellow zone of the matrix and therefore be misleadingly presented as less of a concern than a current, 'very likely' and 'severe' event like a pending legal issue appearing in the red zone. As illustrated in Section 6, a revised version of the matrix is required to allow management to more realistically apprehend red-zone risks, be it under normal business conditions or under worst-case scenarios.

3.4 Inherent risk vs. worst-case scenario

Ultimately, the only things that matter in ORM are actual and recurrent incidents as well as plausible worst-case risk scenarios – and how both categories are being mitigated. By contrast, the frequently used notion of 'inherent risk' should arguably be abandoned and replaced by formal worst-case scenario analysis.

'Inherent' risks are typically defined in traditional ORM frameworks as 'raw' risks which would exist irrespective of (or before) any control.

Experience however demonstrates that using the notion of inherent risks presents three limitations:

- First, managers requested to reflect on inherent risk intuitively realize that assessing the likelihood and impact of a risk event in a totally hypothetical scenario constitutes a fairly improbable task.
- Soon, they also note that assessing fully theoretical risks does not really create valuable information for decision making and action. For existing processes, which represent the vast majority of cases, some controls (even minimal ones) are already in place and should not be ignored. And even for totally new project or initiatives, where control should be defined from scratch, reflecting on totally theoretical risks (which by definition would create the most extreme damage) does not help, and a more relevant approach is to determine a plausible worst-case scenario.
- At the end of the day, the key reason to abandon the concept of 'inherent risk' is that the idea of 'risks without any control' is to a large extent a fiction. The following riddle helps to realize this: *What is the risk of gold bars being stolen in a bank's safe left permanently open without any guard or security arrangement?* Actually, if such would ever be the case, our strangely absent-minded bank would face no risk but rather an absolute certainty: one day or another, and probably sooner than later, somebody would steal that gold.

4. ORM as overarching framework

When setting out to introduce their relatively new discipline, risk managers typically face the daunting challenge of explaining in simple terms why and how ORM, far from replacing or competing with approaches traditionally used for specific categories of risks and controls, actually creates unique value.

Indeed, in all central banks, a large number of policies, procedures and instruments establish a general framework for governance, compliance and internal control, and specifically organize the management of the confidentiality, integrity and availability of information, of the physical security of people and premises, and of the continuity of critical business processes.

Over time, central banks have increasingly come to recognize that this initial approach to various categories of operational risk events has been exceedingly piecemeal. In essence, ORM provides the overarching framework which has been historically missing in most institutions and finally

makes it possible to manage operational risks in a consistent and integrated manner.

What is an ORM 'framework'? Though no academic definition exists, it is widely used and understood by central banks as the verbal and visual representation of the interlinked components which are needed to identify, assess, mitigate and monitor their operational risks. One way of summarizing and charting ORM framework components is to use popular analytical grids like McKinsey's 7S model (Strategy, Structure, Systems, Skills, Staff, Style and Shared values).

In the ECB, the ORM framework has been defined around seven more specific contents. Three of them focus on the required umbrella methodology for risk and control issues. They are:

- a common language – the operational risk taxonomy (see Section 5);
- a generic risk management lifecycle (see Section 6);
- explicit strategic orientations stated in the operational risk tolerance of the Executive Board (see Section 7).

The remaining four other components of the ORM framework of the ECB consist of:

- one yearly top-down ORM exercise providing the big picture of risks at the level of the macro-processes of the bank (see Section 8);
- a five-year programme of bottom-up ORM exercises defining specific action plans at the level of each of its individual processes (see Section 9);
- a governance model fostering convergence and integration of all vertical and horizontal disciplines and activities related to operational risks and controls (see Section 10);
- new developments in the area of ORM reporting and key risk indicators (see Section 11).

5. Taxonomy of operational risk

From a technical perspective, central banks all complement their verbal definition of risk with a taxonomy, i.e. a systematic way of categorizing various risk items. An internal survey on ORM practice within Eurosystem central banks (see report to the Organizational Working Group, June 2007) confirmed previous findings of a IORWG study (see acts of the 2006 conference of the IORWG in Madrid): most institutions have opted to adapt to their own needs the taxonomy of risk events proposed in the Basel II papers. However a few central banks, including the ECB, found it useful to develop

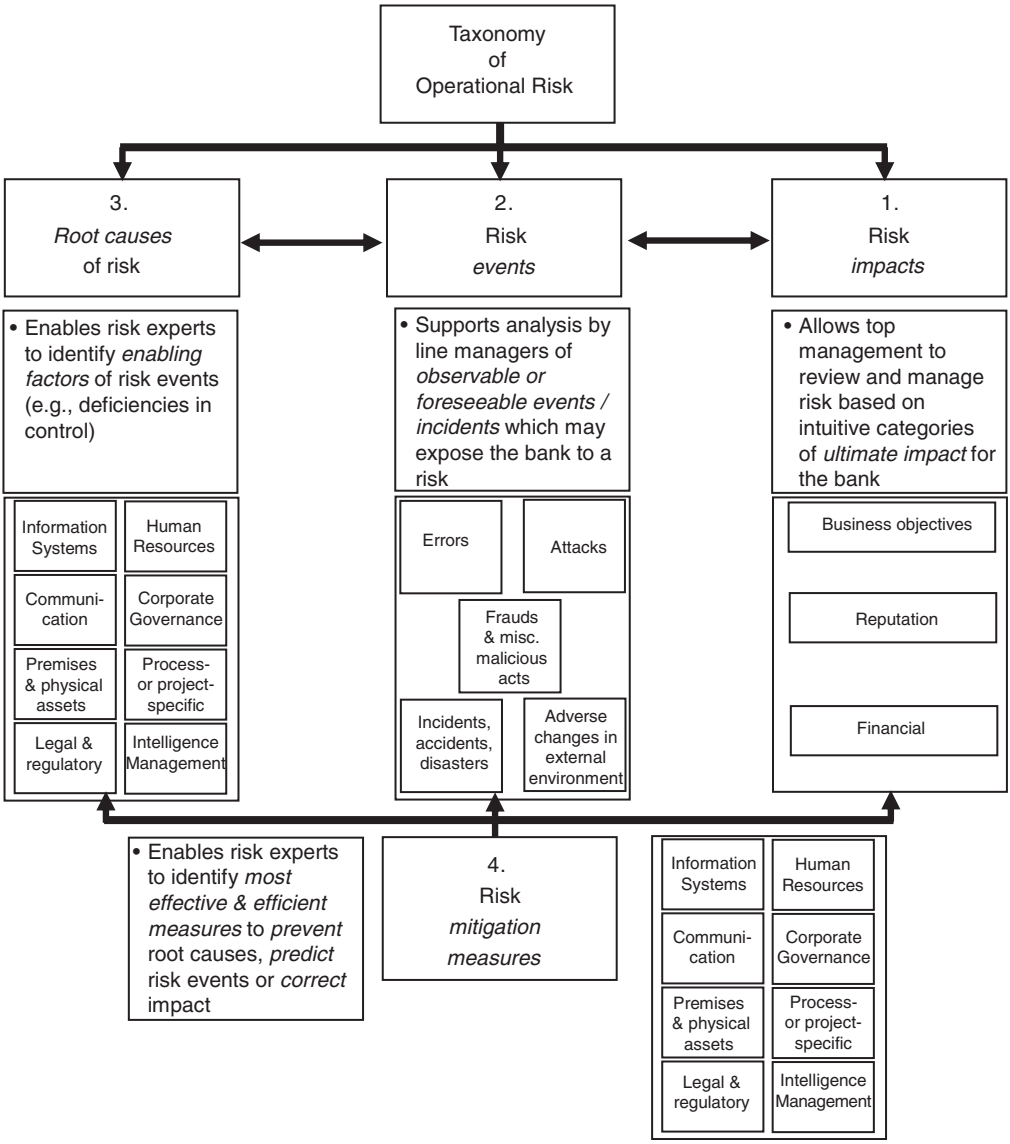

Figure 13.1 Taxonomy of operational risk.

a more comprehensive taxonomy to describe the full 'causality chain' of operational risks, including a categorization of their *root causes*, of observable *risk events*, of controls or other *risk treatment measures* and of ultimate *risk impact* (see Figure 13.1).

The three objectives of this taxonomy are to provide a clear and common language for all risk, control and security stakeholders of the ECB, to support

the quality of risk analyses via robust, mutually exclusive and commonly exhaustive categorizations, and to allow for consistency in risk reporting.

Mapping the full causality chain also helps overcome frequent misunderstandings about the term of 'risk': indeed, for reasons of simplicity in daily communication, the latter is typically (mis)used to express fundamentally different notions such as the given root cause of an event (e.g. in expressions such as 'legal risk', 'HR risk', 'information security risk', 'political risk' etc.), one type of undesirable event which may ultimately generate a negative impact (e.g. in expressions like 'risk of error', 'risk of fraud' etc.) or the nature of such an impact (e.g. in expressions such as 'business risk', 'reputation risk', 'financial risk', 'strategy risk' etc.). Experience at the ECB demonstrates that a comprehensive taxonomy of operational risk can remain simple and user friendly. In practice, it constitutes a modular toolbox used by all risk stakeholders on a flexible, need-to-know basis. Typically:

- risk impact categories are mostly relevant for management reports, as they highlight the type of ultimate damage for the bank;
- risk event categories are extremely useful to structure management or expert discussions regarding the frequency or plausibility of certain risk situations;
- categorizations of root causes and of risk treatment measures are used on a continuous basis by the relevant business and functional experts, in order to detect risk situations, monitor leading risk indicators or select the most effective risk treatments.

Within each of these four categories, a tree structure (simple level one list of items, further broken down into more detailed level two and three categories) allows risk stakeholders to select the level of granularity required for their respective needs.

6. The ORM lifecycle

A second useful element in an ORM framework is a generic representation of the risk management lifecycle. The aforementioned 'standards' on risk managements use largely identical, yet also partly specific and partly contradicting concepts, tools or approaches – and summarize them in heterogeneous representations of the activities and outputs of risk management. To facilitate overall coordination, consistency and transparency, the ECB has mapped the approach used by all existing risk disciplines to a standard lifecycle comprising the five following phases: (1) Risk identification; (2) Risk assessment; (3) Design and planning of risk treatment measures;

(4) Implementation of risk treatment measures and (5) Ongoing risk monitoring, review, testing and reporting.

At the ECB, the three initial phases of the lifecycle are implemented in an integrated manner in the context of ORM top-down and bottom-up exercises and complemented wherever required by specialized risk assessments (e.g. analysis of the business criticality of specific physical or intangible assets; specification of related new or additional security or business continuity requirements).

The last two phases of the lifecycle are conducted under the primary responsibility of business areas. The final phase, in particular, continues to require significant contribution from the various control specialists of the bank. For all of them, it includes tasks as diverse as continuously checking the status of the bank's key risks; verifying that the latter remain in line with the bank's risk tolerance; ensuring that required corrective action plans are implemented and are progressing according to agreed schedules; scanning the business environment to detect emerging new risks and regularly aggregating the information on risks and mitigation responses in coordination with the central ORM team.

In order to encourage business areas to fully disclose incidents or near-losses, candidly discuss emerging threats and define relevant measures, internal auditors do not participate in self-assessments workshops, nor are they associated in their actual implementation or in the preparation of risk reports. Still, internal audit is by principle entitled to full access to the various outputs of ORM processes – be they in the form of collected information on risk events, results of self-assessments, action plans and/or final reports. Such material gathered in a more standardized and homogeneous manner than in the past provides auditors with invaluable insights to plan, reprioritize, execute and monitor risk-based audit programmes, as required by international standards.

7. Operational risk tolerance policy

The third and most critical element of an umbrella ORM methodology is a formal, ex ante definition of the bank's operational risk tolerance – which can be defined as 'the amount or level of operational risk that a central bank is prepared to accept, tolerate or be exposed to at any point in time'. With a view to overcome limitations of their traditional approach in this area, centrals banks have started to design more formalized instruments.

As presented in earlier chapters of this book, a central bank's tolerance for *financial* risk is generally defined and approved at the highest level of the organization. In the case of the ECB, this takes place annually at the level of Governing Council in the case of foreign reserve assets, and at the level of the Executive Board of the ECB for the management of the bank's own funds. Yet, as discussed above, central banks by design cannot determine similar quantitative thresholds for their operational risks, as the latter are primarily of a non-financial nature.

In line with similar practice in other not-for-profit institutions, all central banks analysed by a working group of the IORWG in 2006 (see acts of the 2006 conference of the IORWG in Madrid) reported that they express their operational risk tolerance using an indirect approach:

- ex ante, through the definition of qualitative 'guiding principles'; and
- ex post, via ad hoc instructions by their Board of Directors or other relevant committees on ways and means to react to and/or mitigate risks which had materialized.

The limitations of such an approach are well known by all line managers. At the end of the day, how can anybody reasonably request or even expect them to assess the effectiveness of existing controls or to prioritize alternative risk treatment measures without defining the target in the first place?

A few central banks have started to explore ways of reinforcing ex ante guidance on operational risks and controls. In the new ORM framework of the ECB, for instance, the Executive Board formally defines the bank's tolerance for operational risk. The latter consists of a single and fairly detailed risk impact-grading scale, which is linked to high-level guidelines regarding priority risk treatment measures in both normal business conditions and worst-case scenarios.

7.1 Foundation: the risk impact-grading scale

A unique five-level impact-grading scale is used to assess in a consistent manner the severity of business, reputational and financial impact of all types of risk-generating events of the banks. All three categories of impacts are dependent on specifically defined drivers or causal factors (see Figure 13.2).

Most of these causal factors can be expressed according to fairly objective criteria (e.g. non-respect of critical deadline, impact on balance sheet). The level of severity within each risk impact category can be consistently assessed across the bank using a combination of qualitative criteria and quantitative

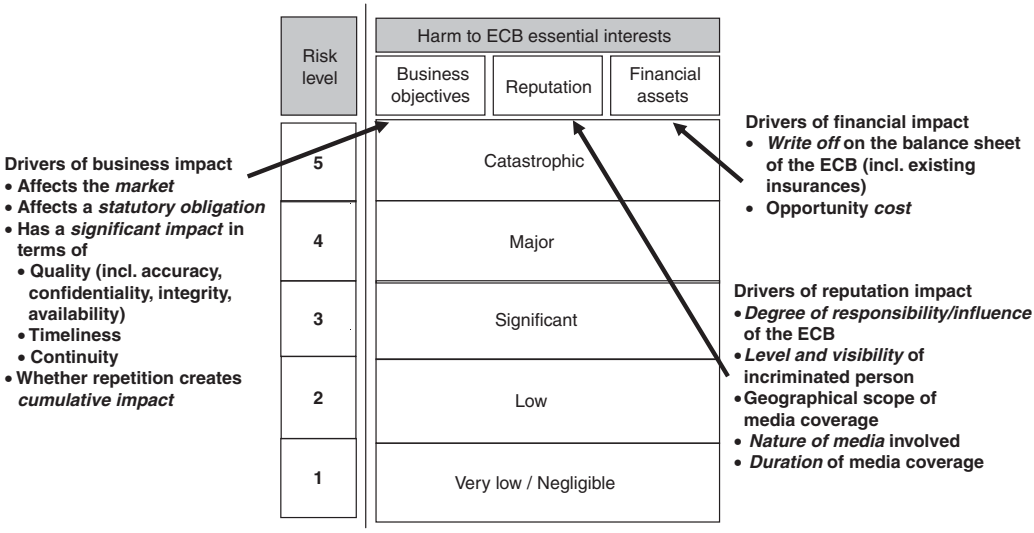

Figure 13.2 Drivers of the risk impact-grading scale of the ECB.

thresholds. In the specific case of impact on reputation, exogenous and subjective elements also play a critical role. As demonstrated in a few much-publicized cases of reputational risk in recent years, perceptions by public opinion tend to prevail over facts – and these perceptions tend to put more emphasis on commonsense and ethical values than on applicable laws and regulations. Legal risk is not represented as a separate risk impact category, as litigation cases ultimately bear a reputational and/or a financial impact. Impacts on reputation related to issues of staff security or confidentiality, availability or integrity of information assets are assessed with consideration of relevant standards and best practices. Business impacts related to staff or information issues can be straightforwardly assessed by considering the most plausible outcomes.

7.2 Implication: risk tolerance guidelines

The operational risk tolerance of the ECB is formalized via a set of ex ante, explicit and high-level guidelines by the Executive Board. The latter provide a prioritization scheme for investment in controls or other risk treatment measures. As shown in Figure 13.3, tolerated levels of risk are expressed considering both *normal business conditions* and *worst-case scenarios*.

The risk tolerance guidelines can be summarized as follows. Risk impacts of level three (under normal business conditions) or four and five (in a

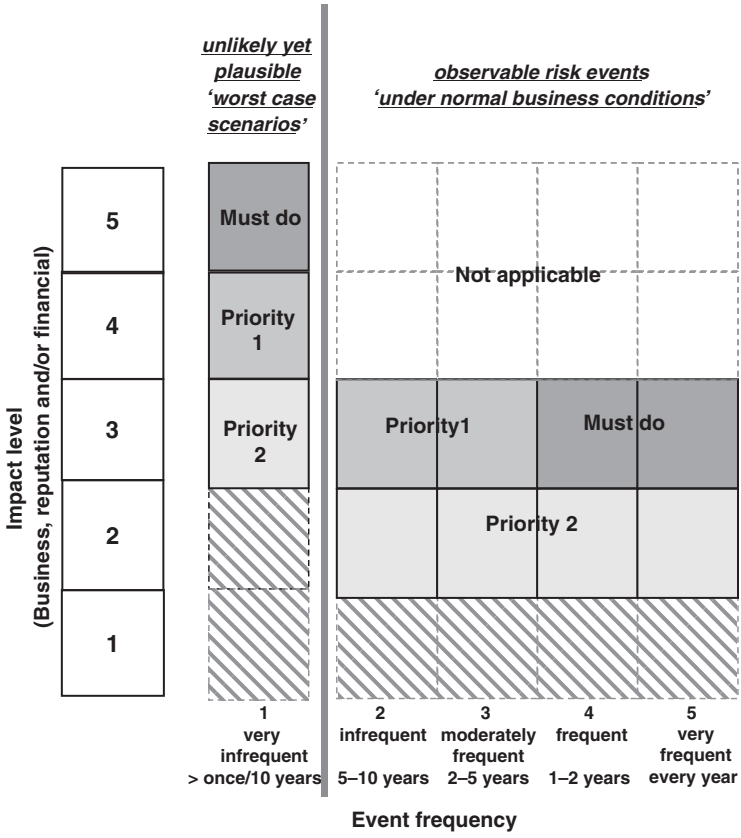

Figure 13.3 Operational risk tolerance: illustrative principles.

worst-case scenario) require implementing priority measures to reduce to the maximum feasible or receiving explicit acceptance by the Executive Board. In order to roughly ascertain the level of severity requiring top management attention, it should be indicated that the financial thresholds used for defining level four and five at the ECB are respectively EUR 1 million and EUR 10 million. Potential risk impacts of level two (in normal business conditions) or three (in a worst-case scenario) require conducting cost–benefit analyses of additional risk treatment measures. And potential risk impacts of level one (in normal business conditions) or one and two (in a worst-case scenario) are considered to be tolerable 'incidents'. From a strict ORM perspective, the latter only require adequate monitoring, yet neither proactive intervention nor reporting. From a broader management perspective, the effectiveness and efficiency of controls related to smaller incidents may justify ad hoc reviews.

8. Top-down self-assessments

How should central banks start implementing ORM? Nowadays, central banks generally recognize that priority should be given to top-down exercises due to their objective and scope. The experience of the ECB is presented to illustrate a possible approach and the related outputs and lessons learned.

8.1 Objective and scope

As demonstrated by a survey of the IORWG (see acts of the 2007 conference in Philadelphia), the vast majority central banks have historically opted to start by a bottom-up approach at the level of individual processes or organizational entities.

After years of implementation, all concerned institutions confirm the benefits of analysing risks and controls in a systematic and detailed manner. However, many of them also stress: (a) the significant cost of conducting bottom-up exercises across all areas; (b) the complexity of aligning and/or reconciling risk information collected on a piecemeal basis; and ultimately (c) the danger of losing sight of the wood for the trees. With the benefit of hindsight, the central bank community nowadays agrees that ORM should start from the top. In essence, top-down approaches achieve two key benefits. They:

- provide an initial and well-calibrated 'big picture' of the critical events bearing the highest risks for the achievement of business objectives, reputation and/or financial assets of the institution; and
- help prioritize subsequent more detailed bottom-up exercises on the most critical processes, functions or organizational entities.

The scope of top-down exercises must facilitate a bird's-eye view on operational risks. In the case of the ECB, the top-down exercise is conducted at the level of the eight core macro-processes (e.g. monetary policy, market operations etc.) of the bank, of its six enabling functions (e.g. communication, IS etc.) as well as for very large projects. The top-down exercise covers all the plausible risk scenarios, be it in a 'worst case' or under 'normal business conditions', which may expose the bank to a risk impact of at least three or more according to the impact-grading scale. From a timing perspective, the top-down exercise is to be conducted each year, as an integral part of the strategy process.

8.2 Approach

At the present juncture, central banks' experience of top-down assessments is probably too recent to describe standard practices and instruments. A notable exception is represented by Bank of Canada, which has accomplished pioneer work in the central banking industry on ways and means of integrating top-down assessments of operational risks with strategic planning. At the ECB, the top-down exercise is centred around two types of workshops: vertical workshops held at the level of each of the core or enabling macro-process of the bank, and horizontal workshops dealing with risk scenarios related to transversal issues of governance (e.g. communication, legal, procurement) and security (information security, physical security, business continuity management).

Defining worst-case operational risk scenarios starts with considering to which extent individual risk items listed in the risk event taxonomy actually apply to a macro-process situation ('*What could go wrong*?' '*Could any of these events ever happen to us*?'). An alternative way to verify whether the universe of worst-case risks considered is comprehensive, is to ponder whether examples of consequences listed in the impact-grading scale would be relevant ('*What would be the worst impact(s) in this area*?') and then 'reverse-engineer' the related worst-case operational risk scenario. In all cases, worst-case scenarios are developed by considering worst-case risk events that have actually happened in partly comparable environments (e.g. governments, public agencies, research centres, faculties, etc.) – thinking of the ECB as a public institution delivering a set of generic functions (e.g. policy making, research/technical advisory, compilation of information, communication of political messages). Based on a mix of primary and secondary research, a database of about 150 relevant worst-case scenarios was compiled by the central ORM team to support the initial top-down assessment and has been continuously updated ever since. Worst-case scenarios are finally tailored to the specific environment of the ECB after due consideration of parameters such as the specific business objectives of the bank (e.g. not for profit dimension), important features of its control environment (e.g. historical 'zero-risk' culture) and predictable changes in the business environment (e.g. transition from Target 1 to Target 2 platform in the area of payment systems). A standard template is completed to describe each worst-case scenario in a comprehensive manner. It provides:

- historical evidence of external catastrophic events which have been considered to establish the plausibility of the worst-case scenario;

- a summary of all the parameters which substantiate the plausibility of the worst-case scenario in the particular environment of the ECB – i.e. specific assumptions pertaining to root causes, failures of mitigation measures, other specific circumstances (e.g. shortage of staff due to holiday season), and ultimate consequences (e.g. leakage/fraud becoming public), etc.;
- an assessment by concerned senior managers of the potential business, reputational and financial impact which the ECB might have to face under the scenario described, including a detailed qualitative justification of this evaluation.

Experts from the central ORM team check consistency of input provided across all business areas and, if required, suggest slight readjustments of assessed risk impacts to ensure overall consistency. During the initial assessment, covering the full universe of about eighty worst-case scenarios and setting the corresponding baseline required a significant one-off effort over a three-month horizon. Fortunately, worst-case scenarios tend to be fairly stable over the medium term – requiring only limited updating during the successive yearly top-down exercises. This should come as no surprise: beyond a few fundamentally new trends in the business and technological environment, and beyond unpredictable events and hazards in the economic and financial conjuncture, the base parameters of operational risks are indeed fairly stable. '*Why do you keep robbing banks*?', a somewhat obstinate criminal was once asked. '*Because it is where the money is*', was the naïve and profound answer.

8.3 Output and lessons learned

In the approach implemented by the ECB, a final report on the top-down ORM assessment is to be produced by the middle of the year. The latter includes an updated heat map charting the status of the bank's key operational risks, a qualitative summary of the key findings, and an appendix including all the compiled worst-case scenario sheets. Expectations regarding initial top-down exercises must be kept at a realistic level. The experience of the ECB indeed shows (or confirms) that only very few totally new risk scenarios emerge from high-level historical analyses, expert brainstorming and management workshops. At first sight, top-down heat map appears to only document widely shared assumptions regarding the concentration of key operational risks in certain macro-processes (e.g. market operations) and horizontal risk categories (e.g. pandemic, massive attack on IS systems). Yet,

very soon, the real benefits of a top-down exercise become much more tangible. ORM workshops with senior management significantly reinforce management awareness of worst-case scenarios – beyond traditional and in-depth knowledge of recurrent incidents. They foster management dialogue and help align fairly diverging individual perceptions regarding the plausibility and potential severity of certain risks (e.g. leak of information) and their relative importance in the global risk portfolio of the bank. And they give new impetus to critical initiatives (e.g. enhance the quality of mission critical IS services to mitigate worst case scenarios related to information confidentiality, integrity and availability; refine business continuity planning arrangements to more proactively address pandemic, strike or other scenarios causing extended unavailability of staff; develop non-IT-dependent contingencies to remedy various crisis situations; leverage enabling technologies such as document management to address risks of information confidentiality, integrity and availability; enhance reputation management through pre-emptive and contingency communication strategy and plans).

9. Bottom-up self-assessments

As a necessary complement to their recent developments regarding top-down exercises, central banks continue to conduct bottom-up exercises at the level of their individual business processes (e.g. 'liquidity management') as well as of horizontal risk and control categories (e.g. 'hacking of information systems'). This section discusses the objective and scope of bottom-up exercises. The experience of the ECB is presented to illustrate a possible approach and to analyse the relationship between bottom-up risk assessments and Business Process Management (BPM) and Total Quality Management (TQM).

9.1 Objective and scope

Central banks have a long tradition of conducting bottom-up exercises to identify and assess current operational risks, define new or enhance existing controls or risk mitigation measures and prioritize related action plans. In the case of the ECB, the scope of these exercises includes all observed or potential operational risk events which may expose the bank to an impact of at least level two according to the impact-grading scale – be it in a worst-case scenario or under normal business conditions. A rolling five-year

programme of bottom-up exercises is prepared by the central ORM team in close cooperation with business areas and approved each year by the Operational Risk Committee in line with the budget life cycle. This programme ensures that all the bank's processes, horizontal risks and related controls of the ECB (including those bearing lower impact levels) will formally be assessed at least every five years, and that key business processes and horizontal risks and related controls, as identified during the top-down exercise, will be assessed in a coordinated manner over the next twelve months. In practice, the programme derived from top-down analysis fosters rational sequencing in ORM implementation and helps avoid that considerations of technical complexity prevail over risk management rationale. Indeed, a few core central banking processes (e.g. related to economic analysis and research) must by nature operate under significant uncertainty, use to some extent incomplete and qualitative information and generally heavily rely on human judgement. As a consequence, these, as well as a few critical management processes (e.g. related to decision making and project management), are typically much more complex to address than transactional processes (e.g. payments, IS operations) and are frequently less covered in early years of ORM implementation.

9.2 Approach

In comparison with top-down exercises, the methodology used in the context of bottom-up exercises typically includes additional elements and generate more granular information.

At the ECB, the step of risk identification includes a quick review of existing processes and underlying assets (people, information systems and infrastructure). The required level of detail of process analysis (i.e. focus on a 'level one' overview of key process steps as opposed to granular 'level three' review of individual activities) is to some extent left to the appreciation of relevant senior managers depending on resource constraints and assessed benefits. The central ORM team ensures the respect of minimal standards (including the use of a standard process documentation tool). The frequency and impact of process incidents is examined by experts and managers. No subjective self-assessment is required for risk events in normal business conditions, as is the case in traditional ORM approaches. By definition, historical facts and/or evidence must have been observed – even though the latter, in most of the cases, are not yet formally compiled in databases.

Plausible worst-case scenarios at process level are defined according to the same methodology as used in the top-down assessment. Specific opportunities to bring normal and worst-case risks in line with the risk tolerance policy are finally discussed. Wherever possible, all these analyses incorporate results from other existing risk management activities to avoid redundancies and relieve management of unnecessary burden. During the next step, a cost–benefit assessment of all identified risk treatment opportunities is performed, using a simple ABC prioritization scheme. Finally, the conclusive steps of a bottom-up assessment include classical action-planning activities as described in various risk management standards. Full documentation of the bottom-up self-assessment via standard templates ensures consistency and re-usability of performed analyses.

9.3 Bottom-up risk assessments vs. BPM and TQM

Many components of the bottom-up exercises are well established and widely shared by the central banking community, and inter alia well documented by the IORWG. Still, two specific aspects of the methodology used at the ECB are worth mentioning, as they underscore the specific value of ORM vs. disciplines such as BPM and TQM.

ORM differs from operations management or BPM. Selectively organizing synergy between all these functions is certainly a good idea. Mixing them up into all-purpose process reviews is a frequent and fatal mistake – ultimately making bottom-up self-assessments costly and cumbersome and hindering the cultural acceptance of risk management. Regarding the specific area of controls, the focus of the ECB is therefore to assess the effectiveness (and to a lesser extent the efficiency) of new or enhanced controls, not the efficiency (and to lesser extent the effectiveness) of all existing controls. The latter approach is a traditional, COSO-based practice which is technically required as a consequence of the following logic flow:

‘*current risk*’ = ‘*inherent risk*’ *minus* ‘*reduced risk through existing controls*’

Yet, as mentioned above, the ECB framework focuses on actual (and potential) risks and how to remedy them. Such an approach by definition takes into account (and thereby implicitly ‘assesses’) the global effectiveness of existing controls and of the general control environment.

Where required, a specific assessment of the general control environment can be performed through use of compliance check lists reflecting relevant process or functional standards.

For specific objectives pertaining much more to process optimization than to ORM, more granular information on the effectiveness and efficiency of individual controls may indeed be required. As benchmarking shows, two approaches are possible in this respect:

- The first, serious and fact based, is traditionally used for instance by internal audit or organization departments. It consists of conducting process or procedure walkthroughs on testing samples to verify ex post how many incidents or anomalies are being detected through given types of verifications or controls.
- The alternative approach, frequently mentioned in traditional ORM frameworks, is arguably always a case of artistic invention: using qualitative self-assessment questionnaires, experts or managers ascertain whether a given control is 'unsatisfactory', 'partially satisfactory' or 'satisfactory'. The problem in this approach is not only that subjective (and naturally partial) opinions of concerned staff should always be challenged by a neutral third party. More fundamentally, the question on control effectiveness itself is meaningless in all frameworks where the objective/target (i.e. the 'risk tolerance') is not specifically predefined. And at any rate, even when the risk tolerance is defined, the question of effectiveness by nature can only be satisfactorily addressed at the level of the full control environment of the institution. By contrast, scoring models used to assess the relative and incremental contribution of controls x, y or z to the current risk situation must by design rely on weighting factors reflecting totally subjective and unverifiable assumptions.

ORM is not total quality management. As a consequence, at the ECB, the management of minor incidents is left out of proactive ORM. Benchmarking evidence show that many central banks have already started a few years ago to compile databases on internal incidents, loss and near-loss events. Over time, incident databases always help improving the reliability and output quality of daily process operations. Even though they almost never produce a sufficient basis for quantitative modelling, they also provide useful reference data points to challenge manager's intuition and to examine key patterns in smaller issues which may as well apply to catastrophes. Yet, systematic and massive capture of incident data has a very significant cost. Reflecting on alternative investment priorities for ORM, it may be useful to keep in mind that daily problems within departments or in interaction with supplier and customer entities in essence constitute cost and quality issues, not risk topics. This explains why, as seen before, the

operational risk tolerance policy of the ECB does not require proactive intervention nor reporting on level one incidents and why the latter are let out of scope of bottom-up self-assessments.

10. ORM governance

Risk management requires a well-defined and integrated governance model. Like most financial organizations, central banks generally make a distinction between the management of financial and operational risks and usually have separate management structures for dealing with these two types of risks.

Regarding the latter, ten *sound practices for management and supervision of operational risk*s have been defined in a seminal paper by the Basel Committee on Banking Supervision (BCBS 2003). Ever since, a few of these practices have been widely adopted by the central banking community (e.g. general sponsorship and oversight function to be assumed by the Executive Board; independent evaluations by internal and external audit functions; key responsibility of line management to implement ORM). For reasons mainly pertaining to individual central banks' size or history, other practices are being implemented under slightly diverging arrangements (e.g. composition of the committee specifically in charge of ORM; establishment or not of a dedicated ORM officer; relative positioning of the central ORM function vs. the business continuity function; precise level of decentralization of ORM activities in business areas etc.). Overall, in most central banks, a key challenge is still to organize the convergence of all disciplines related to operational risks and control (including business continuity, physical security, information confidentiality etc.) and allow for an integrated management of the related risk portfolio.

The new ORM governance model adopted by the ECB in September 2007 comprises the following elements: an Operational Risk Committee (ORC), staffed with seven senior managers of the bank, deals with strategic/medium-term topics. The key mission of the ORC is to stimulate and oversee the development, implementation and maintenance of all disciplines related to operational risks. To that effect, the specific responsibilities of the ORC are to endorse the relevant policy frameworks and strategies; assess the portfolio of risks and the effectiveness and efficiency of treatments of operational risks across the ECB; plan and monitor all related activities; foster the development of risk management culture in the ECB as well as ESCB – and

Eurosystem-wide through appropriate measures and to inform the Executive Board periodically about the status of ORM. Required input for strategic decision-making by the ORC is prepared, at a more tactical level, by an informal network of operational risk managers and risk experts which work on request of the ORC in the form of ad hoc taskforces. Efficiency of ORM decision making is enhanced by addressing dossiers at the first competent level and limiting representation in taskforces to key business and functional stakeholders.

A central and integrated ORM and BCM team, hosted by the Organizational Planning Division of the bank, acts as knowledge broker in charge of cross-fertilizing best practices across business areas and specialized risk disciplines. On top of coordinating all relevant activities (including response to incidents), the team assumes classical central activities pertaining to external benchmarking and cooperation, methodological maintenance (e.g. library of controls) and development (e.g. integration of ORM databases and tools), proactive monitoring of and advisory to business areas, consolidated reporting, and secretariat of the ORC.

The responsibility and accountability of line managers in the implementation of ORM in their respective business areas is confirmed and a decentralized function of ORM coordinators is further formalized – without creating additional resource requirements let alone new positions. Beyond participation to mandatory top-down and bottom-up exercises, line managers, with the support of ORM coordinators, manage their operational risks as part of daily operations. In particular, they are expected to proactively consider the specific risk implications of defined trigger-point events (e.g. assumed new service responsibility; significant staffing or management change; recent centralization/de-centralization of business process or technology; introduction of new software or hardware; hired new vendor; identified issue during contingency test; specific findings in internal and external audits etc.) where the benefits of ORM analyses are particularly obvious.

11. KRIs and ORM reporting

The ultimate function of ORM is not to report on the status of operational risks but to provide insightful support for management decisions on required actions and investments. The present section discusses the gap between theory and practice in this respect and presents current developments in the ECB.

11.1 Theory vs. practice

Key risk indicators (KRIs) and ORM reporting represent a particularly challenging area – arguably one in which sound and nice-sounding advice is commonplace, yet where most central banks are still struggling with fundamental issues.

Take a glance at handbooks, risk management standards and consultant presentations. Each of them includes a more or less compelling lecture and colourful charts on what everybody knows: key risk indicators are essential. They *should be* SMART (Specific, Measurable, Achievable, Relevant, Timebound). They should be tightly linked to the institution's strategy and risk policy. They should be simple without being simplistic. They should primarily be 'leading' indicators (i.e. have early-warning qualities), trigger adequate actions based on well-tested 'safe', 'cautionary', and 'warning' thresholds, be multi-dimensional, have a positive 'value-to-burden' relationship, be easy to benchmark, etc. – just add to the list. The following section on ORM reporting is typically just as enlightening: integrated management reports *should* focus on key risks relevant to various management layers, provide brief and relevant progress reports on action plans and avoid extraneous detail. Periodicity should be adapted to various risks types etc.

And yet meet ORM practitioners of central banks and candidly discuss their experience. As documented in the material compiled by the IORWG, first-generation initiatives understandably could not meet the aforementioned, daunting expectations: while most central banks track a few indicators in a few parts of their organization, they still have not been in a position to put in place a consistent concept nor a formal implementation programme for KRIs. A few banks may have started ambitious cooperation initiatives to co-develop KRI libraries – yet they are now mystified by a database of hundreds of potential risk indicators. In many cases, redundant regular and ad hoc reports on risk, security and control issues, using heterogeneous risk taxonomies and grading scales are randomly produced at various levels of the organization. Board reports frequently include more than fifty or even a hundred risk items – and a list of the top five, ten or twenty top risks of the bank is available in only in very few central banks.

11.2 Current developments at the ECB

Drawing lessons learned from many central bank colleagues, the ECB opted to start working on KRIs and ORM reporting only after all other elements of

the framework had been properly developed and tested and once insights from the top-down exercise helped specify priority content for senior management. Current developments try to transpose the few aspects where the central banking community has reached common conclusions.

Regarding KRIs, the ECB opted to focus on the few metrics which, judging by other banks' experiences, appear to capture most of the value of early risk prediction or detection. Some of these KRIs are for instance:

- indicators of HR root causes of errors and frauds: e.g. ratio of screened applications for predefined sensitive jobs; ratios of predefined jobs with critical skills with appropriate succession planning; trends in consumption of staff training budget, in job tenure, in staff turnover, in overtime, in use of temporary staff, in staff satisfaction etc.;
- indicators of process deficiencies: e.g. trends in number and type of errors in input data and reports; transactions requiring corrections or reconciliation; average aging of outstanding issues; unauthorized activities; transaction delays; counterfeiting rates; customer complaints; customer satisfaction ratings; financial losses; aging structure of pending control issues etc.;
- indicators of IS vulnerability: e.g. trends in system response time; trouble tickets; outages; virus or hacker attacks; detected security or confidentiality breaches etc.

An incident-tracking database tool, feeding into relevant KRIs, will be implemented as from 2009, starting in priority in transaction-intensive areas (e.g. market operations, payment systems, IS). This tool will be used to gather adequate experience in the constitution and management of incident databases and to provide an intermediary solution, until market solutions for ORM (including capture, assessment, monitoring and reporting of operational risks) deliver true value for a medium-size, not-for profit institution like the ECB. In the area of physical security, where prediction and detection of significant external threats is of prominent importance, an approach limited to KRIs is clearly insufficient. As a consequence, this function continues to develop, maintain and continuously implement more advanced monitoring instruments (e.g. intelligence management databases, scoring systems pertaining to the capacity and motivation of potential external aggressors etc.). As far as ORM reporting is concerned, the initial focus of efforts is on top-management reporting. Best practices in the private sector, which allow for representations of quantitative concentrations of financial losses in operational risk portfolios, often confirm and help visualize manager's intuition: ORM follows a Pareto law. About ten to

twenty operational risks, which truly require Board attention, represent a significant proportion (e.g. 50 per cent) of the value-at-risk linked to ORM. Monitoring a second cluster of about eighty to ninety additional risk items, through reporting drill-downs at business area level, is typically sufficient to cover most (e.g. 80 per cent) of the institution's value at risk. By contrast a myriad of existing or potential smaller incidents, which accounts for the remaining 20 per cent value at risk, only justify tracking at base levels of the organization.

In the reporting scheme currently under development, top operational risks will be reported to the Executive Board on a bi-yearly basis and to the Operational Risk Committee on a bi-monthly basis.

From a content and format perspective, these streamlined ORM reports will include:

- a global picture of the current level of operational risks (similar to the 'heat map' produced as a result of the top-down exercise), listing the top ten risks (Executive Board report) and top thirty risks (ORC report) and showing the respective trends vs. the previous period;
- the status of twenty key risk indicators of the ECB;
- a rating of the general control environment by macro-process and horizontal risk category;
- an overview of progress in ORM implementation roll-out by macro-process and horizontal risk category;
- a qualitative synthesis of achievement improvements in key areas of controls.

In line with best practices, this dashboard will be supported by a tool allowing for user-friendly visual representations and simulations.

Equally as from 2009, standard and 'light' yearly business area reports on ORM will be implemented. If deemed advisable after detailed examination of the pros and cons, the latter will be complemented by a management assertion letter. The latter is a best-practice instrument which was originally created in the context of the Sarbanes-Oxley legislation and has since been adopted by public institutions like the European Commission. In short, the assertion letter is an annual declaration whereby senior managers individually attest the effectiveness and efficiency of the key internal controls in light of the key risks identified in their business area, point out reservations, and mention implemented or planned improvement measures. Frequently reported benefits of such a scheme are increased awareness on operational risks and controls and reinforced management responsibility and accountability. Indeed, external practice also suggests that a horizon of

at least one year of ORM implementation is required to provide managers with a sound base of information, knowledge and experience on key operational risks and controls.

12. Conclusions

Over the past twenty years, most central banks have historically developed and built separate risk management frameworks for various business and functional risks, then generally adopted frameworks like COSO to introduce some homogeneity, and more recently attempted to selectively transpose more sophisticated quantitative models of the commercial banking sector.

More recently, after having achieved very significant progress in specific areas (e.g. defining a taxonomy of operational risk events, conducting a number of bottom-up self-assessments, transposing the sound practices of ORM governance), central banks have dramatically increased inter-professional benchmarking and cooperation. In various forums, they now launch next generation developments with a view to reduce subjectivity in risk assessments, integrate risk reports to support management decisions and alleviate costs of ORM implementation.

With the benefit of accumulated hindsight and lessons learned from our central banking colleagues, and reviewing the more recent developments and provisional achievements in the ECB, we can only confirm that a paradigm shift is both necessary and possible in this area.

Nowadays, there is only little merit to reformulate consultants' ritual recommendations such as 'getting top management commitment', 'putting first things first', 'keeping it simple', 'managing expectations', 'delivering value to the customers' or 'achieving quick wins'. Regrettably, such principles prove to be less actionable key success factors to guide action ex ante than simple performance criteria to evaluate results ex post. In our view, what ORM managers and experts perhaps mostly need is to use a sound combination of analytical rigour, common sense, courage, discipline and diplomacy. Only such virtues can help them carefully steer their institutions away from conservatism ('*Why change? Bank X or Y does just the same as us*') and/or flavour-of-the month concepts and gimmicks ('*The critical success factor is to implement KRIs* – or: *a balanced scorecard / a management dashboard / fully documented processes and procedures / an integrated ORM solution / a risk awareness programme / a global Enterprise Risk Management perspective* etc.').

Looking ahead, the critical challenge appears to be, as often the case in management matters, one about people and values. From senior management down to the grass-roots level, new ORM champions and role models are required to develop and nurture a new organizational culture and respond to three key demands: Serving the needs and aspirations of highly educated and experienced service professionals, ORM cannot impose intrusive transparency, but must credibly encourage individuals and teams to openly disclose own mistakes and near misses. Faced with an increasingly complex and uncertain business environment, ORM cannot just 'build awareness' on operational risks but must foster proactive attitudes of risk detection, prevention and mitigation. And spurred by new constraints of effectiveness and efficiency, ORM must fundamentally reorientate the traditional zero-risk culture of central bankers towards a culture of explicit risk tolerance and of cost–benefit assessments of controls.

The ORM journey, it seems, is only starting.

References

Acharya, V., Bharath, S. T., Srinivasan, A. 2003. 'Understanding the recovery rates on defaulted securities', CEPR Discussion Paper 4098.

Acworth, P., Broadie, M. and Glasserman, P. 1997. 'A comparison of some Monte Carlo and quasi-Monte Carlo techniques for option pricing', in P. Hellekalek, H. Niederreiter (eds.), *Monte Carlo and Quasi-Monte Carlo Methods 1996.*, Lecture Notes in Statistics vol. 127. New York: Springer-Verlag, pp. 1–18.

Akeda, Y. 2003. 'Another interpretation of negative Sharpe ratio', *Journal of Performance Measurement* 7(3): 19–23.

Alexander, C. 1999. *Risk management and analysis: Measuring and modeling financial risk.* New York: Wiley.

Alexander, G. J. and Baptista, A. M. 2003. 'Portfolio performance evaluation using value at risk', *Journal of Portfolio Management* 29: 93–102.

Almgren, R. and Chriss, N. 1999. 'Value under liquidation', *Risk* 12: 61–3.

Altman, E. I. and Kishore, V. M. 1996. 'Almost everything you wanted to know about recoveries on defaulted bonds', *Financial Analysts Journal* 52(6): 57–64.

Altman, E. I., Brady, B., Resti A. and Sironi, A. 2005a. 'The link between default and recovery rates: Theory, empirical evidence and implications', *The Journal of Business* 78(6): 2203–28.

Altman, E. I., Resti, A. and Sironi A. (eds.) 2005b. *Recovery risk: the next challenge in credit risk management.* London: Risk Books.

Altman, E. I., Resti, A. and Sironi, A. 2004. 'Default recovery rates in credit risk modelling: A review of the literature and empirical evidence', *Economic Notes* 33: 183–208.

Amato, J. D. and Remolona, E. M. 2003. 'The credit spread puzzle', *BIS Quarterly Review* 12/2003: 51–63.

Amihud, Y. and Mendelson, H. 1991. 'Liquidity, maturity and the yields on U.S. Treasury securities', *Journal of Finance* 46: 1411–25.

Andersson, F., Mausser, H., Rosen, D. and Uryasev, S. 2001. 'Credit risk optimisation with conditional Value-at-Risk criterion', *Mathematical Programming, Series B* 89: 273–91.

Ankrim, E. M. and Hensel, C. R. 1994. 'Multicurrency performance attribution', *Financial Analysts Journal* 50(2): 29–35.

Apel, E. 2003. *Central banking systems compared: The ECB, the pre-euro Bundesbank and the Federal Reserve System.* London and New York: Routledge.

Artzner, P., Delbaen, F., Eber, J.-M. and Heath, D. 1999. 'Coherent measures of risk', *Mathematical Finance* 9: 203–28.

Asarnow, E. and Edwards, D. 1995. 'Measuring loss on defaulted bank loans: A 24-year study', *Journal of Commercial Lending* 77(7): 11–23.

Association of Insurance and Risk Managers. 2002. 'A risk management standard', www.theirm.org/publications/documents/Risk_Management_Standard_030820.pdf.

Bacon, C. 2004. *Practical portfolio performance measurement and attribution*. London: Wiley.

Bagehot, W. 1873. *Lombard Street: A description of the money market*. London: H.S. King.

Bakker, A. F. P. and van Herpt, I. R. Y. 2007. *Central bank reserve management: new trends, from liquidity to return*. Cheltenham: Edward Elgar.

Bandourian, R. and Winkelmann, K. 2003 'The market portfolio', in Litterman (ed.), pp. 91–103.

Bangia, A., Diebold, F., Schuermann, T. and Stroughair, J. 1999. 'Making the best of the worst', *Risk* 10: 100–3.

Bank for International Settlements. 1999. *Implications of repo markets for central banks*. Basel: Bank for International Settlements, www.bis.org/publ/cgfs10.pdf.

Bank for International Settlements. 2005. 'Zero-coupon yield curves: Technical documentation', BIS Papers 25.

Bank of Japan. 2004. 'Guidelines on eligible collateral', www.boj.or.jp/en/type/law/ope/yoryo18.htm.

Bardos, M., Foulcher, S. and Bataille, É. (eds.) 2004. *Les scores de la Banque de France: Méthode, résultats, applications*. Paris: Banque de France, Observatoire des entreprises.

Basel Committee on Banking Supervision. 1998a. 'Framework for internal control systems in banking organizations', Bank for International Settlements 09/1998, www.bis.org/publ/bcbs40.pdf.

Basel Committee on Banking Supervision. 1998b. 'Operational risk management', Bank for International Settlements 09/1998, www.bis.org/publ/bcbs42.pdf.

Basel Committee on Banking Supervision. 2000a. 'Credit ratings and complementary sources of credit quality information', BCBS Working Papers 3, www.bis.org/publ/bcbs_wp3.pdf.

Basel Committee on Banking Supervision. 2000b. 'Principles for the management of credit risk', Bank for International Settlements 09/2000, www.bis.org/publ/bcbs54.pdf.

Basel Committee on Banking Supervision. 2001a. 'The new Basel capital accord', BIS Consultative document, www.bis.org/publ/bcbsca03.pdf.

Basel Committee on Banking Supervision. 2001b. 'The internal ratings-based approach', BIS Consultative Document, www.bis.org/publ/bcbsca05.pdf.

Basel Committee on Banking Supervision. 2002. 'The quantitative impact study for operational risk: Overview of individual loss data and lessons learned', Bank for International Settlements 01/2002, www.bis.org/bcbs/qisopriskresponse.pdf.

Basel Committee on Banking Supervision. 2003. 'Sound practices for the management and supervision of operational risk', Bank for International Settlements 07/2003, www.bis.org/publ/bcbs96.pdf.

Basel Committee on Banking Supervision. 2004. 'Principles for the management and supervision of interest rate risk', Bank for International Settlements 07/2004, www.bis.org/publ/bcbsca09.pdf.

Basel Committee on Banking Supervision. 2006a. *Enhancing corporate governance for banking organisations*. Basel: Bank for International Settlements, www.bis.org/publ/bcbs122.pdf.

Basel Committee on Banking Supervision. 2006b. *Basel II: International convergence of capital measurement and capital standards: A revised framework – Comprehensive version*. Basel: Bank for International Settlements.

Basel Committee on Banking Supervision. 2006c. 'Core principles for effective banking supervision', Bank for International Settlements 10/2006, www.bis.org/publ/bcbs129.pdf.

Basel Committee on Banking Supervision. 2006d. 'Studies on credit risk concentration: An overview of the issues and a synopsis of the results from the Research Task Force project', BCBS Working Paper 15, www.bis.org/publ/bcbs_wp15.pdf.

BCBS. *See* Basel Committee on Banking Supervision.

Berger, A., Davies, S. and Flannery, M. 1998. 'Comparing market and regulatory assessments of bank performance: Who knows what when?', FEDS Working Paper 03/1998.

Berk, J. B. and Green, R. C. 2002. 'Mutual fund flows and performance in rational markets', NBER Working Paper 9275.

Bernadell, C., Cardon, P., Coche, J., Diebold, F. X. and Manganelli, S. (eds.) 2004. *Risk management for central bank foreign reserves*. Frankfurt am Main: European Central Bank.

Bernadell, C., Coche, J. and Nyholm, K. 2005. 'Yield curve prediction for the strategic investor', ECB Working Paper Series 472.

Bertsekas, D. 1999. *Nonlinear programming*. 2nd edn. Belmont: Athena Scientific.

Bertsimas, D. and Lo, A. 1998. 'Optimal control of execution costs', *Journal of Financial Markets* 1: 1–50.

Bester, H. 1987. 'The Role of Collateral in Credit Markets with Imperfect Information', *European Economic Review* 31: 887–99.

Bindseil, U. 2004. *Monetary policy implementation*. Oxford: Oxford University Press.

Bindseil, U. and Nyborg, K. 2008. 'Monetary policy implementation', in X. Freixas, P. Hartmann and C. Mayer (eds.), *Financial markets and institutions: a European perspective*. Oxford: Oxford University Press.

Bindseil, U. and Papadia, F. 2006. 'Credit risk mitigation in central bank operations and its effects on financial markets: The case of the Eurosystem', ECB Occasional Paper Series 49.

Bindseil, U., Camba-Mendez, C., Hirsch, A. and Weller, B. 2006. 'Excess reserves and the implementation of monetary policy of the ECB', *Journal of Policy Modelling* 28: 491–510.

Bindseil, U., Manzanares, A. and Weller, B. 2004a. 'The role of central bank capital revisited', ECB Working Paper Series 392.

Bindseil, U., Nyborg, K. and Strebulaev, I. 2004b. 'Bidding and performance in repurchase auctions: evidence from ECB open market operations', CEPR Discussion Paper 4367.

BIS. *See* Bank for International Settlements.

Black, F. and Litterman, R. 1992. 'Global portfolio optimization', *Financial Analysts Journal* 48: 28–43.

Black, F. and Scholes, M. 1973. 'The pricing of options and corporate liabilities', *Journal of Political Economy* 81: 637–59.

Black, F., Derman, E. and Toy, W. 1990. 'A one factor model of interest rates and its application to the Treasury bond options', *Financial Analysts Journal* 46: 33–9.

Blejer, M. and Schumacher, L. 2000. 'Central banks use of derivatives and other contingent liabilities: Analytical issues and policy implications', IMF Working Paper 66.

Blenck, D., Hasko, H., Hilton, S. and Masaki, K. 2001. 'The main features of the monetary policy frameworks of the Bank of Japan, the Federal Reserve and the Eurosystem', BIS Paper 9: 23–56.

Bliss, R. 1997. 'Movements in the term structure of interest rates', *Federal Reserve Bank of Atlanta Economic Review* 82(4): 16–33.

Bluhm, C., Overbeck, L. and Wagner, C. 2003. *An introduction to credit risk modeling*. London: Chapman & Hall.

Bonafede, J. K., Foresti, S. J. and Matheos, P. 2002. 'A multi-period linking algorithm that has stood the test of time', *Journal of Performance Measurement* 7(1): 15–26.

Bookstaber, R. and Clarke, R. 1984. 'Option portfolio strategies: Measurement and evaluation', *Journal of Business* 57(4): 469–92.

Borio, C. E. V. 1997. 'The implementation of monetary policy in industrial countries: A survey', BIS Economic Paper 47.

2001. 'A hundred ways to skin a cat: Comparing monetary policy operating procedures in the United States, Japan and the euro area', BIS Paper 9: 1–22.

Brennan, M. and Schwartz, E. 1979. 'A continuous time approach to the pricing of bonds', *Journal of Banking and Finance* 3: 133–55.

1982. 'An equilibrium model of bond pricing and test of market efficiency', *Journal of Financial and Quantitative Analysis* 17(3): 301–29.

Brickley, J. A., Smith, C. W. Jr. and Zimmerman, J. L. 2007. *Managerial economics and organizational structure*. Boston: McGraw-Hill.

Brinson, G. P. and Fachler, N. 1985. 'Measuring non-U.S. equity portfolio performance', *Journal of Portfolio Management* 1 1(3): 73–76.

Brinson, G. P., Hood, L. R. and Beebower, G. L. 1986. 'Determinants of portfolio performance', *Financial Analysts Journal* 42(4): 39–44.

Brinson, G. P., Singer, B. D. and Beebower, G. L. 1991. 'Determinants of portfolio performance II: An update', *Financial Analysts Journal* 47(3): 40–8.

British Standard Institutions. 2006. *Business continuity management – Part 1: Code of practice*. United Kingdom: British Standards Institutions.

Bucay, N. and Rosen, D. 1999. 'Credit risk of an international bond portfolio: a case study', *Algo Research Quarterly* 2(1): 9–29.

Buchholz M., Fischer, B. R. and Kleis, D. 2004. 'Attributionsanalyse für Rentenportfolios', *Finanz Betrieb* 7–8: 534–51.

Buhl, H. U., Schneider, J. and Tretter, B. 2000. 'Performanceattribution im Private Banking', *Die Bank* 40(5): 318–323.

Buiter, W. and Sibert, A. 2005. 'How the ECB's open market operations weaken fiscal discipline in the eurozone (and what to do about it)', CEPR Discussion Paper 5387.

Burnie, J. S., Knowles, J. A. and Teder, T. J. 1998. 'Arithmetic and geometric attribution', *Journal of Performance Measurement* 3(1): 59–68.

Burns, W. and Chu, W. 2005. 'An OAS Framework for portfolio attribution analysis', *Journal of Performance Measurement* 9(4): 8–20.

Business Continuity Institute. 2007. *Good practice guidelines*. United Kingdom: British Standards Institutions, www.thebci.org/CHAPTER2BCIGPG07.pdf.

Caballero, R. J. and Krishnamurthy, A. 2007. 'Collective risk management in a flight to quality episode', NBER Working Paper 12896.

Caflisch, R. E., Morokoff, W. and Owen, A. 1997. 'Valuation of mortgage-backed securities using brownian bridges to reduce effective dimension', *The Journal of Computational Finance* 1(1): 27–46.

Calvo, G. and Leiderman, L. 1992. 'Optimal inflation tax under precommitment: Theory and evidence', *American Economic Review* 82: 174–94.

Campbell, J. O. and Viceira, L. M. 2002. *Strategic Asset Allocation*. New York: Oxford University Press.

Campbell, J.O., Lo, A. and Mackinley, A. C. 1997. *The Econometrics of Financial Markets*. Princeton: Princeton University Press.

Campbell, S. D. 2006. 'A review of backtesting and backtesting procedures', *The Journal of Risk* 9(2): 1–18.

Campisi, S. 2000. 'Primer on fixed income performance attribution', *Journal of Performance Measurement* 4(4): 14–25.

2002. 'While we expound on theory, have we forgotten practice?', *Journal of Performance Measurement* 7(2): 7–8.

Carhart, M. M. 1997. 'On persistence in mutual fund performance', *Journal of Finance* 52(1): 57–82.

Carino, D. R. 1999. 'Combining attribution effects over time', *Journal of Performance Measurement* 3(4): 5–14.

Carty, L. V. and Lieberman, D. 1996. 'Defaulted bank loan recoveries', Moody's Investors Service, Global Credit Research special report 11/1996, www.moodys.com.

Catarineu-Rabell E., Patricia Jackson, and D. Tsomocos, 2003, 'Procyclicality and the New Basel Accord-Banks' Choice of Loan Rating System', Bank of England Working Paper No. 181.

CGFS. *See* Committee on the Global Financial System.

Chance, D. M. and Jordan, J. V. 1996. 'Duration, convexity, and time as components of bond returns', *Journal of Fixed Income* 6: 88–96.

Chappell, D. 2004. *Enterprise service bus*. New York: O'Reilly.

Chartered Financial Analyst Institute 2006. *Global investment performance standards (GIPS) handbook*. 2nd edn. Charlottesville: Chartered Financial Analyst Institute.

Christensen, P.O. and Sorensen, B. G. 1994. 'Duration, convexity, and time value', *Journal of Portfolio Management* 20: 51–60.

Claessens, S. and Kreuser, J. 2007. 'Strategic foreign reserves risk management: an analytical framework', *Annals of Operations Research* 152(1): 79–113.

Coase, R. 1937. 'The nature of the firm', *Economica* 4: 386–405.

1960. 'The problem of social cost', *The Journal of Law and Economics*, 3, 1–44.

Cochrane J. H. 2001. *Asset pricing*. Princeton: Princeton University Press, chapter 20.

Colin, A. 2005. *Fixed income attribution*. London: Wiley.

Committee of Sponsoring Organizations of the Treadway Commission. 2004. 'Enterprise risk management – an integrated framework', Committee of Sponsoring Organizations of the Treadway Commission 2004/09, www.coso.org/Publications/ERM/COSO_ERM_Executive Summary.pdf.

Committee on Payment and Settlement Systems. 2000. *The contribution of payment systems to financial stability*. Basel: Bank for International Settlements, www.bis.org/publ/cpss41.pdf.

Committee on Payment and Settlement Systems. 2006. *Cross-border collateral arrangements*. Basel: Bank for International Settlements, www.bis.org/publ/cpss71.pdf.

Committee on the Global Financial System. 1999. 'Market liquidity: Research findings and selected policy implications', Bank for International Settlements, www.bis.org/publ/cgfs11overview.pdf.

Committee on the Global Financial System. 2001. *Collateral in wholesale financial markets: recent trends, risk management and market dynamics*. Basel: Bank for International Settlements, www.bis.org/publ/cgfs17.pdf.

Committee on the Global Financial System. 2005. *The role of ratings in structured finance: issues and implications*. Basel: Bank for International Settlements, www.bis.org/publ/cgfs23.pdf.

Connor, G. and Korajczyk, R. 1986. 'Performance measurement with the arbitrage pricing theory: A new framework for analysis', *Journal of Financial Economics* 15(3): 373–94.

Coppens, F., González, F. and Winkler, G. 2007 'The performance of credit rating systems in the assessment of collateral used in Eurosystem monetary policy operations', ECB Occasional Paper Series 65.

COSO. *See* Committee of Sponsoring Organizations of the Treadway Commission.

Cossin, D. and Pirotte, H. 2007. *Advanced credit risk analysis: Financial approaches and mathematical models to assess, price and manage credit risk*. 2nd edn. New York: Wiley.

Cossin, D., Gonzalez, F., Huang, Z. and Aunon-Nerin, D. 2003. 'A framework for collateral risk control determination', ECB Working Paper 209.

Cotterill, C. H. E. 1996. *Investment performance mathematics: Time weighted and dollar weighted rates of return*. Hoboken: Metri-Star Press.

Counterparty Risk Management Policy Group I. 1999. 'Improving counterparty risk management practices', Counterparty Risk Management Policy Group 06/1999, financialservices.house.gov/banking/62499crm.pdf.

Counterparty Risk Management Policy Group. II 2005. 'Towards greater financial stability: A private sector perspective', The Report of the Counterparty Risk Management Policy Group II 07/2005, www.crmpolicygroup.org/docs/CRMPG-II.pdf.

Cox, L. C., Ingersoll, J. E. and Ross, S. A. 1985. 'A theory of the term structure of interest rates', *Econometrica* 53(2): 385–407.

CPSS. *See* Committee on Payment and Settlement Systems.

Cranley R. and Patterson, T. N. L. 1976. 'Randomization of number theoretic methods for multiple integration', *SIAM Journal of Numerical Analysis* 13(6): 904–14.

Crouhy, M., Galai, D. and Mark, R. 2001. *Risk management*. New York: McGraw-Hill.

Cruz, M. 2002. *Modeling, measuring and hedging operational risks*. New York: Wiley.

Cubilié, M. 2005. 'Fixed income attribution model', *Journal of Performance Measurement* 10 (2): 49–63.

Dalton, J. and Dziobek, C. 2005. 'Central bank losses and experiences in selected countries', IMF Working Paper 05/72.

Daniel, F., Engert, W. and Maclean, D. 2004. 'The Bank of Canada as lender of last resort', *Bank of Canada Review* Winter 2004–05: 3–16.

Danmarks Nationalbank. 2004. *Financial management at Danmarks Nationalbank*. Copenhagen: Danmarks Nationalbank.

Davies, O. and Laker, D. 2001. 'Multiple-period performance attribution using the Brinson model', *Journal of Performance Measurement* 6(1): 12–22.

De Almeida da Silva Junior, A. F. 2004. 'Performance attribution for fixed income portfolios in Central Bank of Brazil international reserves management', in C. Bernadell, P. Cardon, J. Coche, F. X. Diebold and S. Manganelli (eds.), *Risk management for central bank foreign reserves*. Frankfurt am Main: European Central Bank, pp. 315–29.

de Beaufort, R., Benitez, S. and Palomino, F. 2002. 'The case for reserve managers to invest in corporate debt', *Central Banking* 12(4): 79–87.

Deutsche Bank 2005. 'Annual review 2005', annualreport.deutsche-bank.com/2005/ar/servicepages/downloads.php.

Deutsche Bundesbank. 2006. 'Synopsis of the Deutsche Bundesbank's procedure for analysing credit standing', www.bundesbank.de/gm/gm_sicherheiten_downloads.en.php.

Diamond, D. and Dybvig, P. 1983. 'Bank runs, deposit insurance and liquidity', *Journal of Political Economy* 91: 401–19.

Diebold, F.X. and Li, C. (2006), 'Forecasting the term structure of Government Bond Yields', *Journal of Econometrics* 130: 337–64.

Dietz, P. 1966. *Pension funds: Measuring investment performance*. New York: Free Press.

Dowd, K. 2005. *Measuring market risk*. 2nd edn. Chichester: Wiley Finance.

Duffie, D. and Singleton, K. J. 2003. *Credit risk: Pricing, measurement and management*. Princeton: Princeton University Press.

Dynkin, L and Hyman, J. 2002. 'Multi-factor risk models and their applications' in F. Fabozzi (ed.), *Interest rate, term structure, and valuation modeling*. Hoboken: Wiley, pp. 241–94.

2004. 'Multi-factor risk analysis of bond portfolios' in C. Bernadell, P. Cardon, J. Coche, F. X. Diebold and S. Manganelli (eds.), *Risk management for central bank foreign reserves*. Frankfurt am Main: European Central Bank, pp. 201–21.

2006. 'Multi–factor risk models and their applications' in Fabozzi, Martellini and Priaulet (eds.) pp. 195–246.

Dynkin, L., Gould, A., Hyman, J., Konstantinovsky, V. and Phelps, B. 2006. *Quantitative management of bond portfolios*. Princeton: Princeton University Press.

ECB. *See* European Central Bank.

Elton, E. J, Gruber, M. J, Brown, S. J. and Goetzmann, W. N. 2003. *Modern portfolio theory and investment analysis*, Hoboken: Wiley.

Engström, S. 2004. 'Does active portfolio management create value? An evaluation of fund managers' decisions', SSE/EFI Working Paper Series in Economics and Finance 553.

Ernhagen, T., Vesterlund, M. and Viotti, S. 2002. 'How much equity does a central bank need?', *Sveriges Riksbank Economic Review* 2/2002: 5–18.

Esseghaier, Z., Lal, T., Cai, P. and Hannay, P. 2004. 'Yield curve decomposition and fixed-income attribution', *Journal of Performance Measurement* 8(4): 30–45.

European Central Bank. 2004a. 'Risk mitigation measures in Eurosystem credit operations', *Monthly Bulletin* 05/2004: 71–9.

European Central Bank. 2004b. 'The euro bond market study', www.ecb.int/pub/pdf/other/eurobondmarketstudy2004en.pdf.

European Central Bank. 2006a. 'Portfolio management at the ECB', *Monthly Bulletin* 4/2006: 75–86.

European Central Bank. 2006b. 'The implementation of monetary policy in the euro area – General documentation of Eurosystem monetary policy instruments and procedures', General Documentation 09/2006, www.ecb.int/pub/pdf/other/gendoc2006en.pdf

European Central Bank. 2007a. 'Euro Money Market Study 2007', www.ecb.europa.eu/pub/pdf/other/euromoneymarketstudy200702en.pdf.

European Central Bank. 2007b. 'The collateral frameworks of the Federal Reserve System, the Bank of Japan and the Eurosystem', *Monthly Bulletin* 10/2007: 85–100.

Ewerhart, C. and Tapking, J. 2008. 'Repo markets, counterparty risk, and the 2007/2008 liquidity crisis', ECB Working Paper Series 909.

Fabozzi, F. J., Martellini, L. and Priaulet, P. 2006. *Advanced bond portfolio management.* Hoboken: Wiley.

Fama, E. F. and French, K. R. 1992. 'The cross–section of expected stock returns', *Journal of Finance* 47(2): 427–65.

1993. 'Common risk factors in the returns on stocks and bonds', *Journal of Financial Economics* 33: 3–56.

1995. 'Size and book-to-market factors in earnings and returns', *Journal of Finance* 50(1): 131–55.

1996. 'Multifactor explanations of asset pricing anomalies', *Journal of Finance* 51(1): 55–84.

Federal Reserve Bank of New York. 2007. 'Domestic open market operations during 2006', Annual Report to the FOMC, app.ny.frb.org/markets/omo/omo2006.pdf.

Federal Reserve System. 2002. 'Alternative instruments for open market operations and discount window operations', Federal Reserve Study Group on Alternative Instruments for System Operations, Board of Governors of the Federal Reserve System, www.federalreserve.gov/BoardDocs/Surveys/soma/alt_instrmnts.pdf

Feibel, B. J. 2003. *Investment performance measurement.* Hoboken: Wiley.

Fender, I. and Hördahl, P. 2007. 'Overview: credit retrenchment triggers liquidity squeeze', *BIS Quarterly Review*, 09/2007: 1–16.

Financial Markets Association. 2007. *The ACI model code – The international code of conduct and practice for the financial markets.* Committee for Professionalism, cfmx2003.w3line.fr/aciforex/docs/misc/2007may15.pdf.

Fischer, B., Köhler, P. and Seitz, F. 2004. 'The demand for euro area currencies: past, present and future', ECB Working Paper Series 330.

FitchRatings. 2006. 'Fitch Ratings global corporate finance 1990–2005 transition and default study', FitchRatings Credit Market Research, www.fitchratings.com.

Flannery, M. 1996. 'Financial crisis, payment system problems, and discount window lending', *Journal of money credit and banking* 28: 804–24.

Fong, G., Pearson, C. and Vasicek, O. A. 1983. 'Bond performance: Analyzing sources of return', *Journal of Portfolio Management* 9: 46–50.

Freixas, X, Giannini, C., Hoggarth, G. and Soussa, F. 1999. 'Lender of Last Resort: a review of the literature', *Bank of England Financial Stability Review* 7: 151–67.

Freixas, X. 1999. 'Optimal bail out policy, conditionality and constructive ambiguity', Universitat Pompeu Fabra, Economics and Business Working Paper, www.econ.upf.edu/docs/papers/downloads/400.pdf.

Freixas, X. and Rochet, J.-C. 1997. *Microeconomics of banking*. Cambridge (MA): The MIT Press.

Freixas, X., Parigi, B. M. and Rochet, J.-C. 2003. 'The lender of last resort: a 21st century approach', ECB Working Paper Series 298.

Frongello, A. 2002a. 'Linking single period attribution results', *Journal of Performance Measurement* 6(3): 10–22.

2002b. 'Attribution linking: Proofed and clarified', *Journal of Performance Measurement* 7(1): 54–67.

Frye, J. 2000. 'Collateral damage detected', Federal Reserve Bank of Chicago, Emerging Issues Series Working Paper 10/2000 1–14.

Glasserman, P. 2004. *Monte Carlo methods in financial engineering*. New York: Springer-Verlag.

Glasserman, P., Heidelberger, P. and Shahabuddin, P. 1999. 'Asymptotically optimal importance sampling and stratification for pricing path-dependent options', *Mathematical Finance* 9(2): 117–52.

Glosten, L.R. and Milgrom, P. R. 1985. 'Bid, ask and transaction prices in a specialist market with heterogeneously informed traders', *Journal of Financial Economics* 14: 71–100.

Goodfriend, M. and Lacker. J. F. 1999. 'Limited commitment and central bank lending', *Federal Reserve Bank of Richmond Quarterly Review* 85(4): 1–27.

Goodhart, C. A. E. 1999. 'Myths about the lender of last resort', *International Finance* 2: 339–60.

2000. 'Can central banking survive the IT revolution?', *International Finance* 3(66): 189–209.

Goodhart, C. A. E. and Illing, G. 2002. *Financial crises, contagion and the lender of last resort: A Reader*. Oxford: Oxford University Press.

Goodwin, T. H. 1998. 'The information ratio', *Financial Analysts Journal* 54(4): 34–43.

Gordy, M. B. 2003. 'A risk-factor model foundation for ratings-based bank capital rules', *Journal of Financial Intermediation* 12: 199–232.

Gordy, M. B. and Lütkebohmert, E. 2007. 'Granularity adjustment for Basel II', Deutsche Bundesbank, Discussion Paper Series 2: Banking and Financial Studies 01/2007.

Gould, T. and Jiltsov, A. 2004. 'The case for foreign exchange exposure in U.S. fixed income portfolios', Lehman Brothers, www.lehman.com.

Grava, R. L. 2004. 'Corporate bonds in central bank reserves portfolios: a strategic asset allocation perspective', in C. Bernadell, P. Cardon, J. Coche, F. X. Diebold and S. Manganelli (eds.), *Risk Management for Central Bank Foreign Reserves*. Frankfurt am Main: European Central Bank, 167–79.

Grégoire, P. 2006. 'Risk attribution', *Journal of Performance Measurement* 11(1): 67–77.

Grinold, R. C. and Kahn, R. N. 2000. *Active portfolio management*. New York: McGraw-Hill.

Grossman, S. and Stiglitz J. E. 1980. 'On the impossibility of informationally efficient markets', *American Economic Review* 70: 393–408.

Guitierrez, M.-J. and Vazquez, J. 2004. 'Explosive hyperinflation, inflation-tax Laffer curve, and modeling the use of money', *Journal of Institutional and Theoretical Economics* 160: 311–26.

Gupton, G. M., Finger, C. C. and Bhatia, M. 1997. 'CreditMetrics™ – Technical Document', JPMorgan, www.riskmetrics.com.

Hamilton, J. D. 1994. *Time series analysis*. Princeton: Princeton University Press.

Hamilton, J. D. and Varma, P. 2006. 'Default and recovery rates of corporate bond issuers, 1920–2005', Moody's Report 96546, www.moodys.com

Hanson, S., Pesaran, M. H. and Schuermann, T. 2005. 'Firm heterogeneity and credit risk diversification', CESifo Working Paper Series 1531.

Hawtrey, R. 1932. *The art of central banking*. London: Longmans.

Henrard, M. 2000. 'Comparison of cashflow maps for value-at-risk', *Journal of Risk* 3(1): 57–71.

Hirshleifer, J. 1971. 'The private and social value of information and the reward to inventive activity', *American Economic Review* 61: 561–74.

Ho, T. S. Y. 1992. 'Key rate durations: Measures of interest rate risks', *Journal of Fixed Income* 2: 29–44.

Ho, T. S. Y. and Lee, S.-B. 1986. 'Term structure movements and pricing interest rate contingent claims', *Journal of Finance* 41(5): 1011–29.

Ho, T. S. Y., Chen, M. Z. H. and Eng, F. H. T. 1996. 'VAR analytics: Portfolio structure, key rate convexities, and VAR betas', *Journal of Portfolio Management* 23: 90–8.

Holton, G. A. 2003. *Value-at-risk: Theory and practice*. Boston: Academic Press.

Hong Kong Monetary Authority. 1999. 'Policy statement on the role of the Hong Kong Monetary Authority as lender of last resort', *Quarterly Bulletin* 8/1999: 77–81.

Huang, C. and Litzenberger, R. 1988. *Foundations for financial economics*. New Jersey: Prentice-Hall.

Hull, J. and White, A. 1990. 'Pricing interest rate derivative securities', *Review of Financial Studies* 3: 573–92.

1993. 'One factor interest rate models and the valuation of interest rate derivative securities', *Journal of Financial and Quantitative Analysis* 28(2): 235–54.

Humphrey, T. 1986. 'The classical concept of lender of last resort', in T. Humphrey and V. Richmond (eds.), *Essays on inflation*. 5th ed. Richmond: Federal Reserve Bank of Richmond.

IMF. *See* International Monetary Fund.

Information Security Forum. 2000. 'Fundamental information risk management', Information Security Forum 03/2000, www.securityforum.org/assests/pdf/firm.pdf.

Ingersoll, J.E. (1987). *Theory of financial decision making*. Savage: Rowman & Littlefield.

Institute of Internal Auditors. 2004. 'The role of internal audit in enterprise-wide risk management', Institute of Internal Auditors 09/2004, www.theiia.org/download.cfm?file=283.

Institute of International Finance. 2007 'Principles of liquidity risk management', Report 03/2007, www.afgap.org

International Monetary Fund. 2004. *Guidelines for foreign exchange reserve management*. Washington (DC): International Monetary Fund.

International Monetary Fund. 2005. *Guidelines for foreign exchange reserve management: Accompanying document and case studies*. Washington (DC): International Monetary Fund.

International Organization for Standardization. 2002. *Risk management – Vocabulary – Guidelines for use in standards*. Geneva: International Organization for Standardization, International Electrotechnical Commission.

International Organization for Standardization. 2005. *Information technology – Security techniques – Information security management systems – Requirements*. Geneva: International Organization for Standardization, International Electrotechnical Commission.

International Organization of Supreme Audit Institutions. 2004. *Guidelines for internal control standards for the public sector*. Brussels: International Organization of Supreme Audit Institutions, intosai.connexcc-hosting.net/blueline/upload/1guicspubsece.pdf.

International Swaps and Derivatives Association. 2006. 'Guidelines for collateral practitioners', www.isda.org.

Ippolito, R. A. 1989. 'Efficiency with costly information: a study of mutual fund performance, 1965–1984', *Quarterly Journal of Economics* 104: 1–23.

ISDA. *See* International Swap and Derivatives Association.

Israel, R.B., Rosenthal, J.S., Wei, J.Z. 2001. 'Finding generators for Markov chains via empirical transition matrices, with applications to credit ratings', *Mathematical Finance*, 11(2): 245–65.

Jäckel, P. 2002. *Monte Carlo methods in finance*. New York: Wiley.

Jarrow, R. A. and Subramanian, A. 1997. 'Mopping up liquidity', *Risk* 10: 170–3.

Jensen, M. C. 1968. 'The performance of mutual funds in the period 1945 – 1964', *Journal of Finance* 23(2): 389–419.

Jensen, M. C. and Meckling, W. H. 1976. 'Theory of the firm: managerial behavior, agency costs and ownership structure', *Journal of Financial Economics* 3: 305–60.

Jensen, M. C. and Murphy, K. J. 1990. 'Performance pay and top-management incentives', *The Journal of Political Economy* 98: 225–64.

Johnson-Calari, J. and Rietveld, M. 2007. *Sovereign wealth management*. London: Central Banking Publications.

Jorion, P. (ed.) 2003. *Financial risk manager handbook*. 2nd edn. Hoboken: Wiley.

2006. *Value-at-Risk: The new benchmark for managing financial risk*. 2nd edn. New York: McGraw-Hill.

Kahn, R. N. 1998. 'Bond managers need to take more risk', *Journal of Portfolio Management* 24(3): 70–6.

Kalkbrener, M. and Willing, J. 2004. 'Risk management of non-maturing liabilities', *Journal of Banking & Finance* 28: 1547–68.

Kang, J. C. and Chen, A. H. 2002. 'Evidence on theta and convexity in Treasury returns', *Journal of Fixed Income* 12: 41–50.

Karnosky, D. S. and Singer, B. D. 1994. *Global asset management and performance attribution*. Charlottesville: The Research Foundation of the Institute of Chartered Financial Analysts.

Katz, M. L. and Shapiro, C. 1985. 'Network externalities, competition, and compatibility', *American Economic Review* 75: 424–40.

Kilian, L. and Manganelli, S. 2003. 'The central bank as a risk manager, qualifying and forecasting inflation risks', ECB Working Paper Series 226.

Kim, C-J and Nelson, C. R. 1999. *State Space Models with Regime Switching*. Cambridge (MA): The MIT Press.

King, W. T. C. 1936. *History of the London discount market*. London: Frank Cass.

Kirievsky, L. and Kirievsky, A. 2000. 'Attribution analysis: Combining attribution effects over time', *Journal of Performance Measurement* 4(4): 49–59.

Koivu, M., Nyholm, K. and Stromberg, J. 2007. 'The yield curve and macro fundamentals in forecasting exchange rates', *The Journal of Financial Forecasting* 1(2): 63–83.

Kophamel, A. 2003. 'Risk-adjusted performance attribution – A new paradigm for performance analysis', *Journal of Performance Measurement* 7(4): 51–62.

Kreinin, A. and Sidelnikova, M. 2001. 'Regularization algorithms for transition matrices', *Algo Research Quarterly* 4(1/2): 23–40.

Krishnamurthi, C. 2004. 'Fixed income risk attribution', *RiskMetrics Journal* 5(1): 5–19.

Krokhmal, P., Palmquist, J. and Uryasev, S. 2002. 'Portfolio optimization with conditional Value-at-Risk objective and constraints', *The Journal of Risk* 4(2): 11–27.

Kyle, A. S. 1985. 'Continuous auctions and insider trading', *Econometrica* 53: 1315–35.

L'Ecuyer, P. 2004. 'Quasi-Monte Carlo methods in finance', in R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters (eds.), *Proceedings of the 2004 Winter Simulation Conference*. Piscataway: IEEE Press, pp. 1645–55.

L'Ecuyer, P., and Lemieux, C. 2002. 'Recent advances in randomised quasi-Monte Carlo methods', in M. Dror, P. L'Ecuyer, and F. Szidarovszki (eds.), *Modeling uncertainty: An examination of stochastic theory, methods, and applications*. Boston: Kluwer Academic Publishers, pp. 419–74.

Laker, D. 2003. 'Karnosky Singer attribution: A worked example', Barra Inc. Working Paper, www.mscibarra.com/research/article.jsp?id=303.

2005. 'Multicurrency attribution: Not as easy as it looks!', *JASSA* 2: 2005.

Lando, D. 2004. *Credit risk modeling: Theory and applications*. Princeton: Princeton University Press.

Lando, D. and Skødeberg, T. M. 2002. 'Analysing rating transitions and rating drift with continuous observations', *Journal of Banking & Finance* 26: 481–523.

Laurens, B. 2005. 'Monetary policy Implementation at different stages of market development', IMF Occasional Papers 244.

Lehmann, B. and Modest, D. 1987. 'Mutual fund performance evaluation: A comparison of benchmarks and benchmark comparisons', *Journal of Finance* 42: 233–65.

Leibowitz M.L., Bader L.N. and Kogelman S. 1995. *Return targets and shortfall risks: studies in strategic asset allocation*. Chicago: Irwin Professional Publishing.

Leone, A. 1993. 'Institutional aspects of central bank losses', IMF Paper on Policy Analysis and Assessment 93/14.

Lintner, J. 1965. 'The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets', *Review of Economics and Statistics* 47: 13–37.

Linzert, T., Nautz. D. and Bindseil, U. 2007. 'Bidding behavior in the longer term refinancing operations of the European Central Bank: Evidence from a panel sample selection model', *Journal of Banking and Finance* 31: 1521–43.

Litterman, R. 2003. *Modern investment management: An equilibrium approach*. New York: Wiley.

Litterman, R. and Scheinkman, J. 1991. 'Common factors affecting bond returns', *Journal of Fixed Income* 1: 54–61.

Loeys, J. and Coughlan, G. 1999. 'How much credit?', JPMorgan, www.jpmorgan.com.

Löffler, G. 2005. 'Avoiding the rating bounce: Why rating agencies are slow to react to new information', *Journal of Economic Behavior & Organization* 56(3): 365–81.

Lopez, A. J. 2002. 'The empirical relationship between average asset correlation, firm probability of default and asset size', Federal Reserve Bank of San Francisco Working Paper Series 2002/05.

Lord, T. J. 1997. 'The attribution of portfolio and index returns in fixed income', *Journal of Performance Measurement* 2(1): 45–57.

Lucas, D. 2004. 'Default correlation: from definition to proposed solutions', UBS CDO Research, www.defaultrisk.com/pp_corr_65.htm.

Manning, M. J. and Willision, M. D. 2006. 'Modelling the cross-border use of collateral in payment and settlement systems', Bank of England Working Paper 286.

Markowitz, H. 1952. 'Portfolio selection', *Journal of Finance* 7(1): 77–91.

Markowitz, H. M. 1959. *Portfolio selection: efficient diversification of investment*, New York: Wiley.

Marshall, C. 2001. *Measuring and managing operational risks in financial institutions*. New York: Wiley.

Martellini, L., Priaulet, P. and Priaulet, S. 2004. *Fixed-income securities*. Chichester: Wiley.

Martínez-Resano, R. J. 2004. 'Central bank financial independence', Banco de España Occasional Papers 04/01.

Mausser, H. and Rosen, D. 2007. 'Economic credit capital allocation and risk contributions', in J. Birge and V. Linetsky (eds.), *Handbooks in operations research and management science: Financial engineering*. Amsterdam: Elsevier Science, 681–725.

Meese, R. A., and Rogoff, K. 1983. 'Empirical exchange rate models of the seventies: Do they fit out of sample?', *Journal of International Economics* 14: 3–24.

Menchero, J. G. 2000a. 'An optimized approach to linking attribution effects', *Journal of Performance Measurement* 5(1): 36–42.

2000b. 'A fully geometric approach to performance measurement', *Journal of Performance Measurement* 5(2): 22–30.

2004. 'Multiperiod arithmetic attribution', *Financial Analysts Journal* 60(4): 76–91.

Merton, R. C. 1973. 'The theory of rational option pricing', *Bell Journal of Economics and Management Science* 4: 141–83.

1974. 'On the pricing of corporate debt: the risk structure of interest rates', *Journal of Finance* 29(2): 449–70.

Meucci, A. 2005. *Risk and asset allocation*. Berlin, Heidelberg, New York: Springer-Verlag.

Michaud, R. 1989. 'The Markowitz optimization enigma: Is optimized optimal?', *Financial Analyst Journal* 45: 31–42.

1998. *Efficient asset management: A practical guide to stock portfolio optimization and asset selection*. Boston: Harvard Business School Press.

Mina, J. 2002. 'Risk attribution for asset manager', *RiskMetrics Journal* 3(2): 33–55.

Mina, J. and Xiao, Y. (2001), 'Return to RiskMetrics: The Evolution of a Standard'. RiskMetrics.

Mirabelli, A. 2000. 'The structure and visualization of performance attribution', *Journal of Performance Measurement* 5(2): 55–80.

Moody's 2004. 'Recent bank loan research: implications for Moody's bank loan rating practices', Moody's Investors Service Global Credit Research report 12/2004, www.moodys.com.

2003. 'Measuring the performance of corporate bond ratings', Moody's Special Comment 04/2003, www.moodys.com.

Moskowitz, B. and Caflisch, R. E. 1996. 'Smoothness and dimension reduction in quasi-Monte Carlo methods', *Journal of Mathematical and Computer Modeling* 23: 37–54.

Mossin, J. 1966. 'Equilibrium in a capital asset market', *Econometrica* 34(4): 768–83.

Murira, B. and Sierra, H. 2006. 'Fixed income attribution, a unified framework – part I', *Journal of Performance Measurement* 11(1): 23–35.

Myerson, R. 1991. *Game Theory*. Cambridge (MA): Harvard University Press.

Myerson, R. and Satterthwaite, M. A. 1983. 'Efficient mechanisms for bilateral trading', *Journal of Economic Theory* 29: 265–81.

Nelson, C. R. and Siegel, A. F. 1987. 'A parsimonious modeling of yield curves', *Journal of Business* 60: 473–89.

Nesterov, Y. 2004. *Introductory lectures on convex optimization: A basic course*. Boston: Kluwer Academic Publishers.

Nickell, P., Perraudin, W. and Varotto, S. 2000. 'Stability of rating transitions', *Journal of Banking & Finance* 24: 203–27.

Niederreiter, H. 1992. *Random number generation and quasi-Monte Carlo methods*. Philadelphia: Society for Industrial and Applied Mathematics.

Nugée, J. 2000. *Foreign exchange reserves management*. Handbooks in Central Banking vol. 19. London: Bank of England Centre for Central Banking Studies.

Obeid, A. 2004. *Performance-Analyse von Spezialfonds – Externe und interne Performance-Maße in der praktischen Anwendung*. Bad Soden/Taunus: Uhlenbruch.

OECD. *See* Organization for Economic Co-operation and Development.

Office of Government Commerce. 1999. *Procurement excellence – A guide to using the EFQM excellence model in procurement*. London: Office of Government Commerce, www.ogc.gov.uk/documents/Procurement_Excellence_Guide.pdf.

Organisation for Economic Co-operation and Development. 2004. 'Principles of corporate governance, revised', Organisation for Economic Co-operation and Development 04/2004.

Pflug, G. 2000. 'Some remarks on the value-at-risk and the conditional value-at-risk', in S. Uryasev (ed.), *Probabilistic constrained optimization: Methodology and applications*. Dordrecht: Kluwer Academic Publishers, 272–81.

Pluto, K. and Tasche, D. 2006. 'Estimating probabilities of default for low default portfolios', in B. Engelmann and R. Rauhmeier (eds.), *The Basel II risk parameters*. Berlin: Springer-Verlag, 79–103.

Poole, W. 1968. 'Commercial bank reserve management in a stochastic model: Implications for monetary policy', *Journal of Finance* 23: 769–91.

Pringle, R. and Carver, N. (eds.) 2003. *How countries manage reserve assets*. London: Central Banking Publications.

2005. 'Trends in reserve management – Survey results', in R. Pringle, and N. Carver (eds.), *RBS Reserve management trends 2005*. London: Central Banking Publications, 1–27.

2007. *RBS reserve management trends 2007*. London: Central Banking Publications.

Project Management Institute. 2004. *A guide to the project management body of knowledge (PMBOK® Guide)*. Pennsylvania: PMI Inc.

Putnam, B. H. 2004. 'Thoughts on investment guidelines for institutions with special liquidity and capital preservation requirements' in C. Bernadell, P. Cardon, J. Coche, F. X. Diebold

and S. Manganelli (eds.), *Risk Management for central bank foreign reserves*. Frankfurt am Main: European Central Bank, chapter 2.

Ramaswamy, S. 2001. 'Fixed income portfolio management: Risk modeling, portfolio construction and performance attribution', *Journal of Performance Measurement* 5(4): 58–70.

2004a. *Managing credit risk in corporate bond portfolios: a practitioner's guide*. Hoboken: Wiley.

2004b. 'Setting counterparty credit limits for the reserves portfolio' in C. Bernadell, P. Cardon, J. Coche, F. X. Diebold and S. Manganelli (eds.), *Risk management for central bank foreign reserves*. Frankfurt am Main: European Central Bank, chapter 10.

2005. 'Simulated credit loss distribution: Can we rely on it?', *The Journal of Portfolio Management* 31(4): 91–9.

Reichsbank. 1910. *The Reichsbank 1876–1900*. Translation edited by the National Monetary Commission. Washington: Government Printing Office.

Reitano, R. R. 1991. 'Multivariate duration analysis', *Transactions of the Society of Actuaries* 43: 335–92.

Repullo, R. 2000. 'Who should act as lender of last resort: an incomplete contracts model', *Journal of Money, Credit and Banking* 32(3): 580–605.

RiskMetrics Group 2006. 'The RiskMetrics 2006 methodology', www.riskmetrics.com.

Roberts, J. 2004. *The modern firm: organizational design for performance and growth*. Oxford and New York: Oxford University Press.

Rockafellar, R. T. and Uryasev, S. 2000. 'Optimization of conditional value-at-risk', *The Journal of Risk* 2(3): 21–41.

2002. 'Conditional Value-at-Risk for general loss distributions', *Journal of Banking & Finance* 26: 1443–71.

Rodrik, D. 2006. 'The social costs of foreign exchange reserves', NBER Working Paper 11952.

Rogers, C. 2004. 'Risk management practices at the ECB', in C. Bernadell, P. Cardon, J. Coche, F. X. Diebold and S. Manganelli (eds.), *Risk management for central bank foreign reserves*. Frankfurt am Main: European Central Bank, chapter 15.

Ross, S. 1976. 'The arbitrage theory of capital asset pricing', *Journal of Economic Theory* 13: 341–60.

Saarenheimo, T. 2005. 'Ageing, interest rates, and financial flows', Bank of Finland Research Discussion Paper 2/2005.

Samad-Khan, A. 2005. 'Why COSO is flawed', *Operational Risk* 01/2005: 24–8.

2006a. 'Fundamental issues in OpRisk management', *OpRisk & compliance* 02/2006: 27–9.

2006b. 'Uses and misuses of loss data', *Global association of risk professionals [risk review]* 05–06/2006: 18–22.

Sangmanee, A. and Raengkhum, J. 2000. 'A general concept of central bank wide risk management', in S. F. Frowen, R. Pringle and B. Weller (eds.), *Risk management for central bankers*. London: Central Banking Publications.

Satchell, S. 2007. *Forecasting expected returns in the financial markets*. Oxford: Academic Press Elsevier.

Saunders, A. and Allen, L. 2002. *Credit risk measurement: new approaches to value at risk and other paradigms*. 2nd edn. New York: Wiley.

Sayers, R. S. 1976. *The Bank of England, 1891–1944*. 2 vols. Cambridge: Cambridge University Press.

Scherer, B. 2002. *Portfolio construction and risk budgeting*. London: Risk Books.
Scobie, H. M. and Cagliesi, G. 2000. *Reserve management*. London: Risk Books.
Sentana, E. 2003. 'Mean-variance portfolio allocation with a Value at Risk constraint', *Revista de Economía Financiera* 1: 4–14.
Sharpe, W. 1991. 'The arithmetics of active management', *Financial Analysts Journal* 47: 7–9.
Sharpe, W. F. 1964. 'Capital asset prices: A theory of market equilibrium under conditions of risk', *Journal of Finance* 19(3): 425–42.
1966. 'Mutual fund performance', *Journal of Business* 39(1): 119–38.
1994. 'The Sharpe ratio', *Journal of Portfolio Management* 21(1): 49–58.
Shefrin, H. 2007. *Behavioral corporate finance: decisions that create value*. Boston: McGraw Hill/Irvin.
Smith, C. and Stulz, R. M. 1985. 'The determinants of a firm's hedging policies', *Journal of Financial and Quantitative Analysis* 20: 391–406.
Sobol, I. M. 1967. 'The distribution of points in a cube and the approximate evaluation of integrals', *U.S.S.R. Journal of Computational Mathematics and Mathematical Physics* 7: 86–112.
Spaulding, D. 1997. *Measuring investment performance*. New York: McGraw-Hill.
2003. *Investment performance attribution*. New York: McGraw-Hill.
Standard & Poor's. 2006. 'Annual 2005 global corporate default study and rating transitions', www.ratingsdirect.com.
Standard & Poor's. 2008a. '2007 Annual global corporate default study and rating transitions',www.standardandpoors.com/ratingsdirect.
2008b. 'Sovereign defaults and rating transition data: 2007 update', www.standardandpoors.com/ratingsdirect.
Standards Australia. 2004. *Risk management*. AS/NZS 4360. East Perth: Standards Australia.
Stella, P. 1997. 'Do central banks need capital?', IMF Working Paper 83.
2002. 'Central bank financial strength, transparency, and policy credibility', IMF Working Paper 137.
2003. 'Why central banks need financial strength', *Central Banking* 14(2): 23–9.
Stulz, R. M. 2003. *Risk management and derivatives*. Cincinnati: South-Western.
Summers, L. H. 2007. 'Opportunities in an era of large and growing official wealth', in Johnson-Calari and Rietveld, pp. 15–28.
Svensson, L. E. 1994. 'Estimating and interpreting forward interest rates: Sweden 1992–1994', IWF Working Paper 114.
Sveriges Riksbank. 2003. 'The Riksbank's role as lender of last resort', *Financial Stability Report* 2/2003: 57–73.
Tabakis, E. and Vinci, A. 2002. 'Analysing and combining multiple credit assessments of financial institutions', ECB Working Paper 123.
Task Force of the Market Operations Committee of the European System of Central Banks. 2007. 'The use of portfolio credit risk models in central banks', ECB Occasional Paper Series 64.
Thornton, H. 1802. *An inquiry into the nature and effects of paper credit of Great Britain*. New York: Kelley.
Treynor, J. L. 1962. 'Toward a theory of market value of risky assets', unpublished manuscript. A final version was published in 1999, in Robert A. Korajczyk (ed.) *Asset pricing*

and portfolio performance: Models, strategy and performance metrics. London: Risk Books, 15–22.

1965. 'How to rate management of investment funds', *Harvard Business Review* 43: 63–75.

1987. 'The economics of the dealer function', *Financial Analysts Journal* 43(6): 27–34.

Treynor, J. L. and Black, F. 1973. 'How to use security analysis to improve portfolio selection', *Journal of Business* 1: 61–86.

U.S. Department of Homeland Security. 2003. *Reference manual to mitigate potential terrorist attacks against buildings*. Washington (DC): Federal Emergency Management Association, www.fema.gov/pdf/plan/prevent/rms/426/fema426.pdf.

Van Breukelen, G. 2000. 'Fixed income attribution', *Journal of Performance Measurement* 4(4): 61–8.

Varma, P., Cantor, R. and Hamilton, D. 2003. 'Recovery rates on defaulted corporate bonds and preferred stocks, 1982–2003', Moody's Investors Service, www.moodys.com

Vasicek, O. A. 1977. 'An equilibrium characterization of the term structure', *Journal of Financial Economics* 5: 177–88.

1991. 'Limiting loan loss probability distribution', KMV Corporation, www.kmv.com.

Wilkens, M., Baule, R. and Entrop, O., 2001. 'Basel II – Berücksichtigung von Diversifikationseffekten im Kreditportfolio durch das granularity adjustment', *Zeitschrift für das gesamte Kreditwesen* 12/2001: 20–6.

Williamson, O. E. 1985. *The economic institutions of capitalism*. New York: The Free Press.

Willner, R. 1996. 'A new tool for portfolio managers: Level, slope and curvature durations', *Journal of Fixed Income* 6: 48–59.

Wittrock, C. 2000. *Messung und Analyse der Performance von Wertpapierportfolios* 3rd edn. Bad Soden/Taunus: Uhlenbruch.

Wong, C. 2003. 'Attribution – arithmetic or geometric? The best of both worlds', *Journal of Performance Measurement* 8(2): 10–8.

Woodford, M. 2001. 'Monetary policy in the information economy', NBER Working Paper Series 8674.

2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton: Princeton University Press.

Wooldridge, P. D. 2006. 'The changing composition of official reserves', *BIS Quarterly Review* 09/2006: 25–38.

Index

accountability
  active portfolio management, and, 27
  collateral frameworks, and, 344
  lending without collateral, and, 272
  public institutional investors, and, 5
  risk management functions, and, 451
  strategic asset allocation, in, 72
active portfolio management
  academic studies of, 23
  additional costs, 23
  competitive equilibrium, as part of, 24
  discovering arbitrages, 25
  diversifiable risk, cost of, 23
  diversification, and, 25
  industrial organisation, 23
  make or buy decision, 26
  mixed portfolios, 26
  outsourcing, 27
  portfolio management firms, nature of, 26
  portfolio management industry, 25
  public institutions, by, 23
  qualities of managers, 24
  types of funds, 25
  usefulness of, 23
Advanced Measurement Approach (AMA), 446
alpha strategies, 219
application service providers (ASPs), 206
Arbitrage Pricing Theory (APT), 71, 224
asset-backed securities (ABS), 349
asset-liability management (ALM) framework, 57–8

backdated transactions, 186
Banco de España
  credit assessment system, 310
bank loans
  collateral, as, 352
Bank of Japan
  collateral framework, 341
  credit assessment system, 311
banknotes
  central bank profitability, and, 30
  central banks' capital, and, 36
  denominations in circulation, 31
  liquidation risk, and, 31
  potential future decline of, 30
  risks relating to, 31
  seignorage, and, 31
  withdrawal of right to issue, implications, 39
Banque de France
  credit assessment system, 310
Basel Accord
  bank capital, and, 34
Basel Committee for Banking Supervision (BCBS), 446
benchmark portfolios, 158
Business Process Management (BPM)
  operational risk management compared, 481

capital asset pricing model (CAPM), 18
  multi-factor return decomposition models, and, 224
  risk-adjusted performance measures, and, 213
  strategic asset allocation, and, 49
Capital Market Line (CML), 59
central banks
  active investors, as, 26
  agent of the people, as, 6
  ALM approaches, 58
  capital, role of, 34–41
  collateral frameworks, *see* collateral frameworks
  conservative nature of, 118
  credit assessment systems, 309
  credit losses, 30
  credit risk and, 118
  currency appreciation, and, 14
  currency risk, and, 119
  derivatives, use of, 22
  development of market intelligence, and, 9
  diversification of portfolios, 10, 21, 118
  excess reserves, and, 14
  exposure to credit risk, 117

central banks (*cont.*)
financial crisis management, 34, *see* Financial crisis management
financial institution, as, 444
firm, as, 445
foreign exchange policies and reserves, 13, 14, 30, 33
FX markets, and, 90
FX valuation changes, 30
implicit capital, 8
independence of, 8
inflationary policies, and, 41
insider information, and, 9
investment horizon, 70
investment universe of, 117
lending without collateral, 272
operational risk management, *see* Operational risk management
policy related risk factors, 29–34
policy tasks, 10–17
price stability, and, 35
profitability problems, 38
real bills doctrine, 12
reserves, growth of, 119
risk types, 15
segregation of domestic assets, 12
sterilising excess liquidity, 30
supply of deposits, 12
threats to profitability, 29
transparency, and, 8
withdrawal of right to issue currency, 39
Chinese walls, 9
meaning, 450
role of risk management, and, 450
collateral
assessment of compliance with eligibility criteria, 352
asset types, 273
assessment of credit quality, 275, 305, 307–15
availability, 415
available amounts, 276
cash collateral, 279
central bank operations, for, 295
collateral eligibility, 279
collateralization, best practice, 280
cost-benefit analysis of, 284–300
counterparties, choice of, 345
credit risk assessment, 353–7
creditworthiness, 343
cross-border use of, 431
cut-off line, 274
distortions to asset prices, 344
easy pricing, 276
eligibility criteria, 348
Eurosystem approach, 277
frameworks compared, 342–8
haircut determination methods, 318
haircuts, 280
handling costs, 274, 276
inter-bank repurchase markets, for, 295
inter-bank transactions, 279
legal certainty, 275
limits, 337
liquidity, 276
marking to market, 315
mitigation of credit risk, and, 304
monitoring, 274
monitoring use of, 282
ranking of collateral types, 274
transparency and accountability, 344
type and quantity of, 343
valuation, 315
Committee of Sponsoring Organizations of the Treadway Commission, 448
compliance monitoring and reporting, 157
limits, and, 179
concentration, 142, 151
concentration risks 368–74
conditional forecasting
strategic asset allocation, and, 75
corner solutions, 68
corporate governance
structure, 447
counterparty borrowing limits
collateral, and, 356
counterparty Risk Management Policy Group II, 447
credibility
reputation risk, and, 7
credit derivatives
growth of market, 119
credit instruments
rate of return on, 119
credit rating agencies
criticisms of, 124
ECB use of, 174
credit rating tools (RTs)
collateral credit quality, and, 313
credit risk
assessment, 353–7
credit risk, definition, 117
currency risk, and, 119
data limitations, 122
default, and, 120
diversification of risk, 121
liquidity and security, and, 120
market risk models compared, 122
meaning, 303

measuring risk, 119
mitigating, 303
mitigation of downside risk, 120
nature of risk, 117
pay off of credit instruments, 120
resource implications of investment, 120
return distribution of credit instruments, 122
return on investment grade credit, 119
credit risk modelling, 117–74
ECB's approach, 122–42
asset correlation, 141
equity returns, 141
simulating, 360–6
simulation results, 142–52
validation of models, 145
sources of risk, 117
credit spreads
"credit spread puzzle", 120
determinants of, 120
ECB credit risk model, and, 142
idiosyncratic risk, and, 121
limits of diversification, and, 121
currency risk, 119

data transfer infrastructure, 197
debt securities
foreign official institutions holding, 118
decision support framework
strategic asset allocation, and, 72
deflation
interest rates, and, 22
derivatives
central banks, and, 22
Deutsche Bundesbank
credit assessment system, 310
diversification
active portfolio management, and, 25
central banks, by, 10
corner solutions, and, 68
credit risk, 19, 121
limitations of, 121
optimal degree of, 17

emergency liquidity assistance (ELA), 394
enterprise risk system, 198
equity
central banks, and, 22
euro corporate bond market
financials, and, 121
European Bank for Reconstruction and Development (EBRD), 117
European Central Bank (ECB)
approach to performance attribution, 257–67
benchmarks structure, 160
collateral framework, 341
credit limits, 173
credit risk modelling, approach to, 122–42
credit risk, and, 118
distribution of reports, 191
foreign currency assets, 165
foreign reserves, 160
investment framework and benchmarks, 160
investment portfolios, 159
investment process, 161
investment process, components of, 68–74
key risk indicators (KRIs), use of, 485
market risk control framework, 165
operational risk management framework, 462
performance measurement, 219–21
portfolio management, 159–61
tasks of, 6, 11
reporting for investment operations, 193
strategic asset allocation, and, 52
European Investment Bank (EIB), 117
Eurosystem, 159
governance of, 159
Eurosystem Credit Assessment Framework (ECAF), 308, 313
performance monitoring framework, 315
excess reserves
definition, 14

Federal Reserve
Qualified Loan Review (QLR) program, 312
Federal Reserve Board
collateral framework, 341
financial crisis management, 394–438
aggregate excess liquidity injection, 397
availability of collateral, 416
central bank borrowing facility, 417
central bank operational framework, role, 416
central bank's ability to secure claims, 405
constructive ambiguity, 411
cross-border use of collateral, 431
ELA provided by other banks, 407
emergency liquidity injections, 422
emergency solvency assistance, 398
end-of-day borrowing facility, 417
equal access measures, 396, 420–32
individual access measures, 398
individual banks measures, 432–5
inertia principle, 418
intra-day payment system, 417
key lessons from 19th century, 399
moral hazard, 406
motivations for, 403
narrowing spread of borrowing facility, 427
negative externalities of illiquidity, 403

financial crisis management (*cont.*)
reserve requirements, 418
risk taking, and, 34
special lending, rate of provision, 414
special liquidity supplying operations, 397
spread of borrowing facility, 397
superior knowledge of central bank, 405
swap operations, 431
typology of measures, 396–9
widening of collateral set, 397, 428
fixed set up costs
diversification, and, 19
fixed-income performance attribution models, 249
Foreign Exchange Counterparty Database (FXCD), 201
foreign exchange rate policies
implementation, 13
foreign exchange rates
central banks' policy, and, 33
modelling, 86
risk integration, 57
costs of holding, 33
foreign reserves
ALM, and, 58
best practice for, 33, 448
currency composition of, 22
growth in, 13
social welfare, and, 6
sovereign bonds, investment in, 21
US policy on, 14

global investment performance standards (GIPS), 208
gold reserves
foreign reserves, and, 22
governance structures
asset allocation, and, 69
government bonds
fixed income securities, and, 18
liquidity of, 282
government paper
excessive purchases of, 9
government securities
collateral frameworks, and, 344
public institutional investors, and, 5

haircuts, 306–33
average issue size, 329
basic VaR related haircuts, 321
bid-ask spread, 330
credit risk adjusted haircuts, 333
defining liquidity categories, 331
determination methods, 318
effective supply, 329
liquidity risk adjusted haircuts, 323–33
headline risk, 7
hedge funds
diversification, and, 25

independence
central banks, of, 7
index selection, 58
inertia principle
financial crisis management, 417
inflation
banknotes in circulation, and, 33
benchmark rates of, 33
interest rates, and, 32
insider information
central banks, and, 9, 450
insolvency
role of capital, and, 34
integrated risk management, 41–7
best practice, 41
business activities, 43
business model, 42
complete list of risk factors, 46
consistent risk measures, 47
distorted allocations of risk budget, 47
efficient frontier, 42
franchise capital, 44
parameters of risk control framework, 47
policy handbook, 46
profit-loss asymmetries, 45
public investors, for, 43–7
reputation risks, 45
risk budgets, 42
risk factors, 42
risk return preferences, 44
scenario analysis, 47
segregation of risk management, 46
social welfare, and, 46
sources of risk aversion, 44
taxation, 42
interest rates
central banks' inside information, and, 9
deflation, and, 32
real, 33
setting, 12, 32
internal ratings based (IRB) system, 311, 446
International Monetary Fund (IMF)
strategic asset allocation, and, 52
international Operating Working Group (IORWG), 460
investment horizon
strategic asset allocation, and, 70
ISDA's *Guidelines for Collateral Practitioners*
collateralization, and, 280

IT
applications, 200
architecture and standards, 197
build or buy, 204–6
data transfer infrastructure, 197
development support, 200
enterprise risk system, 198
integrated risk management system, 197
outsourcing, 206
application service provider solutions, 206
projects, 203
reporting infrastructure, 198
risk data warehouse, 198
risk management IT team, 199
risk management, and, 196
systems support and operations, 199

Key Risk Indicators (KRIs)
operational risk management, and, 485

lender of last resort (LOLR), 394
limits
risk mitigation tool, as, 337
liquidation risk
banknotes, and, 31
liquidity risk, 20
banknotes, and, 31
meaning, 176
liquidity-related risks
simulating, 366

maintenance costs
diversification, and, 19
market intelligence
central banks, and, 9, 27
market risk, 118
composition of, 162
definition, 162
ECB control framework, *see* European Central Bank
measurement, 164
marking to market, 306
collateral valuation, and, 315
Markowitz portfolio optimization model, 68
Matlab, 202
mean-variance portfolio theory, 58
modern Portfolio Theory, 58
monetary policy
implementation, 12, 272
operations, 35
inter, 32
Monte Carlo method
credit risk estimation, 379
empirical results on variance reduction, 384
importance sampling, 380
Quasi-Monte Carlo methods, 382
multi-factor return decomposition models, 224–8
Arbitrage Pricing Theory, 224
choice of risk factors, 226
parameterizing, 226
multi-factor return decomposition models
empirical multi-factor models, 227

non-alienable risks, 20

Oesterreichische Nationalbank
credit assessment system, 311
open market operations (OMOs)
emergency liquidity injections through, 422
operational risk management, 460
active benchmarking initiatives, 462
bottom up self-assessments, 479
central bank specific challenges, 463–5
ECB framework, 462
ECB governance model, 479
governance of, 483
inherent risk vs worst case scenario, 467
International Operating Working Group (IORWG), 460
KRIs and, 484
lifecycle of, 471
likelihood/impact matrix, 467
normal business conditions vs worst case scenarios, 466
operational risk, definition, 465–8
overarching framework, as, 468
reporting, 484
risk as distribution, 466
risk impact grading scale, 473
risk tolerance guidelines, 474
taxonomy of operational risk, 469, 470
tolerance policy, 472
top down self-assessments, 476–9
optimization, 58
outsourcing
active portfolio management, 27

passive portfolio management
definition, 18
payment systems
unrenumerated liabilities, and, 11
performance attribution, 222–68
active investment decision process, and, 242
Arbitrage Pricing Theory, 224
performance attribution modelling, 223
ECB approach to, 257–67
fixed income portfolios, 228–41

performance attribution (*cont.*)
fixed-income performance attribution analysis, 223
fixed-income performance attribution models, 241–57
multi-factor return decomposition models, *see* Multi-factor return decomposition models
prime objective, 222
range of performance determinants, 242
return-driving risk factors, 223
single period/multiple periods, 243
tailored reports, 242
performance measurement, 207
active performance, 217
active positions, and, 208
benchmark portfolios, 207
Capital Asset Pricing Model, 213
ECB, at, 219–21
extension to value-at-risk, 216
GIPS requirements, 220
information ratio, 217, 220
literature on, 208
passive performance, 215
performance analysis, meaning, 207
reward-to-VaR ratio, 216, 220
risk-adjusted performance measures, 213–19
rules for return calculation, 208–13
Sharpe ratio, 214, 220
total performance, 214
Treynor ratio, 215, 220
private information
diversification, and, 19
public institutional investors
'big' investors, as, 10
accountability and transparency, 5
active investors, as, 26
active portfolio management, 6, 9, 21, 23
credibility, and, 7
diversification of assets, 19–24
foreknowledge, 6
governance standards, and, 7
Government securities, and, 5
headline risk, 7
importance of, 3
independence, and, 7
industry failures, and, 17
investors' preferences, and, 5
large implicit economic capital, and, 8
market intelligence, 9, 27
non-alienable risks, 20
normative theory of investment behaviour, 3
organisational flexibility, 4
outsourcing active management, 27
passive portfolio management, 18
payments to owners, and, 6
portfolio managers, 5
private sector techniques, and, 7
remoteness of activities, 8
reputation risk, 20
risk aversion, and, 17
share of equity, 18
social welfare, and, 6, 27
transparency and accountability, and, 27

Quasi-Monte Carlo methods, 382

rating methodologies, 168
ratings
limitations of, 124
rating aggregation, 169
real bills doctrine, 12
repo portfolios
concentration risks, 368
Credit Value-at-Risk, 376
expected shortfall, 376
liquidity-related risks, simulating, 366
Monte Carlo method, 379
residual risk estimation, 387
risk measurement for, 359–93
simulating credit risk, 360
reporting
accuracy, 190
availability of necessary data, 191
delivery of, 191
ECB investment operations, for, 193
framework for, 190
IT infrastructure, 198
level of detail, 190
objectivity and fairness, 190
operational risk management, for, 484
portfolio managers, 189
risk and performance, on, 189
timeliness, 190
repurchase transactions
cash leg, 303
collateral leg, 304
credit risk, 304
market and liquidity risk, 304
reputation risk
credit exposures, and, 45
definition, 7
integrated risk management, and, 45
non-alienable risk factor, as, 20
public institutions, and, 7
quantifying, 45
transparency, and, 8
residential mortgages
funding of in Europe, 348

residual risk estimation
  credit quality of issuers and counterparties, 390
  Eurosystem credit operations, for, 387–92
  expected shortfall in base case scenario, 388
  liquidity time assumptions, 389
return calculation
  rules for, 208–13
risk control framework, 157–206, 306, 317, 331, 339, 353, 357
  aim of, 157
  coherence of, 162
  components of, 157
  credit risk limits, 166
  defining limits, 161
  ECB's credit limits, 173
  enforcement of, 157
  exposure calculation, 172
  factors driving relevant industries, 169
  inputs to limit setting formulas, 167
  internal credit rating system, 169
  limit compliance monitoring, 179
  limits, 161–78
  liquidity limits, 176
  market measures of credit risk, monitoring, 169
  market risk limits, 162
  rating methodologies, 168
  risk and performance reporting, 189
    strategic benchmarks, maintenance of, 188
    validation of prices transacted at, 182
    valuation at end of day prices, 181
risk data-warehouse, 198
Risk Engine, 201
risk management, 271
  accountability, 451
  adequacy of resources, 454
  Basel Committee for Banking Supervision, 446
  best practices, 445–8
  central banks', 3
  Chinese wall principle, 450
  corporate governance structure, 447
  divisional responsibilities, 455
  independence, and, 448
  interest rate risk, 446
  internal control systems, 447
  middle office functions, 455
  monetary policy operations, 271, 455
  operational risks, 457
  organizational issues, 443, 444
  relevance of in central banks, 444
  risk management culture, 457
  strategic level of decisions, 455
  supervisory process, 446
  transparency, 451
risk mitigation techniques, 277
  haircuts, 278
  limits, 278
  valuation and margin calls, 277
risk-adjusted performance measures, 213–19
RiskMetrics, 166
RiskMetrics RiskManager, 202

semi-passive portfolio management, 208
Sharpe ratio, 214
short selling, 64
spreadsheets, 202
static data
  maintenance of, 187
strategic asset allocation, 49–116
  active portfolio management, 55
  aims, 49
  application of, 99–116
  benchmark allocation, 75
  benchmark process, 52
  calculation of returns, 87
  Capital Asset Pricing Model (CAPM), and, 49
  Capital Market Line (CML), 59
  conditional forecasting, 75
  corner solutions, 68
  credit risk modelling, and, 123
  decision support framework, 58
  degrees of complexity, 54
  delegation of responsibilities, 53
  discretization, 95
  ECB investment process, 68–74
  foreign reserve holdings, 91
  IMF guidelines, 51
  index selection, 58
  integration of different risks, 57
  internalization of process, 54
  investment horizon, 54
  investment universe, 99
  investors' expectations, 63
  length of investment horizons, 63
  level of liquidity, 53
  macro economic scenarios, 104
  macro model for multiple currency areas, 77
  macro-economic variables, 75
  mean-variance portfolio theory, 58
  model for credit migrations, 83
  modelling exchange rates, 86
  multi-currency model, 93
  non-normal scenario, to, 111
  normative considerations, 51
  objective function and constraints, 100
  optimal portfolio allocations, 109
  optimization models for shortfall approach, 89–98

strategic asset allocation (*cont.*)
portfolio optimization, 51, 58–68
return distributions, 64
short selling, 64
single market model, 97
starting yield curves, 104
stochastic factors, modelling, 75–89
strategic benchmark, 52
utility functions, 63
viewbuilding, 56
yield curve model, 80
yield curve projections and expected returns, 105
sub-prime turmoil 2007, 297

taxation
integrated risk management, and, 42
time-weighted rate of return (TWRR), 209
tolerance bands
ECB calculation of, 184
Total Quality Management (TQM)
operational risk management compared, 481
transaction costs
diversification, and, 19
transparency
central banks, and, 9
public institutional investors, and, 5
reputation risk, and, 7
Treynor ratio, 215

unsecured bonds
collateral, as, 349

Wallstreet Suite, 179, 201