

WILEY FINANCE

Synthetic and structured assets

ERIK BANKS

Synthetic and Structured Assets

A Practical Guide to Investment and Risk

Erik Banks



John Wiley & Sons, Ltd

Synthetic and Structured Assets

For other titles in the Wiley Finance Series
please see www.wiley.com/finance

Synthetic and Structured Assets

A Practical Guide to Investment and Risk

Erik Banks



John Wiley & Sons, Ltd

Copyright © 2006 John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester,
West Sussex PO19 8SQ, England

Telephone (+44) 1243 779777

Email (for orders and customer service enquiries): cs-books@wiley.co.uk

Visit our Home Page on www.wiley.com

All Rights Reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London W1T 4LP, UK, without the permission in writing of the Publisher.

Requests to the Publisher should be addressed to the Permissions Department, John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex PO19 8SQ, England, or emailed to permreq@wiley.co.uk, or faxed to (+44) 1243 770620.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The Publisher is not associated with any product or vendor mentioned in this book.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the Publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Other Wiley Editorial Offices

John Wiley & Sons Inc., 111 River Street, Hoboken, NJ 07030, USA

Jossey-Bass, 989 Market Street, San Francisco, CA 94103-1741, USA

Wiley-VCH Verlag GmbH, Boschstr. 12, D-69469 Weinheim, Germany

John Wiley & Sons Australia Ltd, 42 McDougall Street, Milton, Queensland 4064, Australia

John Wiley & Sons (Asia) Pte Ltd, 2 Clementi Loop #02-01, Jin Xing Distripark, Singapore 129809

John Wiley & Sons Canada Ltd, 22 Worcester Road, Etobicoke, Ontario, Canada M9W 1L1

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Library of Congress Cataloging-in-Publication Data

Banks, Erik.

Synthetic and structured assets / Erik Banks.

p. cm. — (Wiley finance series)

Includes bibliographical references and index.

ISBN-13: 978-0-470-01713-5 (cloth : alk. paper)

ISBN-10: 0-470-01713-9 (cloth : alk. paper)

1. Securities. 2. Structured notes (Securities) 3. Derivative securities. I. Title. II. Series.

HG4521.B3463 2006

332.63'2—dc22

2005034995

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

ISBN 13 978-0-470-01713-5 (HB)

ISBN 10 0-470-01713-9 (HB)

Typeset in 10/12pt Times by TechBooks, New Delhi, India

Printed and bound in Great Britain by Antony Rowe Ltd, Chippenham, Wiltshire

This book is printed on acid-free paper responsibly manufactured from sustainable forestry in which at least two trees are planted for each one used for paper production.

Contents

Acknowledgements	ix
About the author	xi
1 Introduction to Synthetic and Structured Assets	1
1.1 Development of structured and synthetic assets	2
1.2 Drivers of market activity	3
1.3 New product design	5
1.4 Overview of the text	7
2 Financial Building Blocks	9
2.1 Introduction	9
2.2 Concepts	9
2.2.1 Risks	9
2.2.2 Time value of money and interest rates	11
2.3 Derivative instruments	16
2.3.1 Derivatives	16
2.4 Host securities/liabilities	38
2.4.1 Public notes/shelf registrations	38
2.4.2 Private placements	38
2.4.3 Transferable loans	39
2.5 Issuing/repackaging vehicles	39
2.5.1 Special purpose entities	40
2.5.2 Trusts	41
2.5.3 Investment companies and partnerships	42
2.6 Financial engineering and product design	43
3 Callable, Puttable, and Stripped Securities	45
3.1 Introduction	45
3.2 Development and market drivers	45
3.3 Product mechanics and applications	47
3.3.1 Callable bonds	47
3.3.2 Puttable bonds	52
3.3.3 Stripped securities	53

4	Mortgage- and Asset-backed Securities	59
4.1	Introduction	59
4.2	Development and market drivers	59
4.3	Product mechanics and applications	62
4.3.1	Mortgage-backed securities	62
4.3.2	Asset-backed securities	90
5	Structured Notes and Loans	99
5.1	Introduction	99
5.2	Development and market drivers	99
5.3	Product mechanics and applications	102
5.3.1	General structural issues	102
5.3.2	Interest-rate-linked notes	104
5.3.3	Currency-linked notes	107
5.3.4	Commodity-linked notes and loans	109
5.3.5	Equity-linked notes	110
5.3.6	Credit-linked notes	113
6	Collateralized Debt Obligations	121
6.1	Introduction	121
6.2	Development and market drivers	121
6.3	Product mechanics and applications	123
6.3.1	General structural issues	124
6.3.2	Cash flow CDOs and market value CDOs	133
6.3.3	Balance sheet and arbitrage CDOs	134
6.3.4	Structured and synthetic CDOs	135
7	Insurance-linked Securities and Contingent Capital	141
7.1	Introduction	141
7.2	Development and market drivers	141
7.3	Product mechanics and applications	143
7.3.1	Insurance-linked securities	144
7.3.2	Contingent capital structures	153
8	Convertible Bonds and Equity Hybrids	163
8.1	Introduction	163
8.2	Development and market drivers	163
8.3	Product mechanics and applications	165
8.3.1	Convertible bonds and variations	166
8.3.2	Bonds with warrants	172
8.3.3	Buy/write (covered call) securities	173
8.3.4	Other equity hybrids	177
9	Investment Funds	181
9.1	Introduction	181
9.2	Development and market drivers	181

9.3	Product mechanics and applications	184
9.3.1	Structure, diversification, and management	184
9.3.2	Open-end funds	186
9.3.3	Closed-end funds	193
9.3.4	Hedge funds	195
9.3.5	Exchange-traded funds	200
10	Derivative Replication, Repackaging, and Structuring	205
10.1	Introduction	205
10.2	Development and market drivers	205
10.3	Product mechanics and applications	207
10.3.1	Synthetic long and short option and swap positions	207
10.3.2	Multiple swap/option positions	214
10.3.3	Asset swaps, liability swaps, and callable/puttable asset swap packages	231
10.3.4	Credit derivatives and synthetic credit positions	238
11	Risk, Legal, and Regulatory Issues	247
11.1	Introduction	247
11.2	Risk and financial controls	247
11.2.1	Market, liquidity and credit risk management	247
11.2.2	Internal financial and operational controls	248
11.2.3	Internal audit	249
11.3	Legal controls	249
11.3.1	ISDA documentation framework	250
11.4	Regulatory and accounting issues	255
11.4.1	Regulatory capital	255
11.4.2	Regulatory review	257
11.4.3	Accounting treatment under FAS 133 and IAS 39	257
	Selected References	261
	Index	263

Acknowledgements

I would like to thank Samantha Whittaker and all at John Wiley & Sons for their invaluable assistance on the project from start to finish. The Wiley team performed to its usual high standards in all respects. Thanks are also due to a number of colleagues and referees from Merrill Lynch, Morgan Stanley, UBS, and Citicorp, who provided useful comments and feedback on key sections.

Most special thanks are due to Milena . . . again!

EB

About the Author

Erik Banks is an independent risk consultant and financial author who has been active in the investment banking sector for 20 years. Erik has held senior risk management positions at Merrill Lynch, XL Capital, and Citibank in New York, Tokyo, London, and Hong Kong, and has written 20 books on derivatives, risk, emerging markets, and merchant banking, including the John Wiley titles *The Simple Rules of Risk*, *Exchange-Traded Derivatives*, *Alternative Risk Transfer*, and *Catastrophic Risk*. He lives in Maine with his wife, Milena, and their horses and dogs.

Introduction to Synthetic and Structured Assets

Financial activities have been a part of the global economic framework for centuries. Lending, borrowing, speculating, investing, and hedging, for instance, have been employed for years by a broad range of institutions in order to achieve specific financial goals. Not surprisingly, as forces of deregulation, technology, and capital mobility have taken firmer root in the landscape of the late 20th and early 21st centuries, the financial marketplace has evolved, becoming increasingly useful, efficient, and sophisticated. It is now common for institutions, which once relied on basic capital-raising and investment instruments, to turn to a range of highly customized, though eminently practical, assets and liabilities in order to achieve desired goals. Our aim in this text is to examine many of these customized instruments, demonstrate how they have developed and evolved, and consider how they function mechanically, and in practice.

Our target sector can be classified in a number of ways. For our purposes, we consider the markets and products in two broad forms: structured assets and synthetic assets.

- *Structured assets.* We define the class of structured assets to include instruments that are created, decomposed, or restructured in some fashion in order to redirect or alter underlying cash flows. This may be accomplished by altering the properties of physical assets, such as bonds or equities, through the use of special purpose entities/trusts and/or through the inclusion of one or more derivative contracts, which are off-balance sheet contracts that derive their value from some underlying reference.
- *Synthetic assets.* We define the class of synthetic assets to include instruments that are created exclusively out of one or more derivatives. The package of contracts generates cash flows that correspond with specific end-user requirements.

There are instances when both classes of assets can be used to achieve the same end results. Consider, for instance, that a pool of secondary mortgages can be combined through a trust or special purpose entity (SPE) to create a mortgage-backed security (i.e. a structured asset), while a mortgage swap or total return swap can be created to mimic the flows of the same pool of mortgage-backed securities (i.e. a synthetic asset). In some instances, it may be advantageous to create the asset in structured form, while in other situations it may be beneficial to do so synthetically. These broad classifications serve us well in arranging the discussion and analysis which follows, by allowing us to consider separately those instruments that can be decomposed and restructured through a redirection of cash flows, and those that can be created or replicated using off balance-sheet contracts.

We consider in this introductory chapter key items related to the historical development of synthetic and structured assets, and key drivers that have fuelled market expansion over the past years. We then consider general issues related to new product design, and the essential characteristics that are required for success. In order to frame the material properly, we also provide an outline of the structure of chapters that follow. These topics provide an appropriate macro context for the more detailed market and product material that follows in the remainder of the book.

1.1 DEVELOPMENT OF STRUCTURED AND SYNTHETIC ASSETS

Although the broad class of synthetic and structured assets has gained increasing attention and use since the 1990s, aspects of the market date back many decades. Indeed, some of the most elemental and popular instruments of the financial markets are structured assets dating back to the 19th century. Consider, for instance, that the convertible bond, which is a package of a fixed-income security and an investor equity option to convert the bond into the stock of the issuer, was first launched in the mid-1800s. The first commodity-linked bond, with redemption tied to the price of cotton (i.e. a bond and an investor cotton option), dates back to the same era. Mutual funds, which essentially are single shares of stock representing an interest in a broader portfolio of assets, were developed in the late 1880s and early 1900s and popularized in the 1960s, and have evolved and expanded since that time. Even callable bonds (e.g. a bond and an issuer call option) and puttable bonds (e.g. a bond and an investor put option), have been in existence for several decades, and are now mainstays of the marketplace.

Though the synthetic and structured market traces its roots back a minimum of several decades, it is clear that the greatest amount of financial innovation and growth has occurred in more recent times. Key factors such as derivative valuation methods, technology, legal structuring, market liquidity, cross-border capital flows, and financial creativity have led to the development of increasingly customized and sophisticated assets. By reacting to forces that simply did not exist during earlier times, intermediaries have been able to expand their ability to meet the needs of end-users, including issuers and investors.

For instance, prior to the advent of option valuation models (beginning with seminal work by Black, Scholes, and Merton in the early 1970s), there was little in the way of comprehensive options dealing; since options are a core constituent of many of the instruments we consider below, many new products simply could not be structured.

Similarly, the introduction of more powerful (and inexpensive) computing capacity, starting in the 1990s and accelerating into the new millennium, has led to the creation of increasingly complex products that require intensive simulation-based pricing routines. Networking and communications have also promoted the concept of electronic trading platforms and electronic communication networks, both of which promise to continue the trend towards online OTC product trading – including trading of structured and synthetic assets.

Clarification of the legal environment has also proven significant in the development of the sector. Creation and use of standardized legal documentation (e.g. bank loan agreements, trading confirmations, International Swaps and Derivatives Association (ISDA) Master Agreements) has made it easier for parties to a transaction to agree binding terms and conditions, and to settle disputes or disagreements efficiently. Legal agreement on the right to net credit exposures in the event of bankruptcy has made it possible for participants to manage their risks (and risk capital) more accurately. Legal development of vehicles such as SPEs and trusts, often in tax-friendly and legally secure jurisdictions, has likewise prompted new product development. Though legal uniformity does not exist in all countries or regions, it is certainly prevalent in most of the world's primary financial dealing jurisdictions (e.g. North America, Europe, Japan).

Market liquidity has also been an important factor in the development of various instruments. As many financial assets and contracts have developed a core base of interest among an increasingly broad group of users, liquidity has grown in tandem. This has been especially important for basic financial assets that are used to construct structured and synthetic contracts.

We shall note in Chapter 2 that the fundamental “building blocks” of the sector are all liquid instruments that benefit from active two-way flows. Without such market liquidity, it would be difficult to create, in an economically rational way, the products we discuss in this text. It is worth noting, of course, that the synthetic and structured assets that result from the financial engineering process do not feature the same degree of market liquidity as the underlying assets. Most are far less liquid than the building blocks used in the construction process – a fact that is hardly surprising, since many of the resulting instruments are intended to meet specific end-user needs.¹

The financial creativity of intermediaries has been a catalyst in structured and synthetic asset development. Intermediaries tend to respond to the requests and demands of end-use clients (i.e. the market is demand-, rather than supply-, driven). However, the ability of intermediaries to apply techniques of financial engineering to create entirely new contracts has helped the market develop successfully. Leading intermediaries can use their knowledge of markets, client requirements, and valuation techniques to develop useful, customized assets that meet specific needs. Intermediaries that can couple financial creativity with a significant amount of risk-taking are well positioned to win client business.

1.2 DRIVERS OF MARKET ACTIVITY

Financial instruments develop and evolve in the marketplace in order to serve a specific function. If that function is performed successfully, the instrument gains a following and succeeds; evolutionary iterations may then follow, permitting further expansion. If the function is not performed successfully, the instrument will eventually fade from use. The specific synthetic and structured assets we discuss in this book include those that have proven successful over a period of time; the instruments have achieved a critical mass of interest by addressing the needs of participants properly. In each of our individual product chapters we shall consider a series of market drivers that have fuelled market development and growth. All, however, trace their foundation to a core series of goals that intermediaries and end-users attempt to meet. In this section we consider, in generic form, some of these elemental market drivers.

Institutions are active in the capital markets in order to achieve one or more core goals related to some aspect of financial management; the synthetic and structured assets we consider in the book can help achieve any, or all, of these goals.

Broadly speaking, core financial management goals include:

- *Funding.* An institution that needs to finance its operations in the external capital markets (rather than via internally generated funds) attempts to do so in an optimal fashion. This generally means arranging the lowest cost of funding while maintaining a balanced portfolio of liabilities across markets and maturities. Synthetic or structured liabilities are used routinely to both lower funding costs and provide new or incremental investor/market access.
- *Hedging.* An institution with exposures that can impact inputs or outputs typically tries to protect against potential downside risks in order to minimize the chance of losses. This is often accomplished through a formal or informal hedging program that makes use of appropriate hedges. Once again, synthetic and structured contracts can be used to create the best possible hedge for an exposure.

¹ Even within the overall sector, we can observe differences in market liquidity: some assets, such as senior-rated tranches of collateralized debt obligations or stripped US Treasuries, feature a reasonable degree of market liquidity, while others, such as privately placed credit-linked notes with embedded exotic options, feature much less liquidity (and effectively must be considered “hold until maturity” contracts).

- *Investing/yield enhancing.* An institution (or department within an institution) that exists solely or primarily to invest cash or capital on behalf of internal operations or external parties again attempts to achieve its investment goals in a rational and cost-effective manner by optimizing its risk/return profile. Specific synthetic or structured assets are often an effective mechanism for increasing returns while preserving a desired risk profile.
- *Speculating.* An institution that is responsible for generating asset returns by taking a greater amount of risk will again seek to achieve its goals by implementing its speculation program in a manner that is structured appropriately with regard to concentration, volatility, leverage, and liquidity. Again, many of the synthetic and structured assets we consider in the text permit establishment of maximum speculative positions, including those that are heavily leveraged and/or exposed to complex and volatile risks.

Naturally, these core goals exist because of financial market volatility, a characteristic of the modern financial markets that generates both risk and opportunity. As long as market volatility remains a feature of the landscape, and there is little to suggest that it will disappear or even decline, then these goals should remain intact, helping to fuel further innovation and activity.

Each of these goals, which together comprise the essence of corporate financial activity, can be met through the use of conventional financial assets and contracts. For instance, a company seeking funding may choose to access the Eurobond or the syndicated loan markets. One that is attempting to hedge an interest rate exposure may opt for the listed bond futures markets. Those trying to invest, yield enhance, or speculate can select from a range of cash or derivative instruments. It is also true, however, that each of these goals can often be met more effectively by using synthetic and structured instruments. Thus, the company seeking funding may find it more cost effective to issue a floating-rate note with an attached swap that converts its interest expense back into a fixed cost, or sell a fixed bond with embedded options to lower its all-in funding cost. Similarly, the investor seeking to speculate on a particular index may choose to introduce a leveraged payout in order to increase its risk/return profile.

So, synthetic and structured assets can help achieve core goals in a better way. The act of repackaging, restructuring, or synthetically replicating asset or liability profiles can lead to the same funding, hedging, investing, or speculating profiles – at a lower cost or for a higher return, and almost certainly in a more efficient manner. This brings us back to our earlier statement – financial instruments survive and thrive when they are useful. The assets we discuss in this book have become established in the marketplace precisely because they are useful in helping institutions achieve their fundamental corporate goals.

Certain other forces supplement the items we have noted above, and serve as additional drivers:

- *Regulation and market access.* An institution may want to participate in a specific market – from an asset or liability perspective – but may be unable to do so as a result of regulatory restrictions or barriers to entry. When this occurs, synthetic/structured contracts can often open up the marketplace to relatively free participation.
- *Asset creation.* An institution may seek a very specific asset or liability profile in order to fulfill risk, funding, or investment mandates. If this is not available in the conventional financial sector, structured and synthetic instruments can surface as potential alternatives by allowing the creation of instruments with relevant yield, maturity, currency, return, and/or risk characteristics.
- *Liquidity creation.* An institution may be impacted by lack of market liquidity in select assets or liabilities that form part of its activities; this may be a temporary or permanent

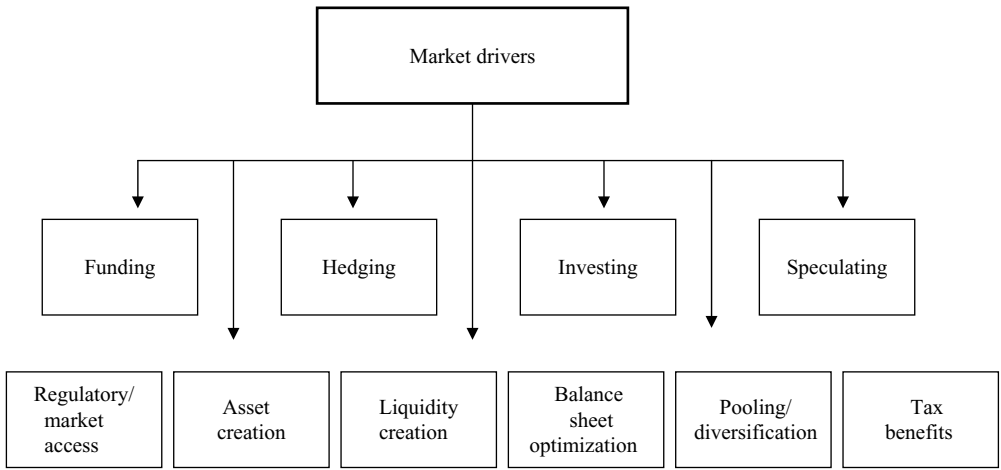


Figure 1.1 Key market drivers

condition that can prevent a firm from arranging transactions in the most economic manner possible (i.e. the illiquidity of market is reflected directly in the size of the bid-offer spread). Structured and synthetic instruments can be used to inject a level of liquidity into the market, returning an institution to a more cost-efficient position.

- *Balance sheet optimization.* An institution that actively manages its assets and funding and capital levels may find it beneficial to use assets or liabilities that help optimize its goals; in some cases, this may involve transferring exposures off balance sheet via synthetic and structured assets.
- *Pooling and diversification.* An institution that is attempting to create a balanced portfolio of risk exposures may find it can do so most effectively by using vehicles that can pool and diversify risks through a single transaction. A variety of synthetic and structured assets can help accomplish this goal.
- *Tax benefits.* An institution seeking to exploit tax differences between marketplaces legally in order to reduce friction costs can do so using certain classes of synthetic/structured contracts.

These primary and secondary drivers, summarized in Figure 1.1, have led to progressively greater expansion and innovation in the financial markets. Although we have framed our discussion in general terms in this section, we shall revisit the topic throughout the book in order to reinforce the point that new products are not created simply to demonstrate financial engineering skills or generate profits – they are developed by intermediaries in order to fulfill the specific needs of institutions in the best way possible. Profit streams can be sustained when products meet client demands.

1.3 NEW PRODUCT DESIGN

Financial intermediaries create new types of product in order to address client requirements. Some new structured/synthetic products are successful, and many others are not; some achieve widespread volume or carve out a niche, while others simply fade. When a product no longer attracts meaningful interest and ceases to be traded, it is abandoned. Most new products that

become accepted in the marketplace feature margin compression over time as more intermediaries and end-users join in the process. This profit compression, coupled with the desire to service client needs, leads intermediaries to create more new instruments.

A new product must feature certain basic characteristics in order to succeed. These characteristics may relate to the product itself, the underlying asset market, or the regulatory/tax environment – or all three. Market evidence suggests that a new product is more likely to be successful when participants (or potential participants) recognize the value it can provide and wish, therefore, to participate; that is, intermediaries do not need to “convince” participants of the benefits of entering the marketplace. Such demand-pull, rather than supply-push, creates a healthier and more sustainable equilibrium. In general, greater likelihood of success exists when:

- The underlying asset is homogenous, storable, and price-volatile, it is in abundant supply and features good price transparency (e.g. is not subject to manipulation).
- The asset and its price performance are transparent enough to attract the attention of investors, hedgers, and speculators (e.g. parties that can help promote liquidity).
- The asset market is developed to the point where there is reasonably strong two-way flow, ensuring a minimum base of liquidity; indeed, the asset should be linked to other cash or hedging markets in order to build on two-way flows.
- Regulatory and tax treatments are equitable, or those featuring differences/discrepancies can be arbitrated. In fact, regulatory issues have been, and are likely to remain, a significant influence in the design of new products and aftermarket activity. If the regulatory environment moves towards a “level playing field” across national boundaries, then the likelihood that more institutions can participate in a new asset when it is launched increases significantly.²
- There is accounting clarity regarding the instrument. The distinction between debt and equity creates significant tax implications, and must therefore be considered carefully. For instance, debt-related structures may feature interest tax deductibility, while equity-related instruments may face double taxation of dividends (e.g. taxation at the issuer and investor level, unless the issuer is a strict pass-through entity).
- Costs are reasonable. Expenses associated with accounting and regulatory requirements, along with stamp duties, clearing/settlement expenses, arranging costs, and other trading-related spreads, cannot overwhelm the economic rationale for structuring or executing a transaction.³ The low-cost providers may emerge as leaders in a highly competitive financial market.

New product development is, unfortunately, an expensive and time-consuming process. For instance, structures that are meant to be listed and traded very widely (rather than as private placements, for instance) must be vetted rigorously from a legal and regulatory perspective; this is especially true if the product is intended for purchase by retail investors. In some instances, this can take several years to accomplish. Intermediaries supporting this type of

² For instance, the original PRIMES and SCORES structures we consider in Chapter 8 were “derailed” by unfavorable rulings from the Internal Revenue Service (IRS). The IRS eventually changed its position and allowed the securities to proceed through a “grandfathered” grantor trust scheme, but no new trusts were ever formed and the product was eventually wound down, to be replaced by Morgan Stanley’s new synthetic version (without the attendant regulatory complications). Similarly, Lehman’s attempt to create unbundled stock units comprised of a coupon bond, growth/income certificate (dividend security), and equity appreciation certificate (call warrant) met with fierce regulatory and tax resistance, and was abandoned before launch. Many other examples exist.

³ For example, Deutsche Bank’s country basket of stocks failed to displace Morgan Stanley’s world equity baskets because the bank attempted to reduce tracking error by adding too many small stocks, which added considerably to the costs of trading, custody, and settlement.

development must be prepared to invest human and capital resources in order to reap benefits; in this sense, they operate just as any other corporation might: allocating capital and human resources to the creation of a profitable venture that may take months or years to develop and market. But their role does not cease with the introduction of a new product. In many instances, financial intermediaries must continue to support the asset by providing ongoing liquidity (market-making) or by assuming a certain amount of credit, market, or liquidity risk. Thus, even after a product has been launched successfully, its ongoing viability may depend in large part on continued participation by the community of intermediaries. This means, of course, that banks, securities firms, and other product creators must be compensated for risks taken in supporting the product. Unfortunately, profit margins on new products can compress quickly as a result of competitive pressures, suggesting that a misbalancing of risk/return may arise.

Participation in new products generally proceeds through evolutionary stages. Activity might begin on a very modest scale, with a few intermediaries and end-users arranging transactions. After some level of experience is gained, changes in the core structure might be implemented in order to resolve problems, reduce costs, or improve efficiencies. Enhancements may then lead to greater product marketing by intermediaries and a gradual accumulation of critical mass. If products are truly customizable, end-users may then begin demanding greater flexibility in risk/return profiles to meet their needs more accurately. The end result for the successful new product is a strong base of demand, leading to improved liquidity and tighter pricing – all while addressing specific client needs. Much of what we have noted above applies to the product development process in any current or immediate period. However, truly innovative intermediaries plan ahead, attempting to estimate or predict future client requirements in a new market environment. By doing so, intermediaries can anticipate, rather than react to, client demand – gaining valuable time over competitors.

1.4 OVERVIEW OF THE TEXT

There are various ways in which a text on financial contracts can be structured: by marketplace, function, product, risk characteristics, geography, and so forth. Each approach has its own merits and, in some instances, shortcomings. We have opted, in this text, to follow a product focus, which allows us to inject uniformity into a discussion that spans multiple asset and liability classes. Accordingly, for each of the product chapters (i.e. Chapters 3 through 10), we include a discussion of market development, growth, and drivers, along with product mechanics and practical applications. Again, though various approaches can be used to separate the broad class of synthetic and structured products, we have selected the following categorization:

- Chapter 3: callable, puttable, and stripped securities, including corporate and government bonds with options or stripped coupons.
- Chapter 4: mortgage- and asset-backed securities, including pass-through securities, mortgage bonds, collateralized mortgage obligations, and receivables/loan-backed securities.
- Chapter 5: structured notes and loans, including interest rate, currency rate, equity, commodity, and credit-linked notes, bonds, and loans.
- Chapter 6: collateralized debt obligations, including cash flow, market value, arbitrage, balance sheet, structured, and synthetic collateralized bond and loan obligations.
- Chapter 7: insurance-linked securities and contingent capital, including catastrophe bonds, noncatastrophe bonds, contingent debt, and contingent equity.

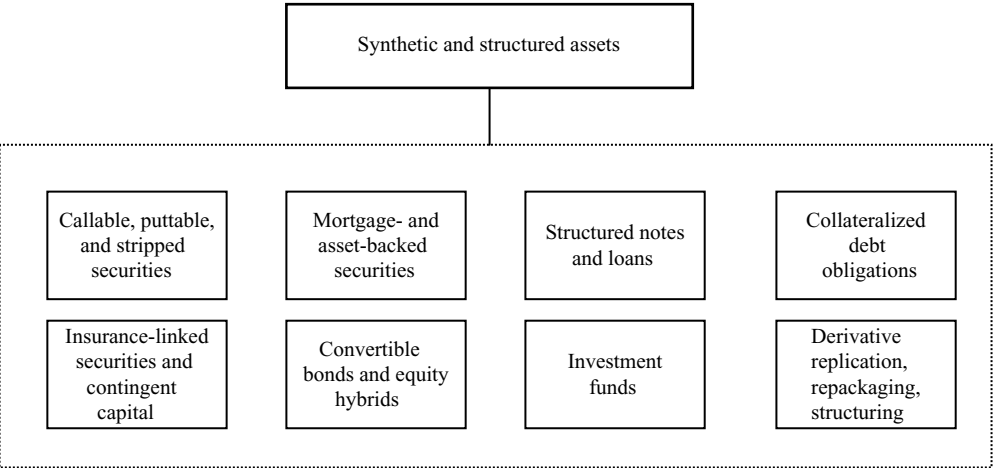


Figure 1.2 Scope of synthetic and structured asset coverage

- Chapter 8: convertible bonds and equity hybrids, including convertible bonds, mandatory convertibles, zero coupon convertibles, reverse convertibles, bonds with equity warrants, and synthetic buy/write packages.
- Chapter 9: investment funds, including open- and closed-end mutual funds, hedge funds, and exchange-traded funds.
- Chapter 10: derivative replication, repackaging, and structuring, including synthetic long and short option and swap positions, multiple option/swap positions, callable/puttable asset swap packages, and credit derivatives/synthetic credit positions.

Figure 1.2 summarizes the scope of our product coverage.
Two other chapters supplement the product-specific chapters:

- Chapter 2: Financial Building Blocks, which considers the essential concepts and tools needed to construct or decompose synthetic and structured assets, including derivatives, host securities, and issuance/repackaging vehicles.
- Chapter 11: Risk, Legal, and Regulatory Issues, which addresses the risk management, financial control, accounting, legal, and regulatory frameworks that surround these unique contracts.

With this background in hand, we are now prepared to begin our discussion of synthetic and structured assets by examining the basic tools that are used to create unique contracts.

Financial Building Blocks

2.1 INTRODUCTION

The creation of synthetic and structured assets depends critically on the existence of various financial and legal tools. In fact, the financial instruments and contracts we discuss in this book exist because a set of building blocks is available to structure and reshape cash flows and risk profiles to meet the needs of end-users and other participants. Before embarking on our analysis of synthetic and structured assets, it is helpful to review these building blocks. As we progress through each product chapter, we will demonstrate how these building blocks lie at the heart of structured and synthetic asset development.

Since actual, expected, and contingent cash flows underpin much of the business of financial engineering, we begin our discussion with a brief review of the concepts of risk, time value of money, and interest rates. We then introduce the essential instruments/vehicles of the synthetic and structured asset world: derivatives, host securities, and issuing/repackaging vehicles.¹ We shall, of course, revisit each of these in greater detail throughout the course of the book.

2.2 CONCEPTS

2.2.1 Risks

Risk, a cornerstone of finance, can be defined and classified in many different ways. In its most basic form, risk can be defined as the uncertainty regarding a future outcome. From a financial perspective, we can expand on this basic definition by noting that financial risk is the uncertainty associated with the future state of a financial market; this uncertainty may lead to a profit or a loss.² Financial risk can be decomposed into various other classes in order to develop meaningful concepts and metrics; common classes include market risk, credit risk, liquidity risk, and operating risk; each of these can, in turn, be subdivided further, as noted below.

The existence of financial risks suggests that parties may experience gains or losses as a result of their exposures. A proper risk management framework allows such gains/losses to be identified and weighed properly, so that optimal risk/return decisions can be made. In some instances this will mean assuming more risks, and in other cases it will mean eliminating or transforming as much risk as possible. There is, of course, no optimal ex ante profile that is applicable uniformly to every end-user or intermediary; each firm must consider risk issues in light of its own operations and goals.

¹ Since our intent in this chapter is to acquaint the reader with the basic tools of structured products, we have chosen not to embark on a financial, mathematical, or legal discourse; rather, we illustrate the fundamental elements of each topic. Readers interested in pursuing in-depth reviews of the mathematical or legal aspects of the topic are urged to consult the works listed in the reference section.

² Risk is often classified in either speculative or pure form, particularly as related to the insurance sector. A speculative risk is one that can yield a loss, a gain, or no loss/gain, while a pure risk can only yield a loss or no loss, but no possibility of a gain. From a financial markets perspective, we deal primarily with speculative risks; indeed, many of the products we consider in the book are created with risk/return profiles that can generate gains or losses for the parties involved. We shall consider the narrower version of pure risk in Chapter 7.

Market risk is the risk of loss arising from the adverse movement of markets or market references. Within the broad class of market risk we may consider various granular definitions, including:

- Directional risk: the risk of loss arising from an adverse movement in the direction of a market/reference, such as an equity, bond, index, currency, or commodity.
- Curve risk: the risk of loss arising from an adverse movement in the shape of a yield curve, such as a steepening, flattening, or inversion of the curve.
- Volatility risk: the risk of loss arising from an adverse upward/downward movement in the absolute or relative volatility of a market/reference.
- Basis risk: the risk of loss arising from an adverse movement in the differential between two asset references, or a futures contract and its underlying deliverable.
- Spread risk: the risk of loss arising from an adverse movement in the differential between a risk-free benchmark and its associated risky counterpart, such as a corporate bond quoted as a spread to a risk-free government bond.
- Correlation risk: the risk of loss arising from a change in the correlation between two or more assets that define a contract or exposure.

Credit risk is the risk of loss arising from the failure of a counterparty to perform on its contractual obligations. The class of credit risks can be subdivided into various components, including:

- Default risk: the risk of loss arising from the outright failure of a counterparty to perform on its liabilities and contractual obligations.³ Default risk exposures may be decomposed into trading risk exposures (from derivative/financing contracts), direct credit exposures (from unsecured loans), settlement risk exposures (from currency/securities settlements), and contingent risk exposures (from future commitments and contingencies that may create credit exposures).
- Sovereign risk: the risk of loss arising from the actions of a government authority on local/nonlocal assets in its financial system, including capital/exchange controls, currency devaluation, asset expropriation, or debt moratorium/repudiation.

*Liquidity risk*⁴ is the risk of loss arising from an inability to obtain funding or liquidate assets in order to cover cash flow needs. Liquidity risk can be segregated into:

- Asset liquidity risk: the risk of loss arising from an inability to sell or pledge assets at, or near, carrying value in order to generate cash to meet liabilities/payments coming due.
- Funding liquidity risk: the risk of loss arising from an inability to obtain rollover financing or incremental unsecured funding to meet other liabilities/payments coming due.
- Joint asset/funding liquidity risk: the risk of loss arising from a cash spiral, where an inability to obtain funding leads to forced asset sales/pledges at below-market prices, leading to further funding shortfalls, and so forth.

Operational risk is the risk of loss arising from the failure of operating procedures, technologies, or processes. Though the class of operating risks is very broad, some of the most

³ It is worth noting that an extreme stage of financial distress, short of default, can create losses in the traded liabilities of a firm; under our taxonomy of risks, this type of loss is captured as a market-based spread risk loss.

⁴ Liquidity risk is often included as a subset of market risk. However, for purposes of clarity, we prefer to consider the exposure on a standalone basis.

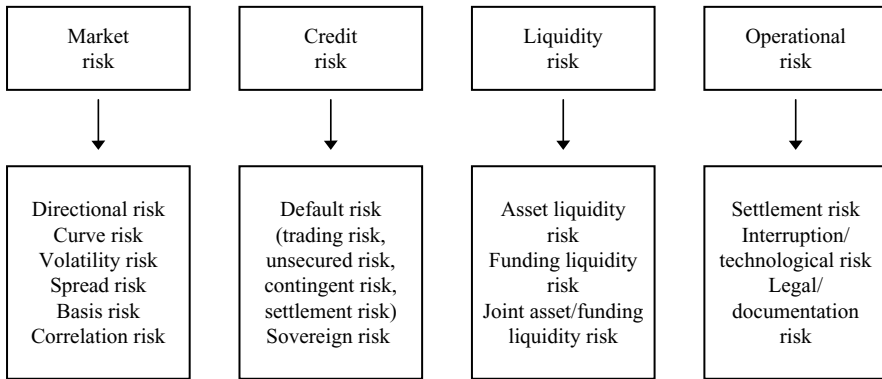


Figure 2.1 Risk classifications

common subclasses include:

- **Payment/settlement risk:** the risk of loss arising from the failure by a firm to deliver cash or securities in fulfillment of a trade/contract (for any reason excluding default).
- **Interruption/technology risk:** the risk of loss arising from an inability to conduct business in a standard operating environment as a result of infrastructure/technology malfunctions, interruptions, or destruction.
- **Legal/documentation risk:** the risk of loss arising from an improperly constructed legal position/defense or flawed legal documentation that does not protect or indemnify adequately.

Figure 2.1 summarizes these basic risk classes.

End-users and intermediaries must also contend with a range of nonfinancial operating risks. For instance, companies must properly manage the input and output risks that affect production and revenue processes, they must contend with liability risks, employment, and environmental issues, and so forth. While these are vitally important, they are beyond the scope of our discussion in this chapter.

The primary point to consider in this brief introduction to risk is that different forms of exposure exist, and can affect the financial outcome/performance of a transaction or line of business. Some firms attempt to eliminate or minimize certain types of risk in order to stabilize the cash flows affecting their businesses. Others, in contrast, prefer to assume or transform certain risk exposures in order to generate a profit and increase business cash flows. It is worth stressing that an exposure that can generate a loss for one party can create a gain for another party; this means speculative profit opportunities exist, allowing for the development of financial solutions that can benefit both parties. Indeed, we shall discover that the synthetic and structured assets discussed in this text play an important role in allowing parties to hedge, transfer, transform, or assume various types of financial risk.

2.2.2 Time value of money and interest rates

Time value of money is the essential concept of present and future cash flows, and is at the heart of all financial engineering. In basic terms, the time value of money indicates that in a marketplace featuring positive interest rates, a dollar invested today will be worth more in the

future; conversely, the value of a future cash flow will be lower in today's market than in the future. The actual value ascribed to present values (PV) and future values (FV) is based on the use of a discount rate. In industrialized "risk-free" markets, such as the US, UK, Germany, and Japan, this discount rate is the rate paid by a government entity with the highest possible credit ratings (e.g. no risk of default); we denote this r_f . In risky markets, an appropriate risk premium, r_p , must be factored in to take account of possible default risk; the overall risky discount rate is thus $r = r_f + r_p$. The larger the discount rate, the greater the discounted cash flows and the lower the resulting PVs. This makes intuitive sense, as the future cash flows of a risky credit are worth less today than the future cash flows of a government credit. Similarly, the larger the discount rate, the greater the FV for a given sum invested today.

The present value of a future cash flow is given as:

$$PV = P \left[\frac{1}{(1 + r)^t} \right]$$

where

P is a cash flow
 r is the discount rate
 t is the time horizon.

The present value of a continuously compounded cash flow can be computed via:

$$PV = Pe^{-rt}$$

Similarly, the future value of a cash flow invested today is given as:

$$FV = P(1 + r)^t$$

The future value of a continuously compounded cash flow is:

$$FV = Pe^{rt}$$

For instance, the PV of \$100 to be received in one year, with a continuously compound discount rate of 5 %, is \$95.13; likewise, the FV of \$100 invested today, at 5 % continuously compounded, is \$105.13.

Time value of money is essential for pricing of securities and contracts with future cash flows, such as notes, bonds, and the structured and synthetic assets we consider in the book. For instance, the current price of a fixed income security depends on both the cash flows that define the security and the discount rate used to PV the cash flows. The cash flow components are fixed for a standard coupon-bearing instrument, and include periodic coupons and redemption value (often at par).⁵ The discount rate, in contrast, changes constantly (e.g. every hour or day, depending on the specific market), meaning the price changes as well.

The price of a bond is inversely related to yield: as yields rise, bond prices fall, and as yields fall, bond prices rise. We can demonstrate this by examining the fundamental bond valuation equation, which is simply an expansion of the PV equation noted above:

$$P = \sum_{n=1}^t \frac{C}{(1 + r)^n} + \frac{M}{(1 + r)^t}$$

⁵ They are also fixed for discount instruments, such as Treasury bills, and include implicit (rather than coupon) interest and redemption value (generally par).

where

- P is now the present value (or price) of the bond
- t is the number of periods until maturity
- C is the periodic coupon
- M is the redemption value of the bond at maturity
- r is the current discount rate.

Let us consider a simple example that demonstrates the valuation process and the inverse price/yield relationship. If a bond has a par value of \$100 and pays an annual coupon of 5 % for two years, the price of the bond when the current market is also 5 % is \$100. If market rates rise tomorrow to 6 %, the price of the bond falls to \$98.17. This makes sense intuitively as the outstanding bond only has a coupon of 5 %; a new two-year bond issued tomorrow at par will have a coupon of 6 %, which is more attractive to investors, meaning the existing 5 % bond will have to trade at a discount in order to be more appealing. In fact, the selling pressure from investors reallocating their capital from the 5 % bond to the new 6 % bond will generate the discount. Conversely, if market rates fall to 4 %, the price of the bond rises to \$101.80. A new two-year par bond issued in a 4 % market environment will be less attractive than the outstanding 5 % bond, which will trade at a premium as investors allocate their capital to the higher coupon security.

An extension to the bond price valuation framework involves the price sensitivity of a bond for a change in interest rates. It is relatively easy to see from the equation above that the level of rates, r , will influence the value of the bond – a slight increase in rates will lower P , while a slight decrease will do the opposite. In order to capture this effect we can use a measure known as duration, which is the approximate change in the price of a bond for a small change in rates; it is, in fact, the first derivative of price with respect to rates, and can be computed in a number of ways. The most fundamental method is to take the first derivative of the bond price equation above, and rearrange the terms to obtain a dollar-based change:

$$\frac{dP}{dr} = -\frac{1}{(1+r)} \left[\frac{1C}{(1+r)^1} + \frac{2C}{(1+r)^2} + \dots + \frac{nC}{(1+r)^n} + \frac{nM}{(1+r)^n} \right]$$

The bracketed term is the weighted average term to maturity of the bond's cash flows (where the weighted values are equal to the PV of the cash flows). An extension of this is the modified duration calculation, which indicates the approximate percentage change in price for a given change in yield:

$$MD = -\frac{dP}{dr} \frac{1}{P}$$

These equations permit an estimate of bond price changes (in dollar or percentage terms) for a small change in rates. For instance, if a bond has a modified duration of 10.5 and rates rise by 10 basis points (bps), the value of the bond will decline by approximately 1.05 %. This is obviously a useful measure for both valuation and risk management purposes, and is employed by intermediaries and end-users to evaluate the relative price riskiness of the assets we consider in the book. It is not, unfortunately, a sufficient measure. If rate moves are especially large, the duration measure does not provide an adequate measure of price or risk. This becomes apparent when we examine the price/rate relationships of a bond curve, which demonstrates a degree of curvature or convexity. Figure 2.2 illustrates the “slippage” that can occur when estimating bond value based on large rate moves.




Figure 2.2 Bond prices and convexity

In order to overcome this hurdle, we introduce the convexity measure, which improves the accuracy of bond price estimates for large rate moves. In fact, convexity is simply the second derivative of price with respect to rates, or the first derivative of price with respect to duration, and can be computed in dollar-based terms as:

$$\frac{d^2 P}{dr^2} = \sum_{t=1}^n \frac{t(t+1)C}{(1+r)^{t+2}} + \frac{n(n+1)M}{(1+r)^{n+2}}$$

The percentage change can be computed as:

$$Cvx = \frac{d^2 P}{dr^2} \frac{1}{P}$$

Alternatively, the percentage change in price due to convexity can be estimated as $\frac{1}{2}(Cvx)(\Delta r)^2$.

To approximate the total price change in a bond for a given move in rates, we use duration and convexity together. For instance, if a bond has a modified duration of 10.5 and convexity of 3.5 for a 200 bp change in rates, then the percentage price change for a 200 bp increase is -17.5% ($-(10.50 \times 2) + (3.5)$). Note that since convexity is positive for virtually all purchased fixed-income instruments, its impact on price is added to the equation.

Forward interest rates

The concept of forward interest rates proceeds logically from a discussion of the time value of money. In financial engineering we need to have a mechanism for estimating the potential future value of rates. As we shall note below, the pricing of many securities and derivative contracts is based on the yield curve, or the representation of interest rates with respect to maturity. In a normal market environment, the yield curve is upward sloping, meaning




Figure 2.3 Sample par, zero coupon, and implied forward curves

short-term rates are lower than long-term rates; flat and inverted yield curves, featuring short-term rates that are equal to, or greater than, long-term rates, are less common. While the construction of a standard, or “par,” yield curve is straightforward, the relative lack of bonds trading at par can create some shortcomings and distortions. In practice, the zero coupon curve is used to price many instruments.

Since financial instruments and contracts are based on future cash flows, the relevant valuation curve is based on forward, rather than par, rates. Fortunately, it is possible to estimate the implied forward rates from a par yield curve. This can be done by constructing a par curve from observable rates (i.e. coupons from instruments that are traded actively in the market),⁶ removing coupon effects by “stripping out” the zero coupon curve, and then using an implied forward algorithm to project a given series of implied forward rates. The end result is an implied forward curve, like the one shown in Figure 2.3. Naturally, since risk-free and risky interest rates change on a daily basis as a result of supply and demand forces, par rates, zero coupon rates, and implied forward rates change as well. Indeed, all three curves are part of a dynamic process.

An implied forward rate is a rate that can be set today for some future period, e.g. a six-month rate that commences three months, or six months, or twelve months, from today, and which makes an investor indifferent to two investment selections. Consider an investor that needs to invest funds for two periods. This investor can either:

- invest directly for two periods by buying a two period bond paying $[1 + r(0,2)]^2$; or
- invest in a one-period bond paying $[1 + r(0,1)]$ and reinvest the proceeds at the implied forward rate $r(1,2)$ between years 1 and 2.

As a result of arbitrage forces, the implied forward rate $r(1,2)$ will be set so that the two choices are identical, $[1 + r(0,2)]^2 = [1 + r(0,1)] * [1 + r(1,2)]$. These choices are illustrated in Figure 2.4.

⁶ Separate par curves can be created for each type of credit quality. For instance, a risk-free par curve can be developed by using AAA-rated government securities, a AA par credit curve can be created through AA credit-risky securities, and so forth. The development of separate credit curves is essential for accurate pricing and discounting of risky cash flows.




Figure 2.4 Implied forward rates

The process can be generalized through the generic formula for a y -period implied forward rate, starting in period x :

$$[1 + r_0(x, y)]^{y-x} = \frac{[1 + r(0, y)]^y}{[1 + r(0, x)]^x}$$

Note that we consider forward rates for non-interest rate assets (currencies, equities, commodities) later in the chapter.

2.3 DERIVATIVE INSTRUMENTS

Intermediaries rely on various instruments and vehicles in order to create the products we consider in this text; these building blocks include derivatives, host securities, and issuance/repackaging vehicles.

2.3.1 Derivatives

Derivatives, or financial contracts that derive their value from other market references, have evolved and expanded rapidly over the past few decades, and are now an integral part of the global capital markets. They have also become an essential building block of the financial marketplace. Though rudimentary derivative contracts have existed for centuries (primarily those on commodity references), the modern era of financial derivative contracts dates back to the early 1970s, when a number of exchanges began offering listed contracts on financial references such as interest rates, currencies, and equities. The customized over-the-counter (OTC) derivative market emerged in the early 1980s, when parallel currency loans⁷ were converted from balance sheet to off-balance-sheet status. These early currency swaps were

⁷ In a standard parallel currency loan, a US company might lend US\$ proceeds to the US subsidiary of a British company, while the British company might lend GBP proceeds to the UK subsidiary of the US company, thereby creating an efficient cross-border transaction. However, such loans, carried on the corporate balance sheet, have the effect of inflating total footings and attracting incremental capital charges.

soon followed by off-balance sheet interest rate swaps,⁸ commodity swaps, and equity swaps; by the early 1990s, various classes of swaps and forwards had become standardized and liquid. Additional structures, including credit derivatives, began appearing in the mid- to late 1990s, and innovation continues to the present time.

To be sure, not all synthetic or structured assets contain derivatives. For example, contracts such as asset-backed securities or stripped government securities are not created using derivative contracts. But many instruments rely on derivatives and derivative technologies. In this section we provide a very brief overview of the broad classes of OTC and exchange-traded derivatives. Since OTC derivatives, allow for considerable customization they are well suited to the product development requirements of the synthetic and structured asset sector. We shall therefore focus most of our attention on the OTC sector, revisiting the contracts periodically throughout the text, and in greater detail in Chapter 10.

Over-the-counter derivatives

OTC derivatives are the backbone of the synthetic and structured asset sector. The instruments facilitate structuring, restructuring, and repackaging of cash flows so that unique assets can be created. In this section we consider the general mechanics of several of the most common types of OTC derivatives, including forwards, swaps, and options; our discussion on options also focuses on certain exotic, or complex, contracts that have emerged over the past decade.

OTC derivatives, as the name suggests, are arranged and traded between two parties on an off-exchange basis. Each contract represents a customized negotiation of terms and conditions, with the parties agreeing to specific characteristics related to notional, term, reference market/asset, payoff profile, and so forth. Since dealing occurs off-exchange (either telephonically or via electronic communications networks), rather than through a formalized exchange, aspects of counterparty credit risk can appear. Indeed, in the absence of any specifically negotiated collateral/margin agreement between the two parties to a transaction, credit exposure arises for one or both parties (depending on the nature of the contract). Interestingly, some synthetic and structured assets have been created expressly to deal with these credit risks. For instance, an investor can enter into an asset swap package to convert a fixed rate bond into a synthetic floating-rate security, but assumes the credit risk of the swap counterparty in doing so; however, if the same package is embedded in a security issued by a repackaging vehicle, the counterparty credit risk is eliminated.

OTC derivatives can be used to achieve various goals. For instance, instruments such as forwards, swaps, and options, whether standalone or embedded in other instruments/vehicles, can be used to:

- hedge or transfer market and credit risks;
- reduce capital allocation charges on particular risks;
- establish a leveraged/unleveraged speculative position in an asset or market;
- diversify a portfolio;
- enhance the return/yield on a portfolio;
- create a customized asset/investment portfolio;
- rebalance the asset weightings in a portfolio;

⁸ In 1982, the US Student Loan Marketing Association (SLMA) imported the cross-currency swap technology that had emerged in the Euromarkets, converting its fixed rate funding into swapped 91-day Treasury-bill financing; others soon followed SLMA's lead, and the market grew rapidly from that point on.

- generate additional liquidity on an existing asset or liability;
- reshape cash flows, duration, and convexity;
- access an investment or funding market that may otherwise be restricted;
- monetize a gain on an asset;
- lower all-in funding costs;
- lock in future financing costs.

Since these applications are vitally important to efficient operation of the financial, corporate, and investment sectors, it is not surprising that OTC contracts have become extremely popular over the past two decades. Indeed, financial engineers in the banking world continue to create new derivative-based structures to capitalize on end-user needs and demand.

Forwards

A forward contract, the most elemental OTC derivative, is a single-period, bilateral contract that allows one party, known as the seller, to sell a particular reference asset at a forward price for settlement at a future date, and a second party, the buyer, to purchase the reference asset at the forward price on the named date. A notional amount is used as a reference to compute the amount payable/receivable at maturity; no initial or intervening cash flows are exchanged between the two parties, and settlement of the contract at maturity may be set in physical or financial terms. If the market price at maturity is greater than the contracted forward price set at trade date, the buyer generates a profit and the seller a loss. If the market price is lower than the forward price, then the seller profits and the buyer loses. These relationships hold true for all price-based forwards, including those involving equities, bonds, indexes, currencies, and commodities. The profit positions of generic price-based long and short forwards are summarized below, and the flows are illustrated in Figure 2.5.

Long forward profit: $(P_{\text{mat}} - P_{\text{forward}}) * \text{Notional}$
Short forward profit: $(P_{\text{forward}} - P_{\text{mat}}) * \text{Notional}$

where

P_{mat} is the prevailing market price at maturity of the forward contract
 P_{forward} is the forward price contracted between the two parties on trade date.

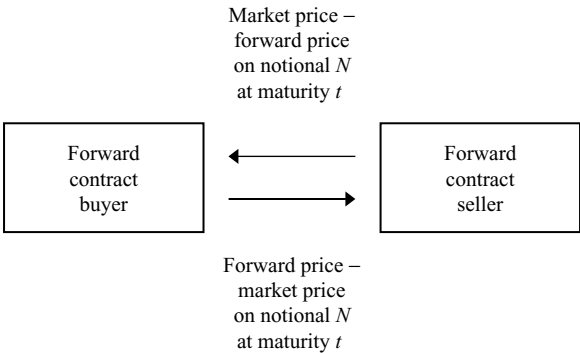


Figure 2.5 Forward contract flows

Credit risk issues are important, as either party can be exposed to potential profits- and thus credit losses in the event of counterparty default. Since no intervening cash flows are exchanged during the life of the contract, value (and therefore exposure) can accumulate steadily. Forward contracts can be written on virtually any asset from any market sector; forwards on equities, bonds, currencies, credits, and commodities are very common, with maturities ranging from one week to many years. Settlement can be in net financial terms or via delivery of the asset.

Let us consider an example of a forward. Company A owns a market asset that it wishes to sell on a six-month forward basis at a price of \$100, and Company B wishes to buy the asset at the forward price of \$100; the two thus enter into a forward contract calling for the purchase/sale of the asset at \$100 in six months. Over the next six months, the market price of the asset fluctuates on a daily basis: as it rises above \$100, Company B experiences an unrealized gain, and as it falls below \$100, Company A benefits from the unrealized gain. However, realized profit/loss comes only at the conclusion of the contract. If the market price in six months is \$105, Company B realizes a profit: it purchases the asset from Company A at the agreed forward price of \$100 and sells it in the spot market for \$105, posting a \$5 gain. Company A, in turn, is obliged to deliver the asset and receive \$100; it posts a loss on a forward basis, as it could have sold the asset, in the market at \$105 absent the forward agreement. If it does not physically own the asset, it must purchase it in the market for \$105, receiving only \$100 in cash proceeds from Company B, thus generating a \$5 loss.

The reverse scenario is also possible: if the market price declines to \$94 at maturity, Company A delivers the asset and receives \$100 (a gain of \$6 versus the market price), while Company B is forced to pay \$100 for an asset that it could have purchased in the spot market at \$94, losing \$6. The issue of whether either party ultimately sustains a net loss depends on whether the forward contract serves as a hedge, or is intended purely as a speculative position. For instance, if Company B has an underlying obligation to cover a short position by taking delivery of the asset at a price of \$100, it may be indifferent to the actual settlement price, as any loss on the forward will be offset by a gain on the underlying position (and vice versa). The same may be true of Company A and its own balance sheet position: if it owns a quantity of the asset that it has sold on a forward basis, a loss on the forward will be offset by a gain on the long position (and vice versa). However, if either or both are using the contract as a speculative position, then a net gain or loss will result.⁹ Figures 2.6 and 2.7 summarize the payoff profiles of the long and short forwards; Figure 2.8 also illustrates the effects of a long forward acting as a hedge against a short asset position; the reverse hedge position can also be created.

The simple examples above apply equally to all price-based forwards. However, given the inverse relationship between interest rates and prices, we first consider the special case of the interest rate forward, or forward rate agreement (FRA), separately.¹⁰ The FRA is a single-period, cash-settled contract with a forward rate set on trade date that is compared to the prevailing market rate on the expiry date to generate a net payment between the two parties. For instance, using market convention terminology, a 2×5 month FRA calls for setting a

⁹ Since the generic forward contract features no initial or intervening cash flows, a long position can be viewed as a leveraged position in an asset, with one party borrowing to buy the asset at some future time. Since the contract is risky but requires no investment, it earns a risk premium. The prepaid forward, an associated structure, is identical to a standard forward, except that a payment is exchanged on trade date (rather than the forward date), suggesting an initial investment is required; however, the asset exchange (or financial settlement) is concluded on the forward date (as in a standard forward). We shall revisit the prepaid forward in Chapter 5.

¹⁰ FRAs are used widely to create synthetic instruments and hedge assets and liabilities that are exposed to short-term interest rate movements, primarily London Interbank Offered Rate (LIBOR) or an equivalent short-term deposit benchmark.




Figure 2.6 Long forward position




Figure 2.7 Short forward position

three-month forward rate (e.g. five months — two months) on trade date, and comparing that rate with the prevailing three-month market rate in two months' time, when the contract expires. Thus, Bank A might agree to pay Bank B on \$100 m notional if three-month LIBOR is above 4 % in two months' time, and receive if LIBOR is below 4 %; this, of course, is simply a form of synthetic floating-rate borrowing or lending. If LIBOR sets at 5 % in two months' time,

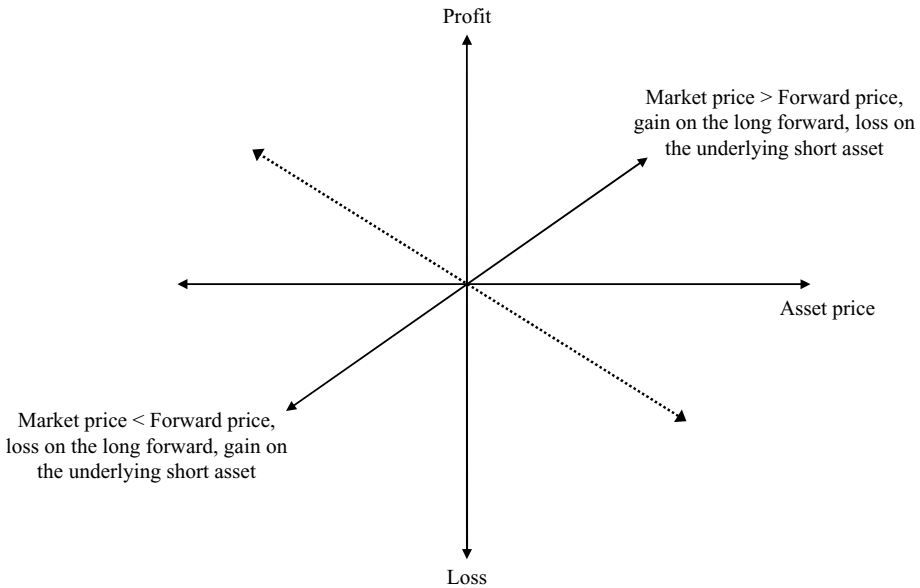


Figure 2.8 Long forward/short asset hedge position

Bank A pays Bank B \$250 000. The net cash settled FRA payments/receipts are computed via the standard formula:

$$NP * \frac{[(MR - FR) * (D/360)]}{[(1 + MR) * (D/360)]}$$

where

NP is notional principal

MR is the prevailing market rate at contract maturity

FR is the contracted forward rate on trade date

D is the number of days in the forward contract.

Though the fair value of any “on market” forward on trade date is 0, the two parties must have different expectations regarding the future movement of the reference (or will need to use the contract for specific hedge purposes). The forward price reflects costs and benefits: costs can include financing, insurance, and/or transportation, while benefits may include yield or dividend, convenience yield, and/or a lending rate. If a quoted forward price is greater than the price implied by the theoretical fair value, an arbitrage opportunity arises (e.g. selling the forward, borrowing and purchasing spot, and lending until the maturity of the forward); the reverse occurs when the quoted price is below the implied fair value price.

Determining a theoretically fair forward price involves the use of the current (or spot) asset price and the net cost of carry (which incorporates the time value concepts introduced above). The net cost of carry can be regarded as a combination of the cost of funding less any return derived from holding the asset (e.g. dividends, coupon). Since the cost of funding can never be negative, the forward price must be higher than the spot price in a freely arbitrageable

market with a minimum of frictions (we consider some exceptions to this rule below). If the cost of funding rises, the price of the forward must rise as well. The forward seller therefore benefits if rates decline (all other variables, including the current price, remaining constant); similarly, if rates rise, the forward buyer will profit. Note that forwards with longer maturities are more sensitive to rate movements than short-term ones, as the cost of carry has a larger impact (that said, many forward positions are funded with short-term floating-rate liabilities). The return that can be earned depends on the asset class. For instance, a forward involving an equity portfolio generates a net return equal to the difference between the cost of financing the equities (i.e. a short-term borrowing rate) and the dividend yield on the portfolio; a bond/interest rate forward creates a net return equal to the difference between the cost of financing and the coupon on the bond; and, a foreign exchange forward yields a net return equal to the difference between the two reference interest rates on the underlying currencies. The return on a physical commodity forward is slightly more complex, as it must take account of additional costs of physically possessing the commodity, including storage, insurance, and transportation; in fact, the net return is the difference between the cost of financing and holding the commodity and the yield generated by the commodity (which is often obtained through the lease market). It is important to re-emphasize that while the net cost of carry is important in helping establish the value of the forward, the current asset price remains the dominant price determinant. Figure 2.9 summarizes the main pricing components of a forward.

We have stated that the forward price should be higher than the spot price because financing costs account for the largest portion of the carry component, and financing costs can never be negative. This relationship, known as contango, occurs when the asset can be borrowed/lent

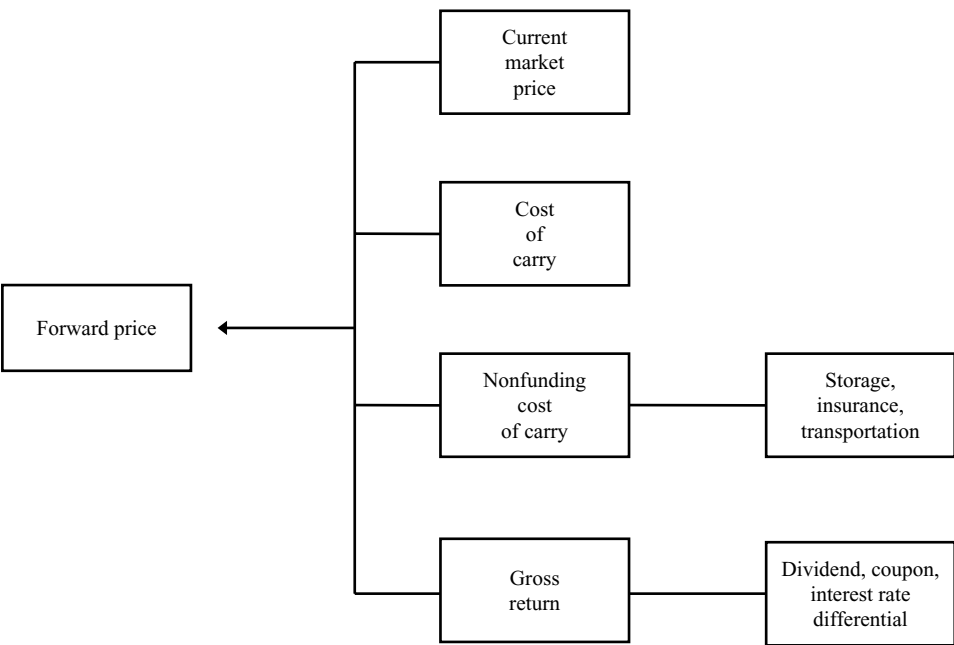


Figure 2.9 Pricing components of a forward

easily and few friction costs or barriers exist. Arbitrageurs help enforce the pricing rule by intervening whenever discrepancies arise. In practice, financial assets, which display good liquidity and can be transferred with ease, are contango markets. However, there are times when the forward price is lower than the spot price; this characteristic, known as backwardation, tends to arise when an asset/market is not freely arbitrageable, and the cost of carry involves more than simple financing. The most obvious examples relate to forwards on energy, agriculture, and metals, where the nonfunding carry involves transportation, storage, and insurance, and where a commodity may not be readily available for buying and selling by arbitrageurs. When this situation exists, the market may become “one-way” for an extended period of time.

With these basic concepts in mind, it is relatively easy to derive forward prices for general classes of assets:

Forward contract on stock that pays no dividends: $F_{0,T} = S_0 e^{rT}$

where

$F_{0,T}$ is the forward price at maturity T

r is the financing rate

S_0 is the stock price.

Forward contract on stock that pays continuous dividends: $F_{0,T} = S_0 e^{(r-\delta)T}$

where

δ is the continuous dividend yield

all other terms are as defined above.

Forward contract on stock that pays discrete dividends: $F_{0,T} = S_0 - \sum_{i=1}^n PV_{0,t_i}(D_{t_i})$

where

$PV_{0,t}(D_{t_i})$ is the dividend payable in period t_i

all other terms are as defined above.

Forward contract on currency: $F_{0,T} = x_0 e^{(r-r_y)T}$

where

r is the financing rate of currency 1

r_y is the financing rate of currency 2

x_0 is the spot rate between currencies 1 and 2

all other terms are as defined above.

Forward contract on commodity with a lease market: $F_{0,T} = S_0 e^{(r-\varepsilon)T}$

where

ε is the lease rate

all other terms are as defined above.

Forward contract on commodity with a lease and carry market, and a convenience yield:

$$F_{0,T} = S_0 e^{(r+\lambda-c)T}$$

where

λ is the storage cost

c is the convenience yield (i.e. the economic benefit derived from having immediate access to a particular commodity)

all other terms are as defined above.

Note that the forward rate on interest rates reverts to our discussion on implied forwards from earlier in the chapter. As a reminder, the y period implied forward rate starting in a future period x is:

$$[1 + r_0(x, y)]^{y-x} = \frac{[1 + r(0, y)]^y}{[1 + r(0, x)]^x}$$

Swaps

A swap contract is the second major component of the OTC derivatives market, and can be viewed as a package of sequential forward contracts that mature at successive periods in the future, until the stated maturity date; we shall discuss this analogy at greater length in Chapter 10.¹¹ Swaps, like forwards, can be written on virtually any asset from any market sector. While interest rate swaps still account for the largest share of the market, active dealing occurs in equity, currency, commodity, and credit swaps. An emerging market has also started to appear in other types of swaps, including those centered on inflation, macroeconomic indicators (e.g. GDP), real estate, and weather references (e.g. temperature, precipitation); as we might imagine, these are highly specialized and customized transactions with a limited audience. Maturities in the swap marketplace range from approximately one to ten years, though certain 20- to 30-year transactions appear from time to time. Since swaps are bilateral contracts, credit risk issues are again of considerable importance; indeed, these are often intensified by the multi-year maturities associated with standard transactions, though they are partly mitigated by the periodic exchange of net cash flows, which can help prevent the mark-to-market value of a position from accumulating to an excessively large amount. As noted below, currency swaps, which involve the initial and final exchange of notional principal, feature an extra element of risk; the parties face delivery risk exposure on each exchange, which is akin to 100 % risk for a very short period of time (e.g. intraday or overnight, as currency flows settle).

An interest rate swap is a contract where one party agrees to pay a fixed interest rate and a second party agrees to pay a floating rate; the fixed rate payer is said to be “long the swap,” or to have “bought the swap,” because the floating rate is considered to be the “deliverable commodity” (e.g. the fixed-rate payer is paying a fixed rate to receive the deliverable floating rate). Apart from the special case of cross currency swaps, notional principal is not exchanged in the swap; in fact, the notional is used only to compute flows payable and receivable at each intervening evaluation period. On each date, which may be monthly, quarterly, semi-annually, or annually, the fixed and floating rates are compared, and a net payment is arranged between the two parties. If the floating rate (generally a recognized index such as LIBOR), is above the fixed rate, the floating-rate payer makes a net payment to the fixed-rate payer in the amount of

¹¹ It is worth noting that a swap can also be viewed as a package of a fixed-rate bond and a floating-rate note, or a strip of deposit futures; we shall, however, save discussion of these comparisons for Chapter 10.

(floating rate — fixed rate) * notional principal. If the floating rate is below the fixed rate, the fixed-rate payer makes a payment to the floating-rate payer. The process continues to the next evaluation/settlement period, and so forth, until maturity. The payment of accumulated value at each settlement period reduces credit exposure.

Consider, for instance, Company A, which wants to lock in a fixed-rate payment on \$100 m of its floating-rate liabilities over the next five years, and Bank B, which is willing to pay a floating rate in return. The two parties enter into a five year swap transaction on \$100 m notional, where A pays a fixed rate that is established on trade date in exchange for semi-annual LIBOR. Every six months for the next five years, A and B will exchange net payments on the \$100 m notional based on the level of LIBOR in relation to fixed rates. As LIBOR exceeds the fixed rate, Company A will generate a gain for the current period, which it can use to offset the higher cost on its floating-rate liabilities; as LIBOR falls below the fixed rate, it will generate a loss on the swap, but will pay less on its liabilities. By entering into the transaction, Company A faces a known cost of funding (e.g. a fixed, rather than floating, rate), while Bank B enjoys the floating-rate flow it requires for its own operations. Figure 2.10 illustrates the flows of a basic swap.

Interest rate swaps may be arranged for funding, hedging, investing, and speculating purposes. Funding activity based on credit arbitrage that appears between different segments of the capital markets is particularly important; the arbitrage opportunity can appear as a result of various forces, including information inefficiencies, regulatory and tax barriers, and market frictions. The credit arbitrage is simply the credit differential that appears between issuers and counterparties that raise fixed- and floating-rate liabilities. In general, strong investment grade companies (i.e. AA–AAA) enjoy an absolute advantage in raising fixed- and floating-rate funds, and a comparative advantage in raising fixed-rate funds. Lower-rated companies (i.e. BBB and below) often have a comparative advantage in raising floating-rate funds. When the differential suggested by these comparative advantages becomes wide enough, swaps can be arranged to benefit both parties simultaneously.

Let us consider the funding levels of two firms, Company Y (which is rated AAA) and Company X (which is rated BBB–) to demonstrate how they can use the swap market to achieve better all-in funding results. Assume that Company Y can raise fixed-rate funds at 4.5 % and floating-rate funds at LIBOR + 25 bps, while Company X can raise fixed rate funds at 6 % and floating-rate funds at LIBOR + 75 bps. While it is clear that Company Y can raise funds

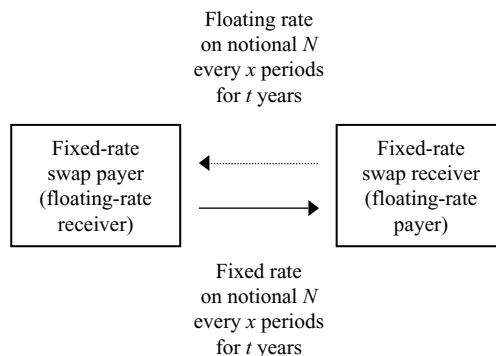


Figure 2.10 Interest rate swap flows

Table 2.1 Credit arbitrage flows

	Company Y	Company X
Rating	AAA	BBB–
Fixed-rate funding levels	4.50 %	6.00 %
Floating-rate funding levels	LIBOR + 0.25 %	LIBOR + 0.75 %
Comparative advantage fixed	+1.50 %	
Comparative advantage floating		+ 0.50 %
Fixed rate from X to Y on swap	+5.25 %	–5.25 %
Floating rate to X from Y on swap	– LIBOR + 0.25 %	+ LIBOR + 0.25 %
Fixed rate to bond investors	–4.50 %	
Floating rate to FRN investors		– LIBOR + 0.75 %
Net swapped funding cost	LIBOR – 0.50 %	5.75 %
Direct funding cost	LIBOR + 0.25 %	6.00 %
Net savings	+ 0.75 %	+ 0.25 %

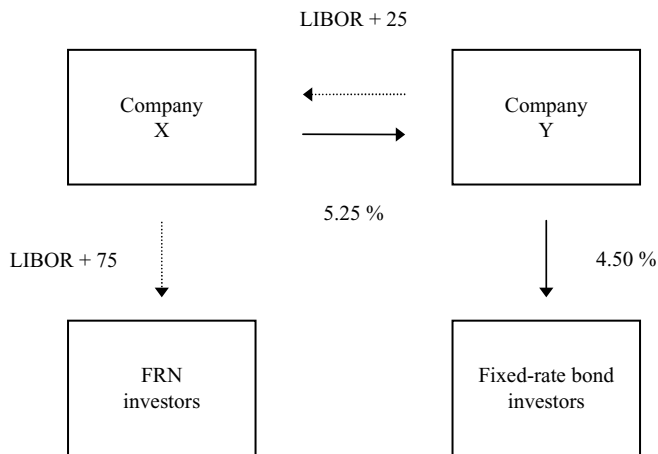


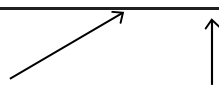
Figure 2.11 Credit arbitrage using interest rate swaps

at a cheaper level in either market (i.e. the absolute advantage), we notice that the fixed-rate funding differential is 150 bps (i.e. 6 % – 4.5 %), while the floating rate funding differential is only 50 bps (i.e. LIBOR + 75 – LIBOR + 25). Company X thus has a comparative advantage in the floating-rate market and Company Y has a comparative advantage in the fixed-rate market, giving rise to a credit arbitrage opportunity of 100 bps (i.e. 150 bps – 50 bps). Both companies can lower their funding costs by using the swap market arbitrage. Assume that the two firms decide to split the advantage 75/25 – Company Y, with the better credit, enjoys more of the benefits, but both companies gain in the process.¹² Company Y thus issues its fixed-rate bond at 4.5 % and Company X launches an FRN at LIBOR + 75; the two then enter into a swap where Y pays X LIBOR + 25 and receives 5.25 % fixed. By doing so, the two crystallize reduced funding costs: Y pays 75 bps per annum less than it would by launching a straight fixed-rate bond, while X pays 25 bps per annum less than it would in the FRN market. These arbitrage results, summarized in Table 2.1 and Figure 2.11, provide a powerful incentive to

¹² In practice, the arbitrage can be split in any number of arbitrary ways.

Table 2.2 Sample swap quotes

UST 5-year bid-offer (price)	UST 5-year offer-side yield (%)	Swap spread (bps)	Semi-annual swap price (%)
101.09–101.10	4.25	65–68	4.90–4.92



Intermediary pays the
bid in a fixed swap

Intermediary receives
the offer in a fixed swap

swap. It is worth noting that while liability-based credit arbitrage is an important driver of swap activity, a similar process occurs on the asset side of the balance sheet; we shall reserve this discussion for Chapter 10 when we consider asset swap packages and callable/puttable asset swaps.

Interest rate swaps can also be used to express an outright view of the market (e.g. entering into a swap to pay fixed rates and receive floating-rates if rates are expected to rise, the reverse if rates are expected to fall). They can also be used to hedge or protect a market exposure. Thus, as in our first example above, if a company raises floating-rate funds, it can enter into a swap to hedge against rising rates; it may do so by receiving floating in the swap and paying fixed; if rates rise, it will pay a higher coupon on its funding, but receive the same higher rate on the floating leg of the swap, effectively hedging its liability. Swaps can reshape the cash flows or duration of an asset or liability portfolio. For instance, if an investment fund seeks to lengthen the duration of an asset portfolio, it can receive fixed and pay floating in the swap market (recalling that this is equivalent to buying a bond (long maturity/duration) and shorting an FRN (short maturity/duration)); if it chooses to shorten the duration, it can pay fixed and receive floating. Naturally, it can select the reverse positions for his liability portfolio. Finally, swaps can be used to reduce costs or enhance yields. For instance, if swap rates are higher than cash market rates, a swap can be used to lengthen asset portfolio maturity or shorten liability portfolio maturity. If cash rates are higher than swap rates, a swap can be used to shorten asset maturity or lengthen liability maturity.

Interest rate swaps are quoted typically as a basis point spread to the yield of a government benchmark, such as US Treasuries, Japanese Government Bonds (JGBs), gilts, or bunds. This quoted spread represents the fixed rate payable or receivable on the swap against LIBOR flat.¹³ Thus, the bid-offer spread on a US\$ interest rate swap begins with the US Treasury bid-offer spread; after converting this price spread to the offer-side yield, the swap spread is added to obtain the swap bid-offer. A swap broker or intermediary pays the bid spread to the fixed rate receiver and demands the offer spread from the fixed-rate payer, as illustrated in Table 2.2.

¹³ Note that the quote provided in the marketplace is on the same day count/payment frequency basis as the underlying government benchmark (e.g. in the US, it is the semi-annual bond equivalent yield). Since the market quotes swaps to LIBOR flat, a company that is securing floating-rate funding that it wishes to swap must adjust the LIBOR spread to obtain the correct all-in funding cost. This can be done by starting with the relevant US Treasury yield, adding in the swap spread, adding in the LIBOR funding spread, and then calibrating by the relevant day count adjustment (e.g. converting LIBOR from Actual/360 to a bond-equivalent yield via a 365/360 scaling on the LIBOR funding spread). For instance, if the seven-year swap spread is UST + 75 bps, and the company's LIBOR funding is +50 bps, the all-in funding cost after scaling for 365/360 is UST + 126 bps. If the seven-year UST rate is 5 %, then the company's all-in swapped funding cost is 8.26 %. The reverse computation – switching from floating- to fixed-rate funding – can be computed in a similar manner, with the all-in cost equal to LIBOR + (fixed-funding spread – swap spread) * 360/365.

Basis swaps, sometimes known as money market swaps, involve the exchange of two floating-rate flows. Such swaps help hedge against, or establish a position in, the differential between the floating rates, and are commonly used by banks, investment funds, and companies to manage money market portfolios and treasury exposures. Common basis swaps include commercial paper (CP) versus LIBOR, Treasury bills versus LIBOR, Prime versus LIBOR, and EURIBOR versus EONIA, among others. In most markets, basis swaps are quoted as the index \pm a spread to a core reference flat (e.g. six-month LIBOR flat). For instance, a two-year US CP/LIBOR swap might be quoted at 25–27 bps, meaning the dealer will pay CP +25 bps versus LIBOR flat, or receive CP + 27 bps versus LIBOR flat. A two-year Prime/LIBOR swap, in turn, might be quoted at $-125/-123$, indicating that the dealer will pay Prime -125 bps versus LIBOR flat, and receive Prime -123 bps versus LIBOR flat. Similar conventions exist in non-\$ markets.

As indicated earlier, cross-currency swaps, involving the exchange of currency flows, were the first OTC swaps in the marketplace. Two parties to a standard currency swap agree to an initial exchange of principal, a series of periodic foreign currency payments, and then a final exchange of principal. As noted, this is distinct from interest rate and basis swap structures, which involve no exchange of notional principal. Currency flows may be swapped on a fixed/fixed, fixed/floating, or floating/floating basis. Fixed/fixed exchange currency swaps are equivalent to a series of forward foreign exchange contracts based on current spot rates (rather than forward rates used in forward foreign exchange contracts). Fixed/floating swaps, which exchange a floating rate in one currency for a fixed rate in a second currency, are more common. For instance, Company A may enter into a five-year \$/yen currency swap with Bank XYZ. On trade date, A delivers yen to XYZ in exchange for dollars, on each semi-annual date for the next five years it pays \$ LIBOR in exchange for a fixed yen coupon, and at maturity in five years it returns the dollar principal in exchange for its original yen principal. Figure 2.12 illustrates these flows.

Currency swaps can be used to alter an existing exposure, reduce costs, or enhance yield. For instance, if a firm wants to swap into a foreign currency, it can receive the foreign currency on its assets, or pay the foreign currency on its liabilities. If it wants to swap out of a foreign currency, it can pay the foreign currency on its assets and receive the foreign currency on its liabilities. Similarly, if foreign currency swap rates are greater than cash market rates, it can swap its assets into the foreign currency, or swap its liabilities out of the foreign currency; if swap rates are lower than cash rates, it can swap its assets out of the foreign currency and swap its liabilities into the foreign currency. Cross currency swaps are quoted as a fixed foreign currency payment against \$ LIBOR flat. Accordingly, changes in the quote level and spread are a function of both interest rates and foreign exchange rates. The receiver of a fixed foreign currency rate receives the bid rate and pays \$ LIBOR; the payer of the fixed foreign currency rate pays the offer rate and receives \$ LIBOR.

Standard commodity, equity, and credit swaps work in a similar fashion. Most structures involve the exchange of a fixed price for a floating price, and feature notional, maturity, payment frequency, and physical/financial settlement terms, just as interest rate or currency swaps do. Alternate structures can, of course, be created, as we shall note later in the text.

All on-market swaps, regardless of type, are priced to a zero net present value (NPV) on trade date. That is, at the inception of a deal, a swap will reflect a zero profit and loss position for each party.¹⁴ Thereafter, the NPV starts to change: as the swap is revalued each day with

¹⁴ Premium/discount swaps are off-market structures and do not, therefore, reflect zero NPV.

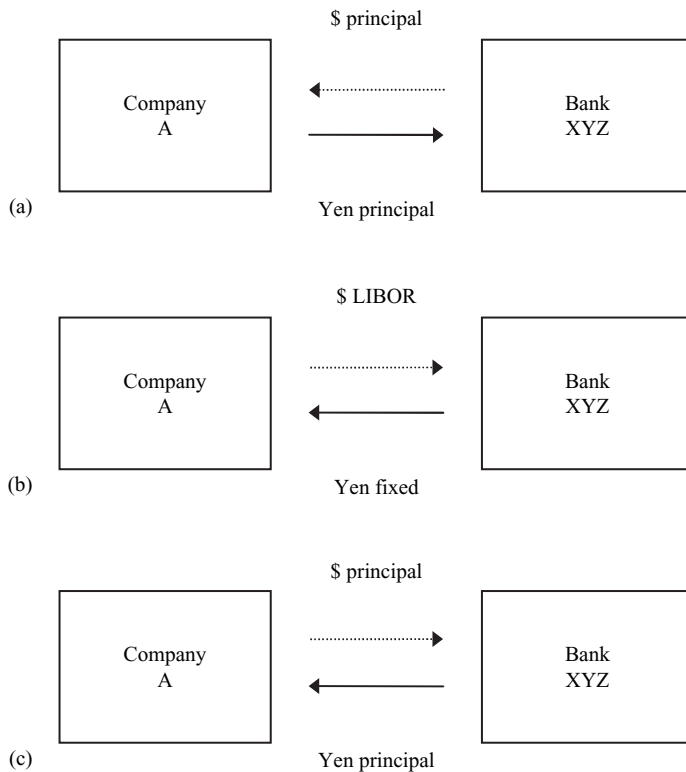


Figure 2.12 Cross-currency swap. (a) initial exchange, trade date; (b) periodic coupon exchange, semi-annual dates; (c) final exchange, maturity

changing rates (or other relevant market variables), the NPV will favor one of the two parties. For instance, as interest rates rise, the cash flows and NPV on the floating leg of an interest rate swap will increase, e.g. the present value of the floating-rate inflows becomes larger than the fixed-payment outflows; the reverse occurs as rates fall.

Options

Options comprise the third major segment of the OTC derivatives market (and also form part of the exchange-traded market, as noted below). The options market can be divided broadly into standard, or vanilla, contracts, and exotic, or complex, contracts. Standard contracts, including call and put options, are the basic instruments of the options marketplace and account for most market activity.

A call option is a unilateral contract that gives the buyer the right, but not the obligation, to purchase from the seller a specified asset at a predetermined strike price. In order to secure this right, the buyer pays the seller an option premium. The right may be exercised at any time until maturity (American option), at maturity only (European option), or at specified points up until maturity (Bermudan option). If the buyer exercises the option (e.g. the price of the underlying reference asset is above the strike price), he delivers the required amount of cash (defined by

strike price * number of units) and receives the underlying asset; he may then liquidate the asset in the marketplace at the higher prevailing market price, or may hold the asset for future use. The seller must be prepared to deliver the specified asset should exercise occur. If the seller does not own the asset, he must purchase the relevant amount from the marketplace at the prevailing market price (this process is known as naked, or uncovered, call writing). If the seller owns the asset, he simply delivers it upon exercise (this is known as covered call writing). If the buyer does not exercise the option (i.e. the price of the underlying reference asset is below the strike price), the contract expires, meaning that the seller has no further performance obligation.

A put option functions in a similar manner, giving the buyer the right, but not the obligation, to sell a specified asset at a predetermined strike price. The buyer will exercise the put option when the market price is below the strike price, delivering the specified asset to the seller in exchange for the required amount of cash. Clearly, the buyer will not exercise the contract when the market price is above the strike price. The seller, in turn, is obliged to acquire the asset upon exercise by delivering cash; if she possesses the cash, she is essentially writing covered puts, and if she does not, she is writing naked puts.

Figures 2.13 through 2.16 illustrate long and short call and put option payoff profiles. OTC options can be purchased and sold on a broad range of assets and markets. Contracts are arranged regularly on references from the fixed-income markets (including short- and long-term interest rates, interest rate volatility, credit spreads, credit defaults), equity markets (including single stocks, baskets, indexes, equity volatility), currency markets, commodity markets (including energy, softs, metals, agriculturals), and other noncommodities (including weather, inflation, catastrophe). Option maturities can range from overnight to several years, and settlement can be arranged in both physical and financial terms.




Figure 2.13 Long call option




Figure 2.14 Short call option




Figure 2.15 Long put position

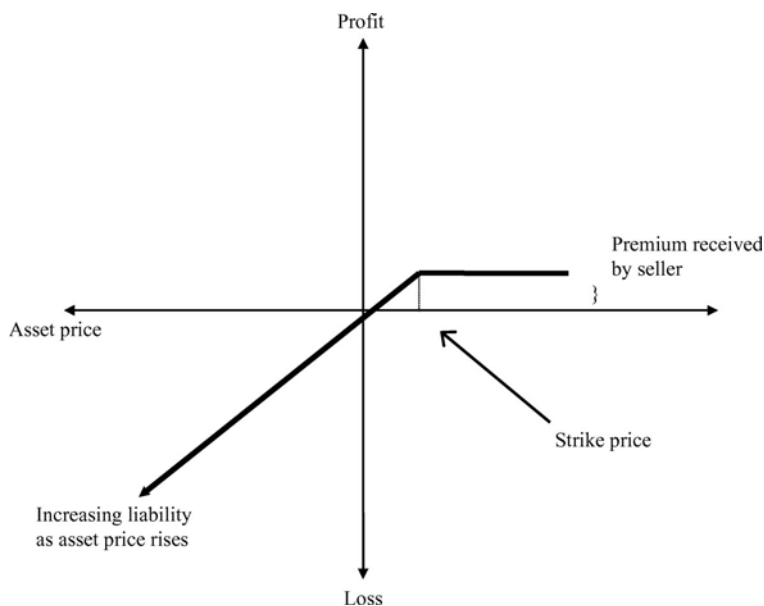


Figure 2.16 Short put position

Equity-related options comprise one of the most popular segments of the OTC option market, with intermediaries and end-users actively dealing in puts and calls on individual stocks, sectors, and broader market indexes. The market has become global, with contracts arranged on equity references in both developed and emerging equity markets. Equity options are used routinely to hedge underlying exposures, or to express a leveraged view of equity market direction and/or volatility.

The interest rate option market is also very active, with puts and calls written and purchased on a range of interest rate references, including money market instruments and government notes/bonds from both developed and emerging markets. Caps and floors are also an important component of the sector, and are often used in the creation of synthetic assets. Caps (which are equivalent to a strip of puts on the price of a floating rate security or deposit futures, or a strip of calls on the rate level) and floors (a strip of calls on the price of the security or deposit futures, or a strip of puts on the rate level) are used for funding, investment, and risk management purposes. Caps and floors generally reference an actively traded floating-rate benchmark such as LIBOR, EURIBOR, TIBOR, Treasury bill rates, bank bill rates, or CP rates; given the close relationship between caps/floors and swaps, most activity is centered in LIBOR and other interbank rates. Let us consider a standard LIBOR cap (which, again, is equivalent to a put on Eurodollar futures)¹⁵: the buyer pays the seller a premium for a cap with a six-month LIBOR strike and a 12-month maturity date (e.g. two evaluation periods, suggesting the creation of two “caplets”). If LIBOR rises above the strike during either or both of the evaluation periods,

¹⁵ There are certain subtle differences between caps and deposit future puts (and floors and deposit future calls). For instance: the holder of the cap receives a cash payment from the seller, while the holder of the put has an asset that varies with market prices; cap payments are generally received in arrears (in accordance with swap market convention); some basis point value differences can arise, depending on the notional amount of the cap; and expiration dates and possible reference fixings can vary (e.g. London a.m. fixing for caps versus London p.m. fixing for deposit futures). Nevertheless, the general structures are sufficiently similar that the two are regarded as fungible.

the seller is obliged to pay the buyer the difference between prevailing LIBOR and the strike, times the notional amount of the cap. The buyer has thus “capped” the maximum LIBOR level at the strike. The reverse occurs in the case of a floor: the buyer effectively locks in a minimum interest rate equal to the strike; if LIBOR falls below the floor, the seller is obliged to make a compensatory payment equal to the difference between the market rate and the strike, times the notional. Caps and floors are quoted in terms of maturity and strike, and dealers generally provide a two-way price at which they will buy or sell a particular option structure. For instance, a three-year LIBOR cap struck at 4 % might be quoted at 40–43 bps, where the bid of 40 bps represents the price the dealer will pay for the cap, and 43 bps represents the price the dealer receives for the cap.¹⁶

Multiple put/call options are often used to reshape the basic risk and return profiles presented in the diagrams above. We shall review some of the most common multiple option strategies in our discussion on derivatives structuring in Chapter 10. While vanilla puts and calls on various asset references dominate the OTC option market, certain exotic options with unique payoff characteristics are also available. Though some of these instruments are difficult to price and risk manage, they produce unique results that can help satisfy specific end-user requirements. As we shall note later in the book, these exotic options are incorporated periodically in other contracts to produce synthetic assets with unique profiles. The list of exotic options is extensive and beyond the scope of our introductory discussion,¹⁷ but some of the most common contracts include:

- Asian options (also known as average options): options that derive their payoff from the average price of the underlying reference asset during a set time period, rather than at a single point in time (e.g. maturity). The Asian option can be structured with an averaging period applied to the strike (average strike option) or price/rate (average price option).
- Barrier options (also known as knock-in or knock-out options): options that are either created (knock-in) or extinguished (knock-out) when the price of the underlying reference asset breaches the barrier. Knock-ins are available in the form of down-and-in calls/puts and up-and-in calls/puts, while knock-outs are available as down-and-out calls/puts and up-and-out calls/puts. Those that knock-in or knock-out with intrinsic value are also known as reverse knock-ins and reverse knock-outs.
- Compound options (also known as nested options): options that permit exercise into a second option. The compound option is available as a call on a put or a call, and as a put on a put or a call.
- Digital options (also known as binary options): options that provide a fixed payout if the strike is breached, rather than a continuum of payouts that depends on the price of the underlying asset in relation to the strike price (as in a standard option).
- Power options options that provide a payout that is based on an exponential leveraging of the strike compared to the underlying asset price; the result is an exponentially increasing gain (or liability) as the option moves further in-the-money.
- Rainbow options multivariate options that provide a payout based on the simultaneous performance of several different assets. The rainbow option can be structured in various forms, including call/put on the best or worst of n assets, spread option, and basket option.

¹⁶ Each individual caplet or floorlet comprising a cap or floor has its own value; the total value of the cap or floor is simply the sum of the individual components.

¹⁷ Readers interested in greater detail may wish to consult Banks (2003).

- Quanto options (also known as guaranteed exchange options): options that convert gains generated by a foreign currency contract back into a home currency, leaving the underlying asset protected from currency movements.

Since the seller of an option must perform if the price of the reference moves above/below the strike price, he requires risk compensation from the buyer in the form of an option premium. The price of an option (e.g. premium paid or received) is obtained through the option valuation process. Premium is comprised of two elements – intrinsic value and time value. Intrinsic value is the difference between the strike price and the current market price, and indicates a contract's degree of "moneyness"; a contract that is "in-the-money" can be exercised or sold for immediate gain. A call option has intrinsic value when the market price is above the strike price; a put option has intrinsic value when the market price is below the strike price. Those for which the market and strike prices are precisely equal are considered to be "at-the-money" and have no intrinsic value; similarly, those for which the market price is below (calls) or above (puts) the strike price are "out-of-the-money" and have no intrinsic value. Obviously, options with intrinsic value are more valuable than those that lack such value, and result in a larger option premium. Time value is the second major element of option premium. Time is useful, as it gives the buyer more opportunity to achieve a positive result on the option contract; as the maturity of a contract is extended, there is more time for the reference asset to move and cause the option to migrate in-the-money (or deeper in-the-money). Accordingly, options with long maturities possess greater time value and command a higher premium; those with short maturities have less time value and are thus cheaper. Time must also be considered a wasting asset: with each passing day, the time value component of the contract erodes; this "time decay" works in favor of the option seller, who stands to benefit by preserving more premium as maturity draws nearer. In fact, time decay accelerates as maturity draws nearer; very short-term options lose value more rapidly than longer-term options.

Option prices are derived from mathematical models based on various inputs. Arbitrage forces help ensure that market prices and theoretical prices track one another closely over time. Though detailed derivation of the seminal Black–Scholes model is beyond the scope of our discussion, we note in general terms the logic behind the framework and the key inputs used in the process. Specifically, the value of an option is dependent on the current price of the reference asset, the strike price, the risk-free rate, the volatility of the reference asset, and the time to maturity. Use of each of these variables makes intuitive, as well as mathematical, sense. For instance, the current asset price and the strike price are both essential in determining the intrinsic value of a contract; the further in-the-money an option, the greater the premium payable/receivable. Time to maturity is also vital for the reasons we have noted above. Volatility is a central input, as it indicates the degree to which a contract may move in-the-money; greater volatility increases the likelihood that an asset will move the contract in-the-money, and therefore increases the contract price. The risk-free rate is essential in order to discount future cash flows to the present. In practice, each of these variables is observable directly in the marketplace, except for volatility. In order to cope with the lack of observable volatility, participants trade options by determining the volatility implied in the prices of outstanding contracts. The option-pricing inputs used to derive theoretical option prices are summarized in Figure 2.17.

The Black–Scholes model, designed as a closed-form solution for European options on non-dividend-paying stocks, is an essential approach to option valuation. Subsequent refinements and approaches (e.g. the binomial model used for valuing American options) have expanded

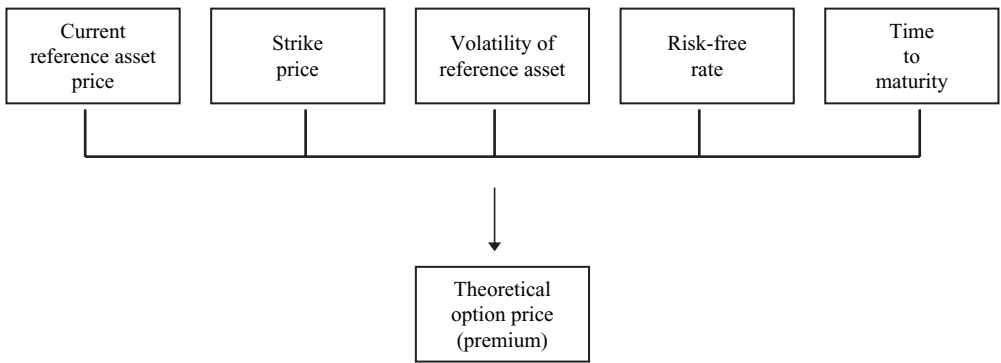


Figure 2.17 Option-pricing inputs

the industry's valuation choices. In basic terms, the Black–Scholes model builds on the general minimum value of an option (market price (S) – strike price (X) for a call, $X - S$ for a put) by considering both the present value of a potential payoff and the probability of exercise. The underlying asset price is assumed to move up and down in very small increments (i.e. no price gaps or jumps). Since the option is characterized by a future exercise/expiry date, the present value of the potential payoff is discounted on a continuous basis; this yields a key component of the model. The next component relates to the probability that the price of the asset will rise to a level where the option becomes exercisable (within the allocated timeframe) – this is simply the probability that the option will pay off, and is described by a value from the cumulative normal distribution function (i.e. $N(d_1)$). The remaining component is simply the probability that the option will be exercised, and is again generated as a value from the cumulative normal distribution function (i.e. $N(d_2)$). As noted below, d_1 and d_2 are both determined by the inputs illustrated in Figure 2.17 above. Assembling these components, the value of a European call (c) is given as:

$$c = SN(d_1) - Xe^{-rt}N(d_2)$$

where

S is the asset price

X is the strike price

t is the time to maturity

r is the risk-free rate

$$d_1 = \frac{\log\left(\frac{S}{X}\right) + \left(r + \frac{\sigma^2}{2}\right)t}{\sigma\sqrt{t}}$$

and

$$d_2 = d_1 - \sigma\sqrt{t}$$

Valuation models and arbitrage forces suggest that the value of a call option less the value of a put option is equal to the present value of the asset price less the strike price. This concept,

known as put–call parity, means that the purchase of a European call option and the sale of a European put option with the same strike and expiry generates a leveraged investment in the underlying asset (less the value of cash payments on the underlying asset over the life of the option).¹⁸ The value of a put option (p) under put–call parity can be shown as:

$$p = c + Xe^{-rt} - S$$

The mathematics of the framework are obviously a great deal more complex, but the logic is intuitively appealing. As a result of these relationships, it is relatively easy to note that the higher the value of the risk inputs, the higher the value of the option. And, while all of the inputs play an important role, volatility of the asset price remains a central input parameter. When the volatility can be measured with a degree of precision, the resulting fair value of the option is accurate.

Exchange-traded derivatives

The broad class of exchange-traded derivatives is the second major segment of the overall derivatives market, and can be used to achieve many of the same goals as OTC contracts (often at a cheaper cost, but with less structural flexibility). Exchange-traded, or listed, derivative contracts predate OTC derivatives by many decades; in fact, certain listed contracts on agricultural commodities have existed in some markets for several hundred years. However, the modern era of listed derivatives is traced more readily to the early 1970s, when the first listed financial contracts were introduced on select interest rates and equities. Since that time, growth and innovation have expanded rapidly, and the most successful instruments have built significant volume and liquidity. In this section we focus our attention primarily on the listed futures contract; though many exchanges also list options and futures options, these contracts are similar to the OTC option contracts described above, and we shall not therefore replicate the discussion. Before reviewing the essentials of futures contracts, it is important to describe certain characteristics of the listed market that distinguish it from its OTC counterpart; these differences relate to the trading forum/mechanism, credit risk exposure, contract standardization/liquidity, and price limits.

All trading activity in the listed market occurs through regulated exchanges. An exchange, which may be publicly, privately, or mutually owned, is counterparty to every trade (either directly or via its clearinghouse); trading may be conducted via open outcry or electronically (though, in practice, most platforms have migrated to an electronic environment). Thus, the buyer of a listed contract faces the exchange as her selling counterparty, while the seller faces the exchange as her buying counterparty. Each buyer or seller dealing with the exchange is required to post initial margin (or security) with the clearinghouse to protect against possible credit losses. The initial margin is evaluated by the clearinghouse every day based on a client's open positions, and additional calls for margin (i.e. variation margin) may be made if a particular level (i.e. maintenance margin) is breached as a result of losses. Since the exchange serves as counterparty, and participants must post margin, the counterparty credit risk that characterizes many OTC derivative contracts is eliminated; this is one of the central advantages of the listed marketplace.

¹⁸ Knowing this, it is possible to use options in different combinations to create a range of synthetic structures; we shall explore the concept of synthetic derivative positions in greater detail in Chapter 10.

All contracts listed on the exchange are standardized. With rare exception granted to certain “flex” options contracts (which allow participants to define strike, exercise style, and maturity within certain parameters), all listed contracts are governed by standardized terms. This means trade details cannot be customized, as in the OTC market; all buyers and sellers dealing through the exchange trade in precisely the same quantity, tick size, maturity, delivery grade, and settlement style. The advantage of this type of standardization is, of course, liquidity. Since all parties trade the same instrument, a critical mass of liquidity can build for each contract month. This is especially beneficial for widely used references in the interest rate, equity, currency, and commodity markets, where hedgers and speculators are attracted by the tight dealing spreads that characterize very liquid contracts. Maximum daily price movements also govern many listed contracts. This means that if buying or selling pressures become too extreme, the exchanges can impose price bands lasting for a period of minutes or hours (or even an entire trading session); they may also increase margin requirements to ensure sufficient security. No such price limits exist in the OTC market.

Futures

The futures contract, which is the fundamental instrument of the listed market, is a bilateral contract between two parties that allows one party, the seller, to sell a particular reference asset at a forward price for settlement at a future date, and the second party, the buyer, to purchase the reference asset at the forward price on the named date. Based on this definition, the futures contract appears to be identical to the forward contract discussed earlier; in fact, the payoff profiles of the two instruments are precisely the same. However, as a result of the exchange mechanism, the futures position calls for a daily settlement of cash flows, rather than a single settlement of cash flows at maturity: each futures position is revalued by the clearinghouse at the close of each trading day, and a net settlement is made in favor of the party holding the day’s gain. Thus, if the market price at the end of the day is greater than the price at the start of the day, the buyer posts a profit and the seller a loss (and vice versa). Like forwards, futures can be settled in physical or financial terms; the precise settlement mechanism is specified as part of a contract’s standardized features. Contract maturities vary, but typically span one month to one quarter. The most active futures references may be listed simultaneously on a monthly or quarterly cycle out to several years (e.g. Eurodollar futures, S&P 500 futures, Brent crude futures).

Options and futures options

As noted, the listed market features options that function much like the OTC options described earlier in the chapter. In fact, there is no difference between the two contracts apart from the general characteristics that distinguish listed products from OTC products (i.e. listed options come with a set number and range of strike prices and maturities, they are margined every day through the clearinghouse, and so forth). A unique subset of the options market centers on futures options, which are simply options that give the buyer the right to buy or sell an underlying futures contract, and which require the seller to accept or deliver the futures contract upon exercise. The futures option can therefore be viewed as a structured derivative comprised of an option and a futures contract. The actual positions of the long and short futures options positions are summarized in Table 2.3.

Table 2.3 Futures options positions

Position	If exercised
Long futures call	Buyer acquires futures contract at strike price
Long futures put	Buyer sells futures contract at strike price
Short futures call	Seller obligated to sell futures contract at strike price
Short futures put	Seller obligated to acquire futures contract at strike price

2.4 HOST SECURITIES/LIABILITIES

Host securities/liabilities are another essential component of the structured and synthetic asset market. As we shall note in many of the chapters that follow, securities play an integral role in conveying the economic risks and returns that intermediaries and end-users seek. Though some synthetic instruments are created on a standalone basis (i.e. certain derivative structures, as described in Chapter 10), most are formed using a host debt security. For simplicity, we include in our discussion tradable/transferable loans; these are not, strictly speaking, securities, but have some of the marketability features of securities, and therefore form part of our discussion. Host securities are issued directly by companies and other corporate/sovereign/supranational issuers, either through public notes/shelf registrations, private placements, or transferable loans (equity rights offerings, or secondary equity placements, are not used as host vehicles; equity-linked securities, such as those discussed in Chapter 7, are combinations of debt securities and equity derivatives). Note that other securities are issued via specialized issuing/repackaging vehicles, which we consider at greater length in the section below.

2.4.1 Public notes/shelf registrations

Public note programs, including those created as shelf registrations, have terms ranging from less than 30 days to over 30 years. They represent the senior, unsecured liabilities of an issuer, placed with investors through underwriters/intermediaries or on a direct basis. The standard note program, which can be structured in the form of commercial paper (CP), Eurocommercial paper (ECP), medium term notes (MTNs), or Euro/global MTNs (EMTNs), allows an issuer to register the facility once every few years with the relevant financial regulator and then float tranches of notes as needed. This stands in contrast to standard registered notes and bonds, which must be scrutinized individually by regulators, a process that creates additional issuance costs and timing delays. In the US, public note programs were popularized in 1982, when the Securities and Exchange Commission (SEC) passed Rule 415, which permitted the use of the shelf registration structure; issuers need only register an issue once every two years, and can float securities at will during that period, with only minor updates to the disclosure with each tranche flotation. The benefit of this process is, of course, that individual tranches can be floated to accommodate market-funding opportunities, as well as the specific demands of investors. Not surprisingly, public note programs have become a popular means of creating specific synthetic and structured assets as they can accept embedded derivatives readily.

2.4.2 Private placements

Private placements are securities that are exempt from registration with a national securities regulator, meaning that the specific disclosure and listing requirements that characterize public

issues need not be followed. Issuers choosing to raise capital through a private placement generally must lodge an information memorandum with investors, providing detail on financial condition and terms of the deal; this information is, however, less detailed than that required for a standard public transaction. The limited disclosure and lack of registration mean that private placements can only be distributed to a small number of sophisticated institutional investors (qualified institutional buyers, or QIBs). These investors must have a minimum level of net worth, and must be able to absorb a greater amount of risk than typical investors in a registered deal. In addition, and unlike public transactions, resalability is very limited. Most private placements are designed as “hold to maturity” transactions, with no possibility of selling the securities on a secondary basis. In some instances, private placements can qualify for limited resalability. In the US, this occurs under Rule 144A, which permits an unregistered security to be traded on a selective basis between the original purchaser, the underwriter/dealer, and other recognized QIBs. Though this rule has injected some liquidity into the market, the population of QIBs authorized on a given deal is generally quite small, so deals must still be considered relatively illiquid. Despite this feature, structured assets are arranged routinely through private placements.

2.4.3 Transferable loans

Banks historically have arranged loans with the intent of keeping them on their books until the underlying borrowers repay them. This mechanism remains prevalent to the present time, suggesting that the lending bank is the lender of record on a given deal, and must seek specific approval from the borrower if it intends to sell the exposure on a secondary basis to another party. There are, of course, well-accepted exceptions to this process: when a bank arranges a new credit facility for a borrower, it may create a syndicate of associated banks to take a share of the loan; each syndicate bank thus becomes a lender to the borrower, though the lead bank may retain the agency function associated with collecting and distributing principal and interest payments. In other instances, a loan can be transferred to other banks or investors on a secondary basis. Such transferable loans, which do not require the specific approval of the original borrower, add a dimension of liquidity to the traditional loan product, and make them more suitable for a range of structured asset transactions. It is important to note that transferable loans generally require notification to the borrower that a portion of the loan has been sold to another party, but do not typically require the borrower’s prior consent to transfer.

2.5 ISSUING/REPACKAGING VEHICLES

Our discussion in this text will touch on many examples where cash flows are altered by placing assets, liabilities, or other cash flows into an entity that then issues its own liabilities to investors. This is particularly, though not exclusively, true for structured assets such as mortgage- and asset-backed securities, collateralized debt obligations, structured notes, and certain types of investment funds. Accordingly, a brief overview of certain common issuance or repackaging vehicles is useful; the major classes we consider include special purpose entities, trusts, and investment companies. In each case, the entities may be established on-shore (e.g. in a particular domestic market) or offshore (e.g. outside national borders); in fact, offshore vehicles dominate the market, as they can convey important tax and regulatory advantages to sponsors and/or investors. For instance, structured notes issued from an offshore vehicle may be free of withholding taxes, and may not have to adhere to financial disclosure




Figure 2.18 Program issuer SPE

requirements – these are powerful advantages that can lower friction costs and improve transaction economics; they may also, however, add a certain degree of opacity to the financial engineering process.

2.5.1 Special purpose entities

Special purpose entities (SPEs), also known broadly as variable interest entities (VIEs), are the single most popular vehicle for structured asset activities, primarily because they can be established with relative ease and at low cost, they can be domiciled in tax- and regulatory-friendly jurisdictions with a strong legal framework¹⁹ (e.g. The Netherlands Antilles, Jersey, Cayman Islands), and can accommodate a range of structures that range from basic to highly complex. Since SPEs reside off balance sheet (e.g. sponsors hold a nominal amount of equity – enough to transfer risk, but not enough to consolidate on the balance sheet), their actual activities with, and relationships to, sponsors are often unclear, which adds a degree of regulatory concern.²⁰ Nevertheless, their use in the structured asset sector remains popular.

SPEs can be created as single- or multiple-purpose entities. Single purpose SPEs, as the name suggests, are discrete entities that are created for each individual structured transaction. Multiple-purpose entities, in contrast, are designed to accommodate multiple issuance or dealings; while this involves additional legal structuring (e.g. ring-fencing, or isolation, of assets and liabilities within the overall SPE to ensure no crossover between transactions), it leads to quicker and cheaper deal execution. Thus, after funding the initial fixed costs to establish the entity, the SPE must only bear small incremental variable costs for each new deal.

Multiple-purpose entities, which have become a standard in the market in recent years, can be created as program issuers or multiple issuers. A program issuer is a single legal SPE that is authorized to float multiple series of liabilities, each backed by its own asset/derivative/cash flow base. Investor recourse within the program is limited through nonrecourse agreements to specific assets associated with a particular issue – there is no ability for investors to attach other nonrelated assets. Figure 2.18 illustrates the program issuer SPE.

The critical factor for program issuers is ensuring that the legal structure is created properly; if it is not, then a default under one part of the program may cause investors to seek claims

¹⁹ In order to give investors comfort, the legal framework must be undoubted; this means applying well-accepted standards related to contract law and bankruptcy, among others.
²⁰ Corporate scandals, such as those involving Enron, demonstrated how SPEs can be used to mislead investors and regulators about inter- and intra-company dealings. Though certain regulations have been put in place to increase transparency, these entities still maintain a considerable amount of flexibility and discretion in the way they conduct their activities.

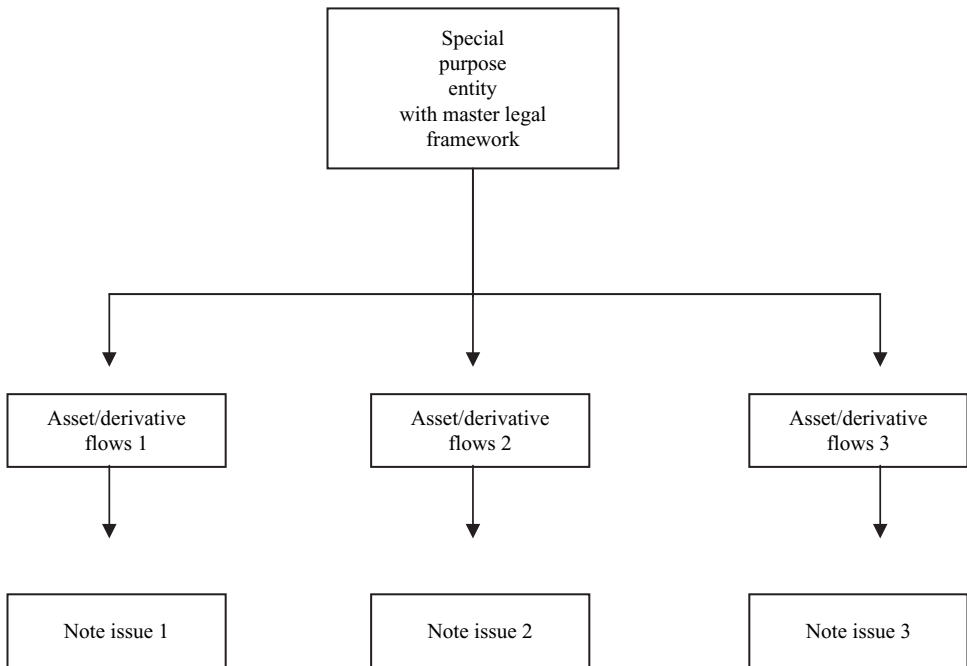


Figure 2.19 Multiple-issuer SPE

from other, unrelated, parts of the program, leading to structural collapse of the SPE. In fact, program issuer SPEs are not used in jurisdictions where segregation is doubtful.

Multiple-issuer SPEs use a master legal/documentation structure along with separate legal “subvehicles” established for each issuance, as illustrated in Figure 2.19. As a result of this arrangement, asset commingling and breaches of the nonrecourse tenet cannot occur, suggesting greater legal comfort for all parties. However, the additional layer of legal complexity in the framework leads to higher costs. Multiple-issuer SPEs generally are favored when high-risk assets are involved (e.g. those with a high probability of default, where cash flow disruption will ensue).

Structured liabilities issued by any SPE take the form of securities (rather than trust receipts, as discussed below); the securities may be publicly registered or privately placed, and are often individually rated by the rating agencies.

2.5.2 Trusts

Trusts, like SPEs, perform many of the structuring and issuance functions we describe throughout the book. From a legal perspective, trusts are not owned by the sponsoring institution, but by a charitable foundation established expressly for the purpose; an independent trustee is appointed to oversee the trust’s operations and ensure property management for the benefit of investors. Trusts, used for asset consolidation and restructuring, issue liabilities known as trust receipts, which represent a claim on specific assets or cash flows held within the trust. Trusts are used commonly in certain asset-backed structures, such as credit card receivables

transactions (e.g. the master trust structure), and may also be used with various types of credit- and currency-linked issues and investment funds (e.g. exchange-traded funds).

Trustees play a key role in many structured asset transactions (though trustees are appointed on every trust-based transaction, they also play a role in SPE and investment company/partnership transactions, in order to inject an additional layer of security and independence; in fact, the trustee represents the essential link between issuers and investors, holding and exercising the legal rights of investors). We distinguish, under common law, between legal ownership (a party holding title to certain assets and recognized as the owner) and beneficial ownership (a party that does not hold title, but gains the benefits associated with legal ownership); thus, under common law, a party can legally own property for “holding purposes” only – to the benefit of third parties (e.g. investors). The rights between the trustee (as legal owner) and investors (as beneficial owners) are communicated via the trust deed. The terms of the trust itself are agreed between the issuer and the trustee; investors have no input, though terms must be in their favor or else they will have no incentive to participate. Placing a trustee between the issuer and investors eliminates problems associated with perfecting security interests, as interests do not then have to be renegotiated with each secondary sale; the trustee effectively negotiates rights for the benefit of any future investors.²¹ The trustee also works in close partnership with other service providers, which may include a listing agent, servicer, administrator, paying agent and/or custody agent²² (in some instances, the trustee also serves as the custody agent).

2.5.3 Investment companies and partnerships

Investment funds, including those we consider in Chapter 9, can be established as trusts (as noted above), or as investment companies or investment partnerships. In practice, the investment company structure tends to be most popular. The investment company framework requires registration with the relevant national financial regulator, and adherence to minimum disclosure standards and regulations (e.g. strict accounts and records must be maintained, and all customer assets must be held in custody). As a result of the regulatory oversight, investment companies generally are allowed to make public offerings of securities to retail and institutional investors, and are allowed to market their services and products to the public at large.

The partnership, comprised of general and limited partners, is an alternative structure. General partners typically are the fund managers of the company, who bear the risk of unlimited liability characteristic of most other types of partnership; in practice, general partners often place their interests into limited liability companies in order to protect against unlimited claims. The limited partners, in turn, generally consist of the investors committing capital to a transaction. As the name suggests, their risk is limited to the capital they have invested in a particular venture, and general partners have no ability to make additional capital calls on them in the event of losses. The actual conveyance of interest is done in the form of partnership interests, rather than securities; not surprisingly, these generally are rather illiquid, and in many cases can only be resold on a restricted basis to a small number of investors. Dealings

²¹ The trustee also serves as the contact between an issuer and what may be dozens, or hundreds, of anonymous investors (e.g. bearer, rather than registered, security holders). This is useful in the event of default or restructuring negotiations. Here, we must also distinguish between a note trustee, who is responsible for noteholders or trust receipt holders only, and a security trustee, who is responsible for both noteholders/trust receipt holders and secured creditors.

²² The custody agent is responsible for settlements, holding physical securities/receipts, providing corporate action information, executing coupon payments, collecting principal and interest, and so forth.

are governed by partnership agreements, which detail the rights and liabilities of each of the parties.

2.6 FINANCIAL ENGINEERING AND PRODUCT DESIGN

We have introduced various essential building blocks in the preceding sections, but before embarking on the product discussion that follows in the next chapters, we make a few concluding remarks about the process of financial engineering.

Financial engineering relates to the creation of new products that are useful to intermediaries and end-users. It involves identifying a particular need, determining how best to address the need, assembling the necessary building blocks, and delivering the finished product to the user. To be sure, the process is often very complicated; some end-user requirements can only be resolved by using complex instruments and cash flows. However, it is important to remember that any asset, liability, or off-balance-sheet contract can be divided ultimately into a discrete set of risk drivers and associated cash flows – these will always form the foundation of any new synthetic or structured asset. We know, for instance, that a bond is a package consisting of interest rate risk, default risk, liquidity risk, and possibly also convexity risk and currency risk. Knowing this, and the amount of the cash flows involved in the payment of interest and principal, we can construct, reshape, or deconstruct any type of asset around this bond. The same is true for all other assets. Financial engineering simply formalizes the process.

Callable, Puttable, and Stripped Securities

3.1 INTRODUCTION

The first structured assets we consider, callable bonds, puttable bonds, and stripped (zero coupon) securities, are among the most established of the financial markets. Transforming the maturity horizons and/or principal and interest flows payable to investors has proven to be appealing to issuers, intermediaries, and investors for decades, and has led to significant and continuous activity in all parts of the sector. Indeed, issuance, purchase, and trading of these types of structured assets is likely to remain robust in the future, as many of the fundamental motivations that drive activity are genuine and unchanging.

In this chapter we consider the development and mechanics of callable bonds, puttable bonds, and interest only (IO) and principal only (PO) tranches comprising stripped securities; these are summarized in Figure 3.1. For ease of presentation, our focus is on the application of callability, puttability, and stripping/reconstitution to generic government or corporate bonds. We stress, however, that the same features and technologies can be applied to other types of securities, including convertible bonds, structured notes, and so forth.

3.2 DEVELOPMENT AND MARKET DRIVERS

Callable and puttable bonds have existed for several decades, with the first issues appearing in the early 20th century, and true issuer and investor interest commencing in the late 1960s. Rapid growth of the Eurobond market during the 1960s and into the 1970s was particularly instrumental in promoting the callable security structure. During these early days of international capital market development issuers were becoming increasingly attuned to the funding levels they could achieve in different market cycles; the ability to call outstanding bonds as rates declined proved particularly appealing, and the frequency of callable bond issuance soon increased. The structure soon spread to other markets, and domestic issuers in the US, Australia, Canada, and Japan, among others, began floating bonds on the same basis. Puttable structures were, and still are, rather less common than their callable counterparts. Such securities appeared several years after the first callable bonds were floated, and achieved a critical mass of interest among institutional investors seeking to manage their investment and reinvestment risks more actively; they also drew interest from issuers interested in lowering their funding costs via the embedded option. While corporates originally floated callable and puttable securities, financial institutions, sovereigns, and supnationals soon joined them.

The informal market for US strips dates back to the late 1970s, when institutions traded government securities without one or more coupons by physically stripping the bearer securities. This practice was done primarily to take advantage of generous tax advantages, which were reduced in 1982 through the passage of legislation. In fact, the same legislation required the US Treasury to begin issuing its securities in electronic book-entry form, eliminating the

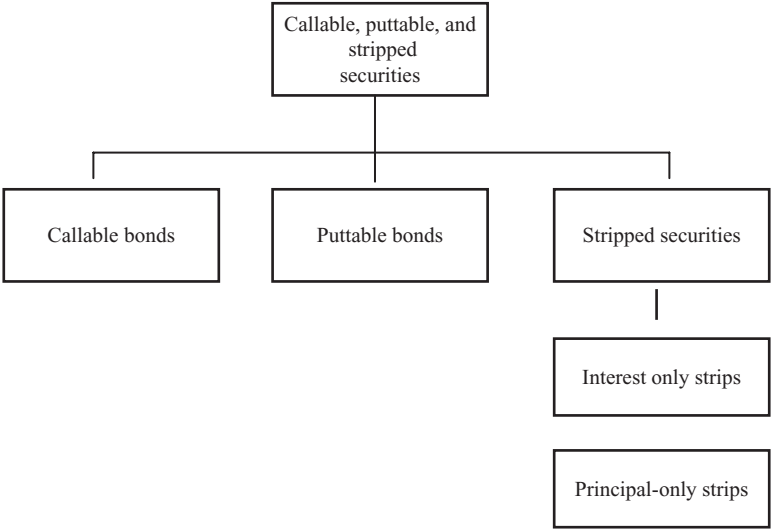


Figure 3.1 Callable bonds, puttable bonds, and stripped securities

practice of physical stripping and ushering in the era of book-entry, or electronic, securities stripping. Shortly thereafter, financial institutions began splitting bond cash flows into separate interest and principal components through book-entry trust receipts. Merrill Lynch and Salomon Brothers were the first to introduce US government bond strip products (TIGRs, Treasury Income Growth Receipts, and CATS, Certificates of Accrual on Treasury Securities, respectively), floating separate trust receipts for each coupon plus the principal. Various other intermediaries followed suit, and the market began to grow reasonably quickly. However, since the proprietary products were not fungible, a critical mass of liquidity failed to accumulate. Accordingly, major intermediaries began stripping and trading “generic” securities with even coupon dates alongside their proprietary products. In 1985, the US Treasury announced the launch of its own program (STRIPs, Separate Trading of Registered Interest and Principal), though it opted not to do any of the stripping itself – leaving that task for the major dealers. In 1987, the Treasury also permitted the practice of reconstituting stripped securities. The STRIP program now dominates activity in the US.

The success of the US program during the 1980s led public and private institutions to sponsor similar efforts in other local government securities markets (e.g. France, UK, Germany, Canada), and in other asset classes (e.g. mortgage-backed securities strips, as described in Chapter 4). For instance, in 1991, the French Treasury became the first sovereign issuer in Europe to authorize stripping and reconstituting of its securities (the Obligations Assuables de Tresor, OAT), via the community of primary dealers (SVTs, who act as strip agents and market-makers). The program has proven very successful, and OAT strips now comprise the single largest government strip program in the Euro-zone. In most countries, issuers of bonds being stripped are involved in the process, either directly or via fiscal agents. In fact, government support is ultimately important if a proper base of liquidity is to develop. Some governments have opted to place limits on the types of instrument that are eligible for stripping, in order to concentrate activity in a relatively narrow group of securities – which helps build a critical mass.

Activity in core structured securities is driven by several key factors:

- Callable/puttable bonds:
 - allow issuers to take advantage of future rate declines to lower funding costs (callable securities), and permit investors to earn an incremental yield in the interim;
 - permit investors to take advantage of future rate increases to boost yield (puttable securities), and allow issuers to lower all-in funding costs in the interim;
- create funding arbitrage opportunities through the monetization of optionality.
- Stripped (zero coupon) securities:
 - permit investors to acquire specific cash flows from a given bond; those investing in high quality government strips face minimal risk of default;
 - create opportunities to match duration and convexity requirements very precisely, allowing for more effective risk and investment management;
 - require a smaller initial capital investment, since all instruments are sold on a discount basis; this can be significant for those trading at a deep discount (e.g. long-dated principal only strips);
 - give investors a certain and fixed redemption value if securities are held to maturity;
 - generate tax advantages for investors that can place the securities in tax deferred accounts.

Investors in callable and puttable bonds include most major institutional parties, including mutual and pension funds, hedge funds, bank intermediaries, and corporations. The option-embedded benefits have become a useful risk management tool for some parties, and an effective way of increasing all-in returns for other parties.

Strip investors tend to be sophisticated institutional investors that use the products to create specific portfolio risk management profiles. The additional convexity and duration that can be obtained from certain securities is very attractive for those running diversified and/or complex portfolios. Small denomination strips are marketed regularly to a base of retail investors (though generally as “hold until maturity” products).

3.3 PRODUCT MECHANICS AND APPLICATIONS

The first two products we consider in this section, callable and puttable bonds, are combinations of bonds and options that are issued directly by companies and sovereigns. The third class, stripped securities, represents structured cash flow interests that can be issued by sovereigns directly (e.g. government strips), or by intermediaries seeking to replicate the same economics. Pure corporate issuance of zero coupon securities also exists (e.g. a bond floated without interest coupons, rather than a bond manufactured by removing the coupons).

3.3.1 Callable bonds

A callable bond is a package of a bond and a long issuer call option on bond prices (or long issuer put option on interest rates) that allows the issuer to call the security at particular price and/or time intervals. In fact, we may distinguish between an American callable bond (which is callable at any time before maturity), a European callable bond (which is callable only once, generally one or two periods before final maturity), and a Bermudan callable bond (which is callable on regular, though not continuous, intervals, such as every six or twelve months during the life of the bond). US issuers tend to issue callables primarily in American form, while European and Japanese issuers often opt for the European or Bermudan form.

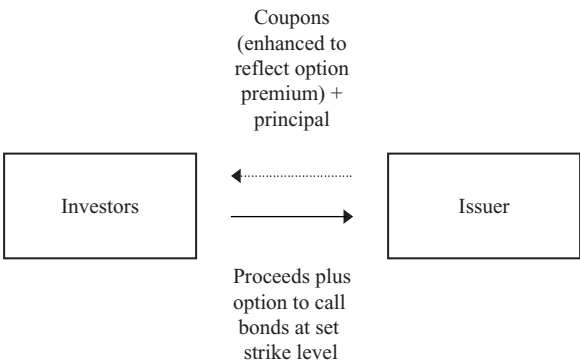


Figure 3.2 Callable bond

In exchange for granting the issuer the call option, the investor collects a premium in the form of an enhanced coupon. If the issuer chooses to exercise the call at the call strike price implied in the security, the investor delivers the bond and receives a principal repayment amount equal to the price-based strike, times the quantity of bonds. For instance, a bond may be issued at par, in denominations of \$10 000/bond, with a call price (strike) of 105. If the price of the bond rises to 107, as an example, and the issuer elects to call, investors will present their bonds for redemption at a price of \$10 500 (and not \$10 700, which represents the market value of the security). Figure 3.2 illustrates the basic callable bond structure.

In practice, the callability of a security is primarily a function of interest rates. Issuers floating callable bonds will choose to call their outstanding debt when rates are declining and opportunities to refinance in the new lower rate environment expand. Knowing from Chapter 2 that rates and prices on fixed income securities are inversely related, this concept is consistent with our simple example above: as rates fall, the outstanding bonds become more valuable (e.g. prices rise). Returning to the payoff diagram we introduced in the last chapter, Figure 3.3 examines the embedded call option from the issuer’s perspective. Assume that at a current rate of 4 %, the bond is worth par; as rates fall to 3.5 %, the price rises to the call price (strike) of 105. At any point after this, the issuer may find it advantageous to call the bonds and refinance in the market (i.e. at a level of 3.5 % or less). While interest rates dominate the option exercise decision, it is worth remembering that credit spreads can also play a role in the process. Since an issuer’s all-in funding cost (as reflected in the coupon) is a combination of the risk-free interest rate level and the relevant credit spread, it is conceivable that static rates, coupled with a significant tightening of the issuer’s credit spread, may also lead to exercise of the call. If the issuer’s credit has improved dramatically in market terms, it will be able to call the security and reissue new debt at the tighter market spread, so reducing its funding costs.

Valuation

We know that a callable bond is simply a combination of a noncallable bond and a long issuer call option, suggesting that from a valuation perspective, we can examine the price of an identical noncallable bond and deduct the price of the call option to determine the theoretical value. Once the theoretical value of the option is determined, the equation can be inverted to see if the noncallable bond price is accurate (e.g. noncallable bond price = callable bond




Figure 3.3 Long issuer call option in callable bond

price + call option price). This exercise is equal to deriving the implied noncallable bond price through observed callable bond prices and theoretical option values.

The value of a callable bond is based, in part, on the value of the option, so we know from the last chapter that option volatility must be one of the pricing inputs. However, determining the callable bond price requires calculating both the noncallable bond price and the call option price, meaning it is difficult to shift from a quoted price on a bond to a yield for a given level of volatility – both variables change as yield changes. Furthermore, since the forward price must be computed, the price of the option depends not on a single yield, but on the entire yield curve. As a result, the market has turned towards the use of option-adjusted yields based on bond price and volatility. The noncallable bond is said to have a fair price when the option-adjusted yield for the callable bond is equal to the yield for a noncallable bond with the same features.¹

We know that bond value is the PV of all future cash flows (principal and interest) and, if the bond is riskless, it is the present value of the replicating portfolio of the risk-free benchmark. The replicating portfolio is valued at the risk-free zero coupon value. If the bond is risky, some spread must be introduced to compensate for the risky nature of the cash flows. The resulting static spread is the spread that makes the PV of the cash flows from the bond, when discounted via the risk-free zero coupon curve, equal to the bond's price.

Unfortunately, this static spread fails to take account of future interest rate volatility that could affect cash flows on a callable (or puttable) bond: the greater the interest rate volatility, the greater the likelihood that the call (or put) embedded in the security will be exercised (though the decision is still dependent on other factors, including remaining time to maturity, call (or put) strike, and so forth). The future path of interest rates thus determines option exercise

¹ An investor can compute the yield to call and the yield to maturity, using the differential as a relative valuation measure in order to determine the relative attractiveness of a callable bond. The investor may also compute the yield to each discrete call date; this, along with the yield to maturity, generates the lowest possible yield, or yield to worst.

and, by extension, potential value. Under the option-adjusted yield framework, we can develop a zero coupon curve and assumed spread for each possible interest rate path. The average of all PVs can then be computed; if the average PV is equal to the market price of the bond, then the spread added to the zero coupon rates equals the option-adjusted spread (OAS); if it is not equal, a new path is computed. An OAS can therefore be interpreted as the average spread over the risk-free zero coupon curve based on future interest rate paths. Its value depends critically on assumptions about interest rate volatilities, which are often determined through a simulation process. The procedure is more computationally intensive than standard static spread valuation, but generates a more accurate valuation as it assigns proper value to the embedded optionality.² That said, OAS relies on certain assumptions of its own: it assumes that the embedded bond is held until the effective maturity date; that cash flows are reinvested at a yield equal to the yield of a noncallable bond plus the OAS; and that interest rate volatility estimates are accurate and option exercise behavior by the issuer (or investor, for puttables) is rational. These assumptions are, however, quite reasonable.

It is possible to compute associated sensitivities, including duration and convexity, once OAS relationships are developed. As indicated in Chapter 2, these provide estimates of how much a bond's value gains or loses for small/large changes in yield. For instance, a standard duration formula can be recalibrated to take account of the price relationship between the callable bond and its noncallable equivalent, and the delta (price sensitivity) of the embedded call option. Option-adjusted duration can be given as:

$$OAD = \frac{P_{NCB}}{P_{CB}} * Dur_{NCB} * (1 - \Delta_{call})$$

where

P_{NCB} is the price of the noncallable bond

P_{CB} is the price of the callable bond

Dur_{NCB} is the duration of the noncallable bond

Δ_{call} is the delta of the call option (i.e. price sensitivity to small rate moves).

Similarly, option-adjusted convexity can be computed to determine the degree of price performance, or underperformance, for a large movement in rates, by extending and adjusting the standard convexity calculation:

$$OAC = \frac{P_{NCB}}{P_{CB}} * [Cvx_{NCB} * (1 - \Delta_{call}) - P_{NCB} * \Gamma_{call} * (Dur_{NCB})^2]$$

where

Cvx_{NCB} is the convexity of the noncallable bond

Γ_{call} is the gamma of the call (e.g. price sensitivity to large rate movements)

all other terms are as defined above.

In fact, convexity is an important dimension of the price performance of callable bonds, particularly for price movements in a declining-rate environment (before the call strike is reached). In practice, the actual amount of upward price movement in the bond tends to be

² Note that we shall revisit the OAS framework in our discussion of mortgage-backed securities in Chapter 4. The same framework can also be applied to other option-embedded securities, including capped floaters, payment-in-kind bonds, and so forth.




Figure 3.4 Callable and noncallable bonds and price compression

limited given the potential exercise of the call; this phenomenon, known as price compression or negative convexity, means that while the price of a noncallable bond may continue to rise as rates fall, the price of the callable bond will tend to lag. In the extreme, a callable bond that is currently callable (e.g. one where the strike has been passed) but remains outstanding, will feature significant compression, as the probability of the security being called at any moment is very high; this compression is, of course, a manifestation of negative convexity. We can examine various yield scenarios to understand the impact of price compression on a noncallable bond, a callable bond that has been called, and a callable bond that remains outstanding. Figure 3.4 reflects the fact that as yields decline, the probability of call exercise increases, and price compression begins to set in; however, when yields are positioned above the strike, the callable and noncallable bonds feature similar price/yield movements.

Investors in callable bonds are exposed to reinvestment risk. In granting the call option and receiving enhanced coupons, they may find themselves in a situation where they must reinvest their capital in a lower-rate environment. In the example immediately above, the 4 % yield is attractive when rates decline below 4 %; once they reach 3.5 % and the issuer calls the bonds, investors must reinvest their capital in the new rate environment – which is now 50 bps lower than before, for assets with an equivalent term and credit quality.

In practice, callable bonds may feature call prices (strikes) that change over time. For instance, in order not to dissuade investors from committing capital when rates are declining, an issuer may establish a relatively high call price for the first few years of a multi-year issue, helping ensure that the bonds can be placed successfully. Only after the passage of several years might the call price ratchet down to a lower level, increasing the likelihood of callability. The reverse may also occur: a step-up callable bond may have increasingly higher strikes to induce the issuer to call sooner, rather than later, in the life of the bond. Naturally, an issuer is under no obligation to call securities, even if it appears optimal to do so (i.e. even if the option is in-the-money). In fact, there are instances where issuers may prefer to preserve a particular type of financing. This may occur when the issuer wants to keep its investor base intact, or when it believes that a particular form of capital (e.g. maturity, market, coupon) may

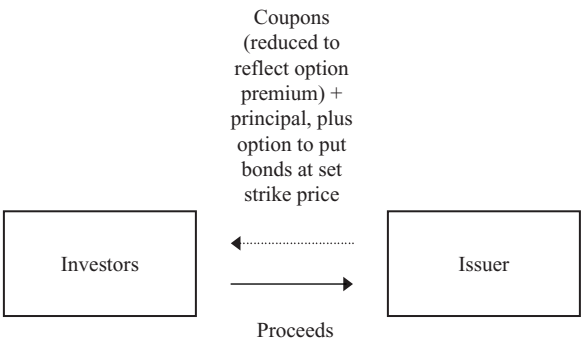


Figure 3.5 Puttable bond

be difficult to refinance in a given environment).³ In general, however, an investor in callable securities must be prepared for a call event.

3.3.2 Puttable bonds

A puttable bond, which is somewhat less common than a callable bond, is a package of a bond and a long investor put option on bond prices (long investor call option on interest rates) that allows the investor to put the security back to the issuer at particular price or time intervals. In exchange for granting investors the put option, the issuer collects a premium in the form of a reduced coupon. If investors decide to exercise the put at the strike price embedded in the security, they deliver the bonds and receive principal repayment equal to the price-based strike, times the quantity of bonds. For example, a bond may be issued at par, in denominations of \$10 000/bond, with a put price (strike) of 103. If the price of the bond falls to 99, as an example, and investors decide to exercise their puts, they present their bonds for redemption at a price of \$10 300 (and not \$9 900, which represents the market value of the securities). Note that unlike the callable bond, which involves a single exercise decision on the part of the issuer, the puttable bond represents discrete decisions by each individual investor holding an eligible security. Figure 3.5 illustrates the basic puttable bond structure.

We can also view the structure in terms of the standard long put payoff profile, as shown in Figure 3.6. In this instance, we note that as the price of the bond falls (e.g. rates rise), the option becomes more valuable to the investor base, which will be motivated to exercise once the strike has been breached. By exercising the option and delivering bonds back to the issuer, investors crystallize a gain and can then reinvest capital in the new, higher-rate environment. As with the callable bond, the puttable bond may be exercisable by investors if rates are static but the issuer’s credit spread widens beyond the strike; under this scenario, it is advantageous for the investor base to put the securities back to the issuer.

The valuation techniques noted above are readily applicable to puttable bonds, since their construction is very similar. Ultimately, the theoretical puttable bond price is simply the value of the nonputtable bond plus the value of the put option.

³ For instance, during much of the 1990s, Japanese companies issuing callable securities (as well as convertibles) chose not to call their bonds as rates declined for fear of alienating their investor base. Overarching relationship decisions thus dominated pure financial decisions.




Figure 3.6 Long investor put option in a puttable bond

It is important to note that since callable and puttable securities are a common form of financing, they are often swapped into fixed or floating rates in order to provide issuers and/or investors with preferred interest flows. However, since the securities feature optionality and some potential for retirement prior to final maturity, any swap that might be attached to a funding or investment strategy must take account of the fact that early termination is a possibility. In fact, a common strategy is to arrange puttable or callable swaps with swaptions, in order to allow one or the other of the two parties to “break” the transaction. We will discuss the use of these derivatives with callable and puttable structures at greater length in Chapter 10 and will consider various arbitrage techniques used to monetize the value of embedded options.

While callable and puttable securities normally are associated with single issues from the corporate or sovereign bond market, they can be found in other sectors as well. Consider, for instance, the guaranteed insurance contract (GIC), a liability instrument issued by insurance companies in various national systems. GICs, which offer investors a specific guaranteed rate of return, can be supplemented or withdrawn at the investors’ option, without penalty. This means that the instrument essentially is a subordinated fixed income security with embedded puts and calls (ranking junior to other policyholder liabilities). Thus, the option to deposit funds at the guaranteed rate will be exercised when the market rate falls below the guaranteed (i.e. strike) rate; this is precisely equal to a call option on an interest-bearing note. Conversely, the option to withdraw will be exercised when rates rise above the guaranteed (strike) rate, so the investor holds a put. In fact, as we shall note later in the book, this structure is equal to a fixed-income instrument with an embedded straddle (i.e. long call and long put options at the same strike).

3.3.3 Stripped securities

The market for stripped securities, or synthetic zero coupon securities, has become significant, with financial institutions and government treasuries routinely creating new supply to meet

demand from investors seeking interest-only (IO) and principal-only (PO) tranches. Though most profit-driven opportunities are based on stripping, market circumstances also allow intermediaries to reconstitute or recombine previously stripped securities into whole bonds.

Strips markets have gained in popularity over the years, as the underlying securities that are manufactured permit the establishment of very explicit cash flows at set horizons. This creates a great deal more cash flow stability than a standard coupon bond, as there is (a) no reinvestment risk, and (b) a certain time horizon for full par value redemption (unless the underlying security is callable, which is rather rare). The strip market can thus be a useful conduit for institutions that seek to match their assets and liabilities or immunize their portfolios (particularly over very long-term horizons).

Stripping processes

A stripped security is a package comprised of a series of IOs and a PO based on cash flows payable by the issuer of the original security. The issuer's obligation to pay interest and principal on the strips does not change with the deconstruction of the security into its component parts. Most national systems that feature strip programs have commenced their activities with physical stripping of future coupons from bearer securities (e.g. "coupon clipping"). As securities have become dematerialized over the past two decades, physical processes have been suspended in favor of electronic processes (i.e. shifting from physically removing coupons from bearer instruments to electronically decomposing book-entry securities).

Stripping securities is, of course, only half of the process. There are times when it is beneficial to reassemble or reconstitute the package of strips, as this can yield a small arbitrage profit (e.g. buying and reassembling the IO and PO components, and selling them in the market for more than the equivalent security). However, reconstituting securities can occasionally be problematic, as it may be difficult to identify all of the holders of the corpus and coupon strips and then persuade them to sell, at economically reasonable levels, into the reconstitution program. In order to circumvent this problem, some national programs attempt to synchronize coupon dates on strippable securities, or permit "like strips" to be substituted in order to complete a package. In fact, strips that have a degree of fungibility generally are more liquid than those that lack the same characteristics, precisely because they can be used for reconstitution. A bond can be stripped into some minimum denomination that varies by national system, but generally is centered at the equivalent of \$10 000 to \$100 000. This means that retail investors are often able to participate in the market directly. It also means, of course, that firms reconstituting bonds that have been decomposed initially into the principal and coupon components, and from there into smaller denominations of each, face more involved reconstitution tasks.

Let us consider several examples of how the stripping process works. In the US, the Treasury auctions government bonds as part of its normal funding/refinancing cycle. Financial institutions bidding at auction and wanting to decompose the securities they acquire (under the STRIPS program) submit a request to the Treasury to have each coupon payment and principal recoded as a separate security in the book-entry system, thus creating the individual securities required by intermediaries and investors; each security is assigned an individual security identifier number (e.g. CUSIP) that can be used for tracking and reconstituting. Every Treasury security with a maturity of ten years or more is eligible for stripping. For instance, a financial institution may identify a "strippable" security that it wants to decompose into individual zero coupon securities – one for each interest coupon and one for the principal. Thus, if a \$20 m, 20-year, 10 % s.a. Treasury is stripped, it yields a \$20 m face value PO strip maturing in 20 years,

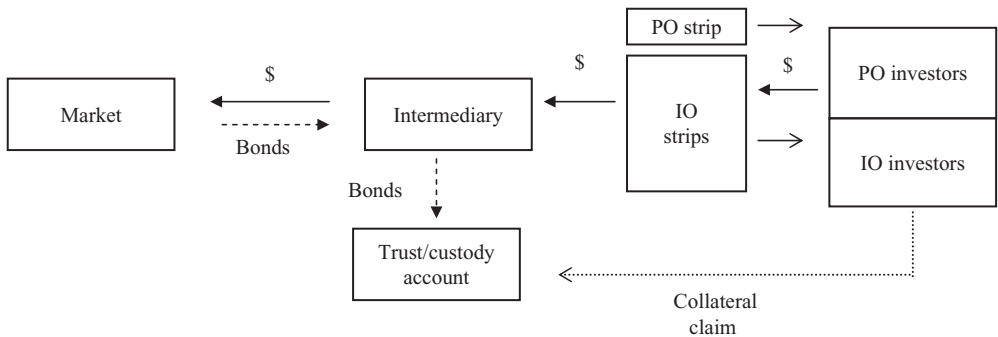


Figure 3.7 Fundamental strip program

and 40 \$ 1m semi-annual IO strips maturing sequentially over the next 20 years. Discounting the PO strip by a 10 % rate over 20 years suggests an initial selling amount of \$2.973 m; this initial investment will grow to \$20m in 20 years. Similar discounting can be performed for each of the individual IO strips; those maturing closest to trade date will trade nearer the \$1 m face value than those maturing in 10 or 20 years.

For securities not forming part of the official STRIPs program, the sponsoring intermediary acquires government bonds, places them into a trust/custody account, and then issues (directly or via an SPE or trust), securities/trust receipts that represent either interest or principal cash flows associated with the government bonds held in trust. These assets, which represent a restructuring of cash flows, are collateralized by the government bonds in the trust account, meaning that investors need not bear the credit risk of the financial intermediary arranging the issue. As each one of the coupon-based IOs is presented for redemption on the appointed coupon date, the intermediary channels the cash coupon payment from the securities in the trust account directly through to the investor, in exchange for the security, which is retired; the same occurs at maturity with the PO. This fundamental structure is reflected in Figure 3.7.

The Canadian program works in a similar manner. The Canadian Depository for Stripping is responsible for the process, and has established procedures for financial institutions and other depositories to strip or reconstitute Canadian government bonds. An eligible depository first creates a book-entry strip bond with a separate security identifier attached to each future coupon and principal flow. The investor can then purchase the position within the depository's ledger account. In some cases, the depository must process the trade by exchanging a position in the underlying security for the strip, crediting one account and debiting another one. For instance, if an investor instructs a depository to strip a C\$100 000 10 % bond with a one year maturity and two C\$5 000 coupons, the underlying bond account is debited for C\$100 000 and credited with the two coupon strips and the corpus strip. The reverse operation occurs in the event of reconstitution. Canada also provides a mechanism for participants to deposit physical strips with the depository; these are considered fungible with book-entry strips created by depositories.

The UK, which has featured a gilt strip program for two decades, operates through processes that are similar to those found in the US and Canada. However, the country also added an additional feature, the gilt stripping facility, in 1995. The facility permits the exchange of coupon gilts for a series of zero coupon strips, or the presentation of constituent strips for reconstituted gilts. The objective of this facility is to create additional market efficiencies that encourage

institutions to participate more actively in the market. The entire process operates using book-entry transfers, so no physical scrip is involved; each potential coupon strip maturing on the same day, along with each unique principal strip, is assigned an individual security identifier (e.g. ISIN and SEDOL). In practice, the coupon-to-strip process can be accomplished directly by any Gilt Edge Market Maker (GEMM, a UK primary dealer), meaning that GEMMs can strip gilts at will, while non-GEMMs must deal through a GEMM. In order to reconstitute a stripped gilt, the principal strip and any outstanding coupon strips must be presented to the facility; once the necessary checks have been performed, a coupon-bearing gilt is released via the GEMM. Requests for stripping or reconstituting can be arranged up to one month in advance.

Government bonds are often traded on a “when issued” (WI) basis; that is, trading commences before a bond issue has been floated, allowing liquidity to build prior to launch. Government bonds that are strippable may trade on a WI basis, meaning that the strips themselves can be stripped and traded on a WI basis. In order to enhance liquidity in the strip product, most advanced national systems allow strips to be used in the repurchase agreement/reverse repurchase agreement markets. This means that positions can be financed (via repos) or excess cash can be invested on a collateralized basis (via reverse repos). Since the price volatility of strips is greater than on equivalent government securities (for reasons we discuss below), the “haircut,” or discount, applied to repo transactions is larger than on coupon-bearing government bonds (e.g., 1–3 % versus $\frac{1}{4}$ – $\frac{1}{2}$ %). Strippable securities are often deliverable under government bond futures contracts, though not in stripped form. That is, any stripped security must first be reconstituted before it can be delivered. Securities that are theoretically strippable but have not yet been stripped are, of course, immediately deliverable.

Supply and valuation

Some countries control the supply of strippable securities, while others prefer a more liberal approach, allowing any security with specific characteristics to be stripped. In some cases, the move from a narrow to broad population of strippable securities is evolutionary, with the government sponsor first wanting to create a critical mass of interest in a select number of issues before permitting expansion into other issues. For instance, the German program, which began in 1997 based on 10 and 30-year Bunds, placed limits on securities eligible for stripping (e.g. two 2007 Bund issues, one 2024 issue, and one 2027 issue were part of the original population of eligible securities). The government’s intent, however, was to expand the population once liquidity built (which it has), but still require eligible issues to feature uniform coupon dates in order to allow for easier reconstitutions (i.e. even coupon dates allow an institution to pair coupons from different issues in order to recreate a Bund).⁴ It is important to note that just because an issue of government securities is strippable, the entire issue need not be stripped; the portion that remains unstripped will continue to trade as standard quoted coupon bonds.

A sufficient supply of bonds is essential in allowing a strip market to develop. In fact, in order to attract institutional interest, large blocks of strips have to be created; this requires a rather significant amount of strippable government securities. Consider, for instance, that a \$5b, 8 % coupon-strippable government bond paying a semi-annual coupon only produces individual IO strips of \$200 m per period (face value), and a much lower amount once the discounted value

⁴ In fact, the German government prohibits reassembling a Bund with “incorrect” components, and attempts to avoid problems by grouping strips with the same maturities under the same security identifier.

is taken into account (especially for the long-dated strips). Though \$5 b would appear to be a significant amount of strippable supply, it is relatively easy to see that, in a market that trades hundreds of millions to billions of dollars per day, the amount is rather small. Strip programs are therefore only suitable for national systems with a significant amount of government debt that is deemed eligible for stripping. Government bond issues that are callable are generally excluded from strip programs, for the obvious reasons associated with disrupting investor cash flow management. Inflation-backed bonds (e.g. index-linked gilts and OATs, Treasury inflation protected securities) may, however, be eligible for stripping, as they can be appealing to investors seeking real, rather than nominal, rates of return.

We know from above that a strip is simply a synthetic zero coupon security that entitles the holder to an interest flow or a principal flow. Armed with this knowledge, it is relatively straightforward to develop a price for a strip. In fact, market convention across national systems is well developed, where the quoted price of the strip is computed as:

$$P = \frac{100}{(1 + \frac{y}{2})^{(\frac{r}{s} + n)}}$$

where

y is the redemption yield

r is the number of days from the settlement issue/date to the next coupon date

s is the number of days in the coupon period in which the settlement date falls

n is the number of remaining coupon periods after the current period.

For instance, if the redemption yield is set at 5 %, the number of days from the settlement date to the next coupon date is 180, the number of days in the coupon period in which the settlement date falls is 360, and the number of remaining coupon periods (after the current period) is five, the fair price of the strip is: 87.30 (i.e. $100/(1.025)^{5.5}$). Though many national markets trade strips on a yield, rather than cash, basis, settlement is arranged in price terms, hence the need to use the formula above for proper conversion.

We noted in the last chapter that financial institutions use zero coupon curves to price their bond and derivative portfolios, and may employ a variety of methods to derive a zero coupon curve from the observed par curve. In fact, a steady supply of actively traded strips along key points of the curve can generate directly a zero coupon curve that may actually be more accurate than a curve derived through bootstrapping procedures. While there should certainly be a close approximation between the actual and theoretical zero coupon curves, some points may not align precisely. This provides institutions with the possibility of arbitrage operations. Figure 3.8 illustrates the minor discrepancies that might occur between actual and theoretical zero coupon curves.

The price sensitivity of IO and PO strips to changes in rates is relatively easy to determine through the duration and convexity tools we introduced in the last chapter. For instance, since strips are zero coupon securities with single cash flows occurring at maturity, the duration is precisely equal to the residual maturity. Thus, while a coupon-bearing security with a final maturity of ten years may have a duration of six or seven, the PO strip will have a duration of ten; convexity is greater as a result. This means, of course, that the PO strip is more price sensitive (i.e. riskier) than the coupon-bearing security. The same is true of IOs versus coupon-bearing bonds of the same maturity. This feature, coupled with the de facto leverage arising from the fact that securities are always purchased at a significant discount to face, provides a




Figure 3.8 Theoretical and actual zero coupon curves

true risk vehicle for those seeking greater risk/return possibilities. As we have noted above, investors holding strips to maturity face no such risks, as their redemption value is certain, regardless of the level or direction of interest rates.

Tax issues

Strips are zero coupon securities, and thus pay no periodic cash flows; all cash flow settlement occurs at the maturity of the individual coupon strips and principal strips. Under most tax regimes, however, the accretion of value that occurs over time is considered a standard taxable gain. For instance, in the UK, all strippable gilts pay interest on a gross basis, and accrued returns are treated in the same way as returns from standard coupon-bearing gilts. The same occurs in the US marketplace. Though strips are often incorporated into investment strategies through the use of tax-deferred accounts, the actual accretion of value is still a taxable event. In fact, national systems with strip programs have favored this type of tax treatment in order to avoid distortions that might appear between coupon-bearing instruments and stripped securities.

While our discussion of stripped securities has obviously focused on government securities, the same technologies can be, and in practice are, applied to other types of fixed-income and hybrid securities. We shall revisit the concept of stripped instruments in the context of mortgage-backed securities in the next chapter, and with regard to convertible bonds in Chapter 8.

Mortgage- and Asset-backed Securities

4.1 INTRODUCTION

The global mortgage-backed securities (MBS) and asset-backed securities (ABS) sectors represent two of the most significant, innovative, and liquid elements of the capital markets. MBS and ABS are complex fixed-income securities backed by diversified pools of mortgages or other assets, which pass through or redirect cash flows from asset seller/servicers to investors. Although MBS and ABS are relatively new creations, having appeared as recently as the 1980s and 1990s, they have already transformed portions of the illiquid real estate mortgage and receivable/loan markets into liquid and tradable securities, freeing the balance sheets of sponsors/originators from capital or regulatory constraints in the process. Their development has also transformed other portions of the financial markets: structured notes, collateralized debt obligations, and insurance-linked securities, which we consider in subsequent chapters, rely on many of the same financial engineering techniques to provide issuers, intermediaries, and investors with desired results.

Many types of assets can be securitized, though some are simpler and more efficient to arrange than others.¹ In this chapter we consider the specific structure and applications of common MBS and ABS structures. Though the two sectors are characterized by fundamental differences, they share many common features, and can therefore be considered jointly. Within the MBS sector, we focus our discussion on residential and commercial pass-through securities, collateralized mortgage obligations, and stripped mortgage products; these apply to both residential and commercial mortgages. In the ABS sector, we concentrate on three key areas, including credit card securitizations, home equity loan securitizations, and other asset-based securitizations. Figure 4.1 summarizes these sectors. Note that our analysis is based primarily on the US market, which has been a pioneer in product development. We shall, however, consider companion products from the European and Asian markets in cases where they have become prominent.

4.2 DEVELOPMENT AND MARKET DRIVERS

The basic MBS, which is a pass-through, or participation, certificate backed by a pool of residential or commercial mortgages, was originally created in the late 1970s by the Government National Mortgage Association (GNMA, or Ginnie Mae). Working in tandem with the other two mortgage agencies – Federal National Mortgage Association (FNMA, or Fannie Mae), and Federal Home Loan Mortgage Corporation (FHLMC, or Freddie Mac) – major financial institutions expanded the secondary mortgage market significantly, adding liquidity to the origination, pooling, and trading process. As a result of these efforts, agency MBS have grown to become one of the largest sectors of the global fixed-income markets.

¹ In fact, the potential universe is broad, and includes mortgages, repackaged mortgage securities, student and auto loans, receivables (trade, credit card), leases (auto, equipment, aircraft, boats), remittances/payments/royalties/rights, tax liens, project finance and export flows, and so forth.

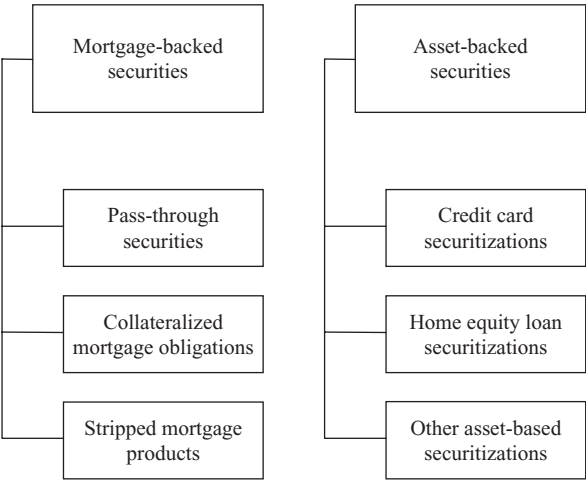


Figure 4.1 General MBS and ABS classes

Following the initial success of “vanilla” residential MBS, financial intermediaries applied the same technologies in other areas, including nonconforming “whole loans” (i.e. those not meeting the size or loan-to-value requirements of the mortgage agencies). This was an important evolutionary step, since many large residential loans, which were otherwise ineligible for inclusion under agency programs, became securitizable and transferable. Pooling and securitization of commercial property mortgages into commercial MBS (CMBS) appeared in the mid-1980s, but grew primarily during the late 1980s, when many US savings and loans were forced to divest their commercial loan portfolios in order to comply with new solvency regulations. The recovery of the US commercial real estate market during the mid-1990s led to even greater primary and secondary activity.

Financial intermediaries next centered on securitizing pools of MBS into instruments known as collateralized mortgage obligations (CMOs). The advent of the CMO in 1983 (via FHLMCs inaugural issue) was again a vital step in the development of the secondary mortgage market, leading to the creation of new structured assets capable of addressing the customized risk/return needs of investors. Specifically, CMOs allowed investors directly and accurately to manage their exposures to prepayment risk, a key mortgage product risk that we discuss at length below. The financial creativity of financial intermediaries led ultimately to the development of many different styles of CMO – some of which proved particularly risky (and financially damaging) in a rising-rate environment (in fact, when US interest rates rose sharply during the mid-1990s, many investors and banks lost heavily on their risky CMO holdings, which led to a temporary curtailment in new securitizations and product development. Following that chastening, the CMO market returned to its original “vanilla” tranching, and has grown steadily since that time). At approximately the same time that CMOs were emerging, intermediaries were also applying the stripping processes described in Chapter 3 to decompose MBS into principal-only (PO) mortgage strips and interest-only (IO) mortgage strips. These stripped mortgage products soon became a popular, if risky, method for expressing a view on rates; they also emerged as important tools for those seeking to hedge other classes of mortgage or fixed-income securities.

The overall US MBS market has grown rapidly over the past three decades, as a result of increased home ownership, property development/financing, and subsequent securitization; these

drivers have been central in creating product supply. In fact, mortgage originators continue to sell large quantities of loans into securitization conduits, in order to free up capital/balance sheets and create new business opportunities. Investor appetite for securities with strong liquidity, low credit risk, and attractive investment returns has helped fuel demand. Standard MBS can yield up to 100 bps more than comparable maturity government benchmarks; the yield pick-up is designed primarily to entice investors to accept a small amount of issuer credit default risk, and a reasonably significant amount of prepayment risk.

Though the MBS market developed in the US, it soon spread to other countries featuring large homeownership and relatively sophisticated financial intermediation processes. The UK mortgage market, for instance, launched securitizations of its own in the late 1980s; the Danish mortgage market developed at about the same time, and Spain, Australia, Germany, France, and Hong Kong followed with their own securitized mortgage products. While each national market features product specifications that reflect the unique characteristics of the local mortgage marketplace, most operate along similar principles, i.e. grouping together portfolios of diversified mortgages so that they can be converted into tradable assets. Naturally, not all financial systems feature local MBS. For instance, Japan and Korea, two countries with significant real estate mortgage markets, do not yet offer investors the chance to participate in the trading of local mortgage assets. Whether regulatory processes will allow such products to be introduced in the future remains to be seen.

Though early asset securitization mechanisms focused primarily on the pass-through structure (e.g. directly passing through principal and interest flows from borrowers to investors), structural developments based on the concept of revolving, or hybrid, flows have allowed medium- and long-term securities to be created, even when the underlying assets have a short tenor (e.g. 3–12 months). This advance led to interest and growth in the ABS market, which is based largely on shorter-term assets. The US ABS market dates back to approximately 1985, when Chrysler securitized a portion of its auto loan portfolio. Credit card securitizations followed in 1987, and by the end of the decade the market had expanded to include securitizations on home equity loans, boat loans, manufactured housing, and student loans; deals on equipment leases (e.g. aircraft leases) followed in the 1990s. Though the market's core is centered on credit card receivables, auto loans, and home equity loans, the sector also supports highly customized transactions that secure music royalties, sports ticket revenues, and film receipts. Portions of the ABS market have also taken hold in certain other countries. For instance, the UK features a growing multi-asset ABS market, while European auto and consumer loan ABS have gained momentum since the late 1990s.

The use of standard securitization technologies has made development of these markets possible. In particular, the structures that we consider in this chapter can be created when four essential factors exist:

- collateral pools are of sufficient size and quality;
- asset originators and servicers are capable and creditworthy;
- proper internal/external structural enhancements exist;
- adequate legal structure and protection are available.

MBS and ABS have proven successful because they provide intermediaries/issuers and investors with significant benefits; in particular, either form of the security:

- gives investors access to unique assets (e.g. mortgages, receivables) in a customized manner reflecting specific preferences/requirements related to risk, return, maturity – this can be done in pure pass-through form, or in a revolving structure;

- creates an efficient mechanism by which to purchase a desired asset – administrative savings result when an investor purchases a tranche of an entire portfolio of assets through a single transaction;
- permits companies and other originators of future cash flows to monetize their stream of forward earnings for use in a current period;
- allows originators to transfer certain cash flow uncertainties (e.g. prepayments, credit defaults) and risk exposures to investors, at a price;
- transforms otherwise illiquid mortgage and receivable assets into a more marketable form, adding greatly to financial sector liquidity – this ultimately helps lower issuer funding costs and creates more attractive and secure opportunities for investors.

Since risk/return characteristics of MBS and ABS are tailored via different tranches, they are suitable for a wide spectrum of investors, primarily from the institutional markets. Accordingly, purchasers of MBS and ABS include financial institutions, insurers, pension, mutual, and investment funds, and corporates. For instance, investors preferring little prepayment risk may be attracted to planned amortization class bonds; those seeking a significant amount of exposure to interest rate movements may purchase interest-only/principal-only tranches, and so forth.

4.3 PRODUCT MECHANICS AND APPLICATIONS

Though mortgage and asset-backed securities share many similarities, the products have enough unique features that we focus on each group separately.

4.3.1 Mortgage-backed securities

As we have noted, the basic MBS is created when an intermediary pools together a group of mortgages, and issues securities or participation certificates. The securities are conveyed to, and issued from, a bankruptcy-remote SPE or trust; the capital obtained from investors through the issuance of securities is used to fund the portfolio of mortgages acquired from the originator. MBS can be divided broadly into pass-through securities, CMOs, and stripped mortgage products; each can be decomposed further into a series of unique products, which we consider below. Before doing so, however, we discuss certain other points associated with the products.

Prepayments

The most interesting, and complex, risk associated with mortgages (and, by extension, MBS) relates to cash flow uncertainty driven by prepayments of borrowed principal. Prepayments, in turn, are a function of refinancings, home sales, partial/accelerated prepayments,² and defaults.³ In a typical fixed-rate, level-pay mortgage, the borrower is obliged to make regular (i.e. monthly or quarterly) payments of principal and interest (we ignore, for simplicity, mortgages that allow

² Curtailments and payoffs must be incorporated in a prepayment model as they impact cash inflows; partial prepayments shorten loan terms and the pool's weighted average maturity (WAM). Such prepayments generally are quite small in the early years of the mortgage (e.g. less than 1 % constant prepayment rate (CPR)), but can become much larger as final maturity draws closer (e.g. up to 20 % CPR). If a borrower cannot refinance, and the mortgage is well in-the-money, then the likelihood of curtailment rises dramatically.

³ Mortgages are granted on the basis of a borrower's credit standing/history, debt to income, loan term, LTV, property type, employment, and so forth. However, these details are not disclosed publicly in the pool. Accordingly, in order to account for defaults, a prepayment model generally must use a generic default curve. For instance, the curve might feature rising default rates for the first three years, followed by a leveling period, and then a gradual decline from year 8–10 onwards (as growing amortization decreases the likelihood that a negative equity position will arise). The default curve can be calibrated to take account of loan concentrations in an economically depressed or weakened region.

for increasing, balloon, or bullet payments of principal). However, interest rates may decline over time, inducing the borrower to repay the original mortgage and refinance through a new one with a lower interest rate (evidence suggests that refinancing becomes attractive when rates have declined by at least 100 bps below the original rate).⁴ Alternatively, the borrower may choose to sell the underlying property, repaying the original mortgage in the process.⁵ Or, the borrower may wish to make greater principal payments during a given month(s) in order to build equity more rapidly. In the worst case, the borrower may be unable to continue paying the mortgage, and declare a default. Each one of these payment/prepayment actions creates cash flow uncertainties that make the valuation of MBS a nontrivial task. Note that adjustable rate mortgages (ARMs) can be impacted by the same variables, as well as the changing value of the interest rate component (which may be readjusted on a monthly, quarterly, semi-annual, or annual basis).⁶ From an MBS investor's perspective, mortgage cash flows can be viewed as the sum of projected monthly interest less servicing fees, projected monthly scheduled principal repayments, and projected monthly additional principal prepayments.

Most prepayment activity is due to home sales when rates are stable. As rates rise, refinancing opportunities disappear and housing turnover slows, meaning that prepayments decline. Conversely, as rates fall, refinancing commences and housing turnover increases, leading to prepayment acceleration; we consider this at greater length below.

Refinancing, fuelled primarily by declining interest rates,⁷ remains the single largest driver of prepayments, and is the central focus of any prepayment model. This means, of course, that projected prepayment experience is extremely dependent on assumptions regarding future interest rates. Refinancing can be considered the equivalent of a call option on interest rates owned by the borrower; the option cannot, unfortunately, be modeled in a standard option pricing framework, due to the behavioral inefficiencies that characterize the housing market. That said, there is sufficient empirical evidence to suggest how refinancing patterns arise during a given cycle, and these results can be incorporated into a model. For instance, most cycles feature a burnout period, meaning that after a burst of refinancings, further refinancing activity tends to slow; each marginal decrease in rates yields less refinancing activity. Friction costs and barriers have declined in recent decades, making the refinancing process simpler, cheaper, and more efficient; the advent of online refinancing technologies, for instance, has allowed a greater number of borrowers to refinance during a particular cycle, meaning that the number of

⁴ When refinancing opportunities first appear, prepayment activity accelerates until a critical mass of activity is done; this phase is known as the "burnout," after which any further refinancing in response to a continued drop in rates is likely to be marginal.

⁵ Housing turnover through home sales is a major driver of prepayments, historically averaging up to 10% in a given year. Housing turnover is, itself, a function of various economic variables, including interest rates, consumer confidence, new construction, seasonality, current and expected economic growth, and so forth.

⁶ In fact, prepayments on ARMs tend to rise after the first six or twelve months as coupons reset (many resets eliminate the first period below-market "teaser" rate); the rising rates following the first period can lead some borrowers to refinance into new teaser rate ARMs (in fact, the number of new ARM-related products and the extremely competitive environment have led to a greater number of ARM refinancing opportunities – this, in turn, has led to a gradual increase in prepayment speeds in recent years). Once seasoned, ARMs settle in a range where prepayments are slightly above fixed-rate mortgages. When modeling prepayments on ARMs, some models distinguish between refinancing into new ARMs and new fixed rate mortgages; the nature of the borrower dictates the type of refinancing that is likely to occur: borrowers seeking to take advantage of teaser rates will refinance into new teaser ARMs, while those that prefer the stability of long-term fixed rates, but may have been unable to gain approval at a previous time, will refinance into standard fixed mortgage products. Other ARM-related products – such as convertible ARMs that change from floating to fixed rates after a specific period of time, and hybrid ARMs that convert from fixed to floating rate after five to ten years – tend to have unique prepayment characteristics of their own. Hybrids have proven especially popular in some markets, and growth has increased steadily as rates are often more competitive than standard 30 year fixed mortgages, and borrowers are able to defer the shock of migrating to a floating rate for up to five years. As a result of these benefits, borrowers often prefer to preserve rather than refinance their mortgages; prepayments are, accordingly, slower than on traditional ARMs.

⁷ A smaller amount of refinancing occurs as a result of credit improvement, e.g. a borrower's financial circumstances and credit standing improve, allowing migration from the subprime category to a standard credit category that commands a lower borrowing rate. When this happens, refinancing is a logical option.

“lagging” refinancers has declined. Refinancing incentives⁸ and mortgage product availability⁹ must also feature in the model. These can be supplemented by other, less tangible, forces – media coverage of the rate environment and refinancing opportunities, promotional campaigns sponsored by lending institutions, “psychological” rate barriers that might drive activity, and so forth. The micro and macro construction of the pool must also be considered. The market features many types of mortgages and rates, some of which promote refinancing action more readily and efficiently than others.¹⁰ In addition, the diversity of borrowers in any given pool leads to a broad range of refinancing options. Slow refinancers will eventually comprise a larger proportion of a seasoned pool as the fast refinancers depart, which will again impact on prepayment speeds and valuation. Refinancings cannot, of course, be viewed in isolation: even if rates are declining to the point where borrowers can achieve real savings, an associated deterioration in housing prices can impact behavior.

Housing sales, the second most significant driver of prepayments, and therefore a vital modeling input,¹¹ can be impacted by overall market turnover, regional/local turnover, seasoning, and lock-in/prepayment penalties.¹² Each of these variables, in turn, can be directly or indirectly influenced by other variables. For instance, overall and regional turnover may be affected by new housing starts and new home sales, economic growth, employment, taxes, interest rates, housing price inflation and propensity for “trading up.”¹³ Seasonal or cyclical effects must also be considered in the model.

Actual prepayment levels are determined by comparing principal cash flows received with those that are scheduled or expected; any difference between the two represents a prepayment, and generally is expressed in terms of the outstanding balance of the mortgage. In order to assess the potential impact of prepayment risk on mortgages – and by direct extension the impact on mortgage pools comprising an MBS – the industry has come to rely on certain simplifying assumptions designed to serve as a proxy for prepayment behavior. The most fundamental computation is based on the single monthly mortality (SMM) gauge, which is a monthly prepayment indicator that computes the fraction of a pool’s balance that prepays

⁸ Refinancing incentives can be regarded, in simplified terms, as the spread differential between the WAC and the mortgage rate; this, however, fails to account for the outstanding loan balance and the term of the loan. A more sophisticated approach compares the present value of all applicable inflows and outflows over the term of the loan:

$$RS = \left(\frac{Bal * r_o}{r_n} \right) * \left(\frac{1 - y^t}{1 - z^t} \right) - (Bal * (1 + RC))$$

where RS is refinancing savings, Bal is the current balance of the loan, r_o is the original loan rate, r_n is the new loan rate, RC is refinancing costs, y is $1/(1 + r_n)$, z is $1/(1 + r_o)$, and t is the remaining term of the loan.

⁹ Models generally distinguish between ARM and fixed-rate mortgage products; the most sophisticated can model changes generated when borrowers move from ARM to fixed-rate mortgages, or from ARMs to new ARMs. ARM pools generally are characterized by high mobility and fast seasoning, and are very sensitive to refinancing opportunities. Thus, if a pool is comprised solely of ARMs (a highly unlikely case, but useful for purposes of an example), the prepayment and valuation results will be very different than if a pool consists only of fixed rate mortgages.

¹⁰ For instance, a 15-year mortgage requires larger monthly payments than a 30-year mortgage, but has a faster amortization profile and may feature a lower rate. The range of new mortgage products available in many markets gives borrowers considerable flexibility. For example, borrowers can choose no-point or low point mortgages, cash-out refinancing loans (a form of home equity lending), and so forth.

¹¹ Though housing sales are an important driver of prepayments, historical evidence suggests that they tend not to create speeds in excess of 10% PSA CPR in a given cycle. Refinancings, in contrast, generally are multiples of that figure.

¹² A lock-in provision or prepayment penalty provision on a loan can also be quantified to determine whether rates will have an impact on prepayment behavior. For instance, if the lock-in rate on a loan is below current market rates, a borrower has an incentive to keep the current property; if current market rates fall below the lock-in rate, the incentive begins to fade. The actual dollar impact can be analyzed by examining the present value impact of current versus expected market rates. A prepayment penalty provision requires that the dollar value of the penalty be compared with the cost savings that can be achieved in a refinancing scenario when market rates fall below the original loan rate; prepayment penalties are unusual in residential mortgages, but they are quite common in commercial mortgages.

¹³ In some markets, activity is driven by the desire or need for existing homeowners to sell their original homes and purchase larger ones; if the trade-up occurs in the same locality, there may thus be purchases and sales impacting two sets of properties per household.




Figure 4.2 PSA CPR speeds

during the month. The constant prepayment rate (CPR) model serves as an annualized version of the SMM. The CPR assumes no difference in seasoning (i.e. the amount of time a loan has been outstanding), meaning that an assumed prepayment rate of 2 % applies to a loan that has been outstanding for six months or five years; this assumption is considered unrealistic in light of available historical data.

In the mid-1980s, the US Public Security Association (PSA) introduced a variation on the SMM/CPR approach to account for seasoning phenomena. The PSA, in examining historical data, determined that prepayment patterns change based on loan life: specifically, new loans have lower prepayment rates than seasoned ones. The PSA model adjusts the CPR by the age of the loan. The base case PSA model, known as 100 % CPR, assumes 0 % prepayment for new loans, 0.2 % prepayment for the first month, increasing by 0.2 % per month for the next 30 months, and then converting to a flat 6 % per year thereafter. This 100 % benchmark can then be increased (e.g. to 150 % or 200 %) to reflect faster prepayment scenarios, or decreased (e.g. to 50 % or 75 %) to reflect slower prepayments. Figure 4.2 highlights various PSA CPR speeds.

It is common in MBS valuation to perform scenario analyses to determine interest rate sensitivity of the CPR and, by extension, the sensitivity of the WAM and yield. For instance, Figure 4.3 illustrates a hypothetical CPR curve for changes in interest rates; this curve, which can be constructed from historical prepayment data, can then be used to generate possible MBS values.

Ultimately, an effective prepayment model must take account of all of the variables that can impact prepayments in order to provide an indication of possible future prepayment speeds, as noted in Figure 4.4, so that a specific MBS can be valued properly. In some instances, models are atomized to project each individual risk variable independently; these must, of course, be internally consistent so that the inputs and outputs among each independent submodel provide rational explanations of behavior (e.g. the effects of housing turnover from one submodel must be consistent with the effects of refinancing behavior from a second submodel). By creating this additional level of detail, an intermediary or investor can gain additional insight into future prepayments and security value.¹⁴ Importantly, the efficacy of a model can be back-tested by running simulations against actual prepayment experience.

¹⁴ The impact of different aspects of prepayment risk – such as refinancing, housing turnover, and so forth – can be ascertained, in part, by the use of partial durations, i.e. the change in the price of a security for a unit change in each of the specific variables through to impact on prepayments.




Figure 4.3 Effect of interest rate scenarios on CPR

Different types of MBS have different prepayment characteristics and sensitivities. Thus, in creating a workable model, attention must be given to the type of security being analyzed/traded; there is no single model that can compute prepayments and valuations properly across multiple types of instruments without some adjustment for the specifics of the asset/pool. For instance, GNMA pools, which we discuss below, tend to feature higher LTVs than commercial pools. High LTVs dampen prepayment speeds, leading to different values than commercial pools with lower LTVs. However, as GNMA loans season (and assuming housing market stability), LTVs drop and prepayments rise compared to commercial pools. Various other differences exist between other classes of MBS.

The modeling of prepayments must be viewed as a dynamic process, with parameters that change over time as market, borrower, and investor characteristics change. For instance, the




Figure 4.4 Prepayment projections from a model

introduction of new mortgage products, fundamental changes in interest rate structure, or differences in borrower behavior can all combine to permanently alter prepayments. A model that fails to take account of these changes will soon be of limited use.

Valuation

The prepayment risks described above are a fundamental input in the valuation process. Indeed, intermediaries and investors attempt to model prepayment behavior so that they can use the output to value an MBS more accurately. There are various ways of valuing MBS; some are elemental and simplistic but easy to implement, while others are computationally rigorous, though may ultimately prove to be more accurate. The potential value/return of an MBS can be analyzed through a simple static yield computation, or the more complex option adjusted spread (OAS) process we introduced in Chapter 3 (based on the concept that an MBS is an option-embedded security).

The static yield calculation is a cash flow yield derived from assumptions about prepayment speeds; it assumes that the investor will reinvest at the yield to maturity rate, and that the MBS is held until the final payoff date based on the prevailing prepayment assumption. The reinvestment assumption is critical, as prepayments (and thus reinvestments) occur every month or quarter; if cash flows differ from the assumption, then the cash flow yield will differ as well. Under a static cash flow yield process, the yield spread of an MBS is equal to the difference between the cash flow yield and the yield to maturity of the risk free benchmark (e.g. a Treasury bond); this, however, fails to account for the term structure of the benchmark and expected interest rate volatility, which will impact prepayments and cash flows.

We extend the OAS concept discussed in the last chapter to the option-embedded MBS. The OAS values an MBS by comparing the security to a mini-portfolio of zero coupon risk free benchmarks that have certain projected cash flows; the portfolio can be adjusted to reflect the MBS default risk and cash flow uncertainty. The resulting spread makes the present value of projected cash flows from MBS, discounted at spot rate and spread, equal to the market price. The OAS process begins with the computation of a yield curve spread as a measure of the return of principal over multiple periods, with each MBS cash flow discounted by the appropriate risk free benchmark and spread using forward rates. However, the yield curve spread still assumes constant rates and cash flows, so it is only appropriate for current scenarios. In the second step, the OAS model computes the yield curve spread through a series of forward rates and then captures the effects of rate volatility to suggest a possible path of future rates.¹⁵ Specifically, the OAS model creates random interest rate paths based on model parameters, and each path is used to project prepayment rates and MBS cash flows. Once computed, the values are discounted to generate an estimate of the average present value.¹⁶ The cost of the embedded option is simply the difference between the OAS and the yield curve spread, and can be interpreted as a measure of the investor's cost of rate volatility. Since OAS models are based on many assumptions, care is required in implementation and interpretation.

¹⁵ In practice, most OAS models use an arbitrage-free term structure model calibrated to the current market. Models may be 1- or 2-factor, with or without mean reversion. Rate volatility must be incorporated properly, as short-term rates are more volatile than long-term rates.

¹⁶ The model can, by extension, be used to compute option-adjusted duration and option adjusted convexity in order to determine the true price sensitivity of a particular security. Duration generates the percentage price change for an x bp parallel shift in the curve, assuming OAS is constant; securities that are sensitive to prepayments feature a duration that is lower than the standard modified duration of other fixed income assets. Convexity acts in a similar manner.

Regardless of the valuation model used, securities feature price compression, or negative convexity – a state where price appreciation is lower, and price depreciation is greater, than on a standard positive convexity fixed-income security. Negative convexity arises as a result of the prepayment option the MBS investor grants the mortgage borrower. For instance, in a falling rate environment, the upside price potential of an MBS is limited: if borrowers refinance in the lower rate environment, MBS investors holding securities that are prepaying very rapidly must contend with the prospect of having to reinvest in a lower-rate environment, and thus suffer from contraction, or reinvestment, risk. Conversely, in a rising-rate environment, the price of an MBS will decline, but at a greater rate than conventional fixed-income securities: rising rates slow prepayments and increase the amount invested in the pool at below-market rates; lack of prepayments means investors do not receive accelerated prepayment cash flows, and cannot therefore reinvest at new, higher rates – thus suffering from extension risk.¹⁷

Legal and structural issues

A mortgage portfolio can be securitized through a bankruptcy-remote SPE or trust via a true sale mechanism or a synthetic process; the true sale tends to be more common, though this varies by country and product. The true sale mechanism involves the transfer or assignment of mortgage assets from the originator to the SPE (or trust), which then issues notes to investors; the originator, however, retains legal title to the assets. The trustee, through a power of attorney from the originator, has the ability to transfer legal title to the SPE, if necessary. The mortgage assets and the cash flows that they generate secure the notes; the cash flows, in turn, are used to service the principal and interest obligations of the note, as well as the fees/costs associated with the structure. Investors acquiring a pass-through security effectively receive an undivided ownership interest in the mortgage pool. Note that an alternative to the true sale is a subparticipation, which transfers risks and legal rights, but leaves the assets on the originator's balance sheet.

In the synthetic mechanism, the SPE (or trust) issues notes to investors with an economic return that is linked to the performance of a pool of reference mortgage assets (e.g. a type of credit linked note, such as we discuss in Chapter 5). The SPE uses the proceeds of the notes to purchase qualifying assets, and simultaneously provides the mortgage originator with de facto risk protection in exchange for a premium. The coupons from the qualifying assets and the premium are used to service the notes, while the qualifying assets are used to secure the investors' interests (note that this approach is used to create the synthetic collateralized debt obligation product discussed in Chapter 6). The synthetic structure avoids difficulties related to taxes and foreclosure laws; the originator retains the actual assets on its balance sheet.

The two approaches feature similarities and differences. For instance, true sale and synthetic deals require proper servicing and trustee administration; both amortize through prepayments, and default-based losses in the underlying asset pools are borne by investors. However, synthetics are not exposed to liquidity shortfalls, as only realized losses are transferred to investors; this obviates the need for a liquidity provider. But the mechanism heightens risks related to security interest/transfer and the performance risk of the servicer and liquidity provider. The synthetic method, which permits greater customization, does not feature any security

¹⁷ Given these features, the optimal time for an MBS investor to enter the market is when securities can be purchased at an apparent discount to the theoretical value.

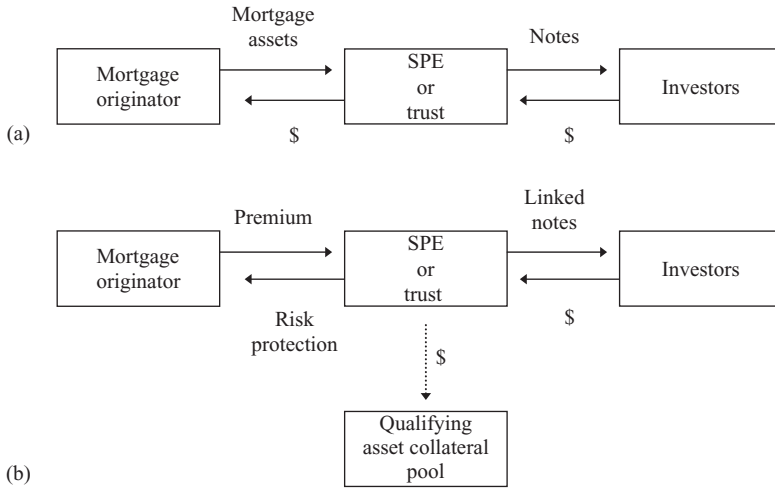


Figure 4.5 Mortgage pass-through structures. (a) True sale; (b) synthetic arrangements

interest/transfer risks (as no asset sales/balance sheet transfers occur), but creates only a weak link between the qualifying asset collateral pool and the notes. In addition, synthetics feature no excess spread, and are often structured in a manner that is quite opaque. True sales can suffer from illiquidity, hence the need for liquidity providers to supply liquidity backing. If a payment is interrupted, the liquidity provider supplies funding, so becoming a senior creditor in the event of losses. The true sale method improves transparency and generates excess spread that can be used to enhance the credit quality of the notes; it also allows the SPE to eliminate any dependency on the performance of the originator. Figure 4.5 illustrates the general flows of true sale and synthetic MBS arrangements.

A servicer is an integral part of the deal process. This party is responsible for collecting and forwarding principal and interest flows from the original borrowers to the SPE/trust on a timely basis. In some cases, the credit originator and the servicer are the same party (i.e. servicing retained portfolios); in other cases, the functions are split (i.e. so-called servicing released portfolios). Since the servicer is responsible for gathering cash flows that will be forwarded to investors, the role is critical. Indeed, the strength and experience of the servicer must factor into any structural decision by the trustee, and ongoing performance must be monitored carefully. The trustee must also designate a backup servicer to perform the task should the original servicer be unable to do so.

Pass-through securities

The pass-through security, as the name suggests, is a mechanism designed to collect and pass through cash flows from an underlying pool of mortgages to the investor holding the security; the SPE or trust deducts a small amount of the cash flow in the form of fees, and then passes through all interest and principal flows to investors. The frequency of the pass-through payment to the investor depends primarily on the payment frequency of the mortgages in the pool; in most cases, borrowers are obliged to make payments on a monthly basis, meaning cash

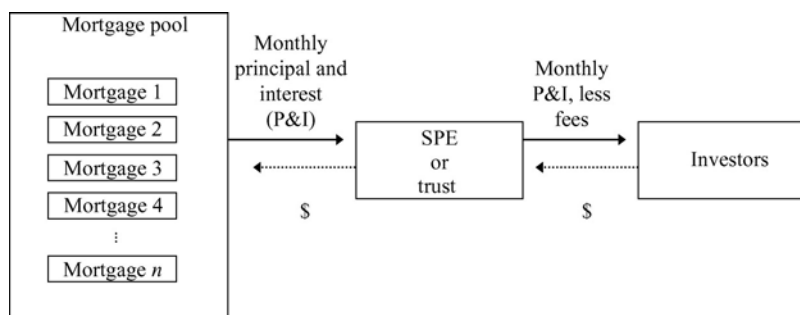


Figure 4.6 Fundamental mortgage pass-through structure

flows passed to investors also occur on a monthly basis; quarterly-pay structures also exist. The basic pass-through framework is reflected in Figure 4.6.

The US market features agency pass-through securities, which are arranged and issued by the three agencies noted above, and private label residential and commercial pass-through securities, which are originated by private sector institutions. Agency pass-throughs include only those mortgages that meet specific size, LTV, and credit criteria established by each agency; these are known as “conforming mortgages.” In general, mortgages included in agency pools are fixed-rate, level-pay, single-family residential loans of less than \$250 000 (and with LTVs below 80 %), and may have explicit guarantees from other governmental agencies.¹⁸ Private label residential pass-throughs, in contrast, feature loans with parameters that diverge from agency criteria (i.e. “nonconforming loans”), including those that exceed the maximum loan threshold, those with higher LTVs, or those granted to subprime borrowers. Commercial MBS are supported by mortgages extended to finance commercial development, including multi-family residential dwellings, hotels, offices, industrial parks, shopping/retail developments, and so forth; such loans are often extremely large.¹⁹ Non-US pass-throughs represent an important component of the overall market; various countries have introduced government-related or private label pass-through structures over the past two decades, some of which have become very liquid (e.g. UK, Denmark).

Agency pass-through securities

The US agency pass-through market forms the core of the MBS market. Agency pass-throughs are created when an authorized agency purchases a pool of conforming mortgage assets from a mortgage originator and issues securities backed by the pool to investors; each pool tends to be comprised of thousands of mortgages with similar features (e.g. single- or multi-family

¹⁸ While such loans account for the vast majority of the collateral included in standard conforming MBS pools, other loans may also be included, such as those with balloon, bullet, or graduated repayment schedules.

¹⁹ It is worth noting that in the US market, real estate investment trusts (REITs) can qualify as pass-through entities if they pay dividends equal to 90 %+ of their taxable income and at least 75 % of their assets are invested in real estate. Equity REITs invest in, and own, properties and derive their cash flows from rental income. Mortgage REITs provide financing to real estate companies and purchase both mortgages and MBS; the income from these sources is used to pay investors. In order to achieve particular credit ratings, REIT investors are protected typically via covenants related to debt/assets (e.g. less than 60 %), unencumbered assets (e.g. at least 15 %), interest coverage (e.g. at least 1.5×), and so forth. In contrast to other MBS pass-through conduits, REITs tend to be much more fluid, dynamically altering their portfolios and adding/subtracting leverage based on market opportunities. Though REITs are often purchased directly by investors, they may be included in collateralized debt obligations.

residential home loans with 15–30 year original maturities, fixed or adjustable rates, maximum loan size and loan-to-value percentage). Pools may feature loans that are new or seasoned; seasoned pools are characterized by relatively stable prepayments.

GNMA, FNMA, and FHLMC issue agency pass-throughs as fully modified or modified assets. Through the fully modified pass-through structure, the agency provides investors with a guarantee of timely payment of principal and interest, regardless of when the agency receives the underlying cash flows from the mortgage pool borrowers. Thus, if a default or delinquency occurs in the pool, the agency will absorb the loss or delay, meaning the investor will suffer no financial consequences. In the modified pass-through structure, the agency guarantees payment of principal and interest, but only the interest is guaranteed to be timely. If a default or delay occurs, the investor receives the interest cash flow but may have to wait a period of weeks or months before receiving the principal component. Each agency pass-through security is defined by issue date, underlying collateral, coupon, original pool balance, current pool balance, pool factor (proportion of the original principal balance outstanding as of a particular date),²⁰ weighted average maturity (WAM, which is the average time that a dollar of principal remains outstanding), and weighted average coupon (WAC, which is the average coupon of the pool, rather than the security; the difference between the WAC and the security coupon is known as the servicing spread). As a result of their special status, agency pass-throughs are exempt from SEC registration requirements.

GNMA, which is explicitly guaranteed by the US government, issues only fully modified MBS under its GNMA I and GNMA II programs; GNMA I MBS include only fixed-rate mortgages, while GNMA II MBS feature fixed mortgages and ARMs. The Federal Housing Administration (FHA), Veterans' Administration (VA), or Farmers' Home Administration guarantee the mortgages underlying GNMA MBS. Most GNMA MBS pools are issued via the GNMA I program, meaning the underlying loans are fixed-rate mortgages with similar rates (i.e. a constant WAC). GNMA II pools allow for different loan rates, meaning the WAC can change over time. GNMA MBS tend to feature lower prepayment rates than other agency MBS, as the pools are comprised of FHA/VA loans extended to lower-income households with lower value properties; these borrowers are generally less likely to sell their homes or prepay their mortgages.

FHLMC is a corporate institution that does not carry the explicit backing of the US government. Given its role in the mortgage market, however, the agency and its modified and fully modified participation certificates (PCs) trade on the basis of an "implied guarantee." FHLMC operates a cash program, where the agency purchases individual single- and multi-family residential mortgages from originators and then pools them; the assets are funded via the issuance of PCs. It also operates the Gold PC program, which allows originators to swap pooled mortgages into PCs. FHLMC pools allow for mortgages with different rates, meaning the WAC, like the GNMA II WAC, can change over time.

FNMA, like FHLMC, is a corporate institution that lacks an explicit US government guarantee, but enjoys some degree of implicit support. FNMA issues fully modified pass-throughs under various programs, including a swap program similar to FHLMC's Gold PCs, several programs covering level-pay, fixed-rate mortgages, and a program related to ARMs; FNMA securitizes single- as well as multi-family residential property loans. FNMA pools also permit mortgages with different rates, meaning the WAC can change over time.

²⁰ The pool factor declines over time as a result of both amortization and prepayment forces.

All three agencies provide investors with the opportunity to swap outstanding pass-through securities for a pro rata share in a large combined pool of securities. The amalgamation of smaller pieces of pass-throughs into single, larger securities helps create more liquid benchmark issues based on a broadly diversified portfolio of mortgages. These swap/amalgamation securities are issued as FHLMC Giants, FNMA Mega-Pools, and GNMA Platinums.

New agency pass-throughs are traded on a “to-be-announced” (TBA) basis, where the buyer and seller agree to general criteria related to the amount, agency, coupon, settlement, price, and par value – but not the specific pool, which is not determined until two days before trade settlement (the so-called “48-hour” rule). The seller in a TBA trade delivers pools that qualify as “good delivery” under established delivery guidelines. All trades are settled electronically, via book-entry.²¹ TBA trading is a vital component of the market, as it allows liquidity to build in new pools based on an assessment of which pools are likely to be delivered.

Secondary trading ultimately migrates to seasoned issues based on specific original maturity or WAM parameters (i.e. the TBA vintage market). Seasoned issues trade actively in the OTC market, with many dealers making two-way markets in a range of issues; liquidity in benchmark issues is excellent. Since prepayment risk is such an important element of pass-through valuation and trading, seasoned securities, which possess a history of prepayment activity by which to judge current and future value, tend to trade at a premium to TBAs. Secondary trading also occurs in specified pools, in which the buyers know all the details regarding the underlying collateral; this characteristic means that specified pools also command a price premium over TBAs. Secondary yields in the agency pass-through market are quoted as a spread to Treasuries, though care must be taken in interpreting the spread as the two securities have different cash flow characteristics (e.g. timing of principal return).²²

Private label pass-through securities

Private label pass-throughs, sometimes referred to as conventional pass-throughs, nonagency pass-throughs, or A/Bs, are issued directly by financial intermediaries, including banks, savings and loans, and securities firms. As noted above, the collateral pool underpinning private label securities includes residential mortgage loans that fail to conform to agency requirements; not surprisingly, the pools typically are far less homogenous, and may include mortgages of varying size, LTVs, and credit quality. Unlike agency MBS, private label securities carry no timely or ultimate guarantee of principal or interest, and must be registered with the SEC. Most need to be rated via the public credit rating agencies, and actually achieving a rating similar to the AAAs of the agency market generally requires some form of credit enhancement.

The largest sector of the private label market is based on jumbos, or loans that are at least three times larger than those granted by the agencies (and with LTVs that are often below 80 %, i.e. more borrower equity is involved). Jumbos are extended to affluent borrowers, so

²¹ In a typical transaction, the buyer instructs the clearing bank to receive the pass-throughs from the seller via the seller’s clearing agent; the instructions and appropriate credits/debits to the accounts are then made. In the US, market trades may flow through the centralized MBS Clearing Corporation (MBSCC) if both parties are participants; MBSCC focuses primarily on TBA, rather than seasoned, trades. Fedwire, operated by the Federal Reserve, clears and settles FHLMC, FNMA, and CMO trades, Depository Trust Company (DTC) clears and settles ABS and nonagency CMOs, Participant’s Trust Company, a division of DTC, clears and settles GNMA, and Clearstream/Cedel clear and settle international MBS. If a fail occurs, i.e. securities are not delivered per the terms of the trade, the buyer still owns the securities and is entitled to receive interest due, but need not pay for them at the agreed price until the delivery is complete.

²² Liquidity in the agency MBS market is reinforced through an active repurchase agreement market, allowing participants to finance their positions on a collateralized basis. Participants can choose to deal in the repo market (requiring the return of the precise securities lent) or the dollar roll market (requiring the return of “substantially similar” securities).

credit default experience generally is much more favorable than on other loans. Jumbos tend to have faster prepayments than conventional mortgages, as borrowers are more attuned to interest rates and refinancing opportunities; the large loan size makes borrowers particularly sensitive to any possible savings that can be obtained through refinancing. Jumbo pools are often concentrated in regions that feature higher real estate prices (e.g. in the US this can include California, the Northeast, and Florida).

Private label pass-throughs may also be created using so-called “alternative A” (Alt A) loans as the backing collateral. Alt As are moderate/high credit quality loans, generally up to 30 % larger than conforming loans (but smaller than jumbos), with LTVs that may exceed agency limits.²³ Alt A loans are often supported by limited borrower documentation, and borrowers may feature debt to income ratios that are below those required by the agencies. It is worth stressing that Alt A loans are not lower credit quality loans, but loans with characteristics that are different from those in conforming pools; true low quality loans form part of the subprime sector, discussed below. The prepayment experience of Alt A pools is generally mixed; borrowers may not initially have access to the same refinancing opportunities as conventional agency borrowers as a result of their unique characteristics, suggesting lower prepayments versus agency pools in a declining rate environment. This, however, appears to be true only in the short term, over the medium term, Alt A pools seem to exhibit greater prepayments than agency pools, as a larger number of “slow” refinancers begins to take action.

The subprime (B&C) loan sector forms the third part of the private label market. As the name suggests, subprime loans are credits extended to borrowers with a prior history of credit problems. In general, loan balances are smaller, and LTVs are lower, than in conforming pools, so that lenders receive additional collateral protection. Pool diversification is considerable in order to spread the risk of borrower defaults. Since lower-quality borrowers tend not to have the same financing options as conforming borrowers, sensitivity to refinancing opportunities is much lower. Indeed, B&C prepayments from declining interest rates are much lower than on standard agency paper: refinancing is driven primarily by credit upgrades of borrowers, which tends to be a much slower process.

Since private labels carry a greater element of credit risk than agency securities, and cannot offer the same guarantee of timely/ultimate repayment of principal and interest, many feature some form of credit enhancement in order to generate AAA or AA ratings and make the securities more appealing to a broader base of buyers. Common forms of credit enhancement include third party guarantees or bond insurance (i.e. from monoline insurers or financial guarantee companies), bank letters of credit, or senior/subordinated structuring (e.g. issuance of senior and subordinated tranches backed by the same pool, with the senior investor holding priority over the collateral and the subordinated tranche absorbing the first loss component). The preferred form of credit enhancement generally is a guarantee from the sponsoring institution, as this can help lower the cost of the transaction; however, this is only achievable when the sponsor is highly rated. In practice, the most common form of credit enhancement is senior/subordinate tranching, with the subordinated, or junior, tranche occupying the first loss position. The actual amount of subordination in a private label pass-through depends largely on the target rating the intermediary/sponsor is attempting to obtain, and the nature of the

²³ The first Alt A securitization occurred in 1994 via Independent National Mortgage Corporation, and the market has grown steadily since that time as new participants have entered the sector. Though the Alt A market traditionally has been dominated by private sector financial institutions, there is considerable evidence that the agencies are encroaching: GNMA, FNMA, and FHLMC have, for instance, introduced new products that include limited documentation underwriting, expanded LTVs, and so forth. This injection of competition ultimately may force a change in Alt A rates and redefine the prepayment characteristics of Alt A pools.

underlying collateral. In general, AAA senior tranches may require subordination ranging from 4 % to 20 %.²⁴ We consider subordination in greater detail in Chapter 6.

Secondary trading activity in nonagency pass-throughs sponsored by major institutions is considered to be quite good, particularly for large issues that receive de facto “benchmark” status. Nevertheless, nonagency securities tend to trade at wider spreads than similarly rated corporate bonds, despite the fact that the securities are supported by a diversified pool of collateral; this is primarily a function of the uncertainties caused by prepayments and the modeling complexities involved in the valuation process.

International pass-through securities

As noted, the market for international pass-throughs, which developed in the late 1980s and early 1990s, is now an established element of the structured asset market. Countries that made an early foray into the sector, including Canada, Australia, New Zealand, UK, Denmark, Germany, and Hong Kong, have achieved considerable success in creating the foundation for securitization of residential and commercial mortgage loans. Many of these efforts were motivated originally by a desire to widen bank funding bases outside traditional retail deposits, where growth rates have reflected a slowing trend for a number of years. More recent entrants, including Italy,²⁵ France, The Netherlands, and Spain, feature less mature markets and relatively modest activity.

Mortgage deals within the European Union are still national in nature; there is no activity associated with Euro-zone or cross-border transactions, primarily because the standards across member mortgage markets are still unique. Most European countries follow standard securitization techniques related to pool construction/diversification, sequential payments,²⁶ and internal enhancement via senior/subordinated tranching²⁷ and use of reserve funds (funded upfront and then financed via excess spread). However, in contrast to the US market, some national systems permit principal repayments on mortgages in a pool to be redirected towards the purchase of additional mortgages during a specified revolving (or substitution) period. In addition, and again unlike the US, pool losses are not written down, but are recorded in deficiency accounts, with excess spread applied to the account. In general, the Euro-zone countries with active mortgage markets tend to originate floating mortgages (or impose significant prepayment penalties), so that prepayment risks are not as great as in the US market.

Interestingly, a number of countries with very large property developments and holdings, such as Japan, India, and Korea, have not yet managed to create meaningful local MBS programs. Reasons for inactivity relate primarily to restrictive regulatory barriers (including lack of securitization laws), insufficient economic incentive for originators to remove mortgages from their balance sheets, and/or inadequate structuring capabilities/analytics at the level of

²⁴ The resulting AAA tranche may then be used as part of a collateral pool supporting a CMO structure, while the subordinated tranche, which may initially be rated in the BB range, can be tranching using senior/subordinated technology in order to generate further credit quality differentiation.

²⁵ Italy, for instance, amended its laws to permit securitization in 1999. The legal changes, along with consolidation within the banking sector, have been instrumental in promoting growth in the Italian MBS and ABS markets.

²⁶ Note that with sequential paydowns, the average cost of the notes rises after the initial tranches are paid off; this reduces the amount of available excess spread flowing to the residual investors. Some transactions attempt to mitigate this effect by allowing a certain amount of pro rata, rather than strict sequential, redemption. This, however, is only permitted once performance triggers have been met (e.g. a minimum level of credit enhancement on senior notes, a fully funded reserve account, and so forth).

²⁷ European deals often feature variable pay term notes (VPTNs), an additional type of tranche where the investor is obliged to pay remaining value at a future date. The issuer thus uses the proceeds to redeem an existing outstanding class on a proximate coupon date, exchanging the VPTN for a standard pass-through at that time.

originators, intermediaries, and/or investors. In addition, countries lacking housing agencies (e.g. FNMA, FHLMC, GNMA, Canadian Mortgage and Housing Corporation) to help promote origination, transfer/securitization, and trading of mortgage loans are effectively missing key conduits in the process.

United Kingdom

The UK, which features the second largest mortgage market in Europe after Germany, traces its inaugural MBS deals back to 1987. The UK experienced steady increases in residential homeownership and a gradual rise in property prices throughout the 1990s and into the new millennium; these factors have helped fuel interest and growth in residential MBS. Loans underlying pools of UK MBS carry maturities of up to 25 years and may be structured with variable rates (capped or uncapped) or hybrid fixed/variable rates. Traditional mortgages are either interest-only loans or endowment loans (with an investment plan or endowment policy funded by regular payments from the borrower that liquidates the principal balance at maturity); the market in recent years has also begun to feature a greater amount of repayment mortgages, which are akin to the amortizing, level-pay mortgages common in the US. There is no government agency involved in the local mortgage market; private sector lenders originate all loans, and all issuance is structured and arranged by financial intermediaries.

Mortgages in a pool are conveyed via a true sale of assets²⁸ into an SPE or, following changes to tax laws in 2000, a master trust,²⁹ which then issues the pass-through securities (or trust receipts) to end-investors. Deals often allow for the inclusion of additional mortgages during a revolving period; these must, of course, adhere to existing pool criteria. Most pass-through tranches are issued with floating-rate coupons, consistent with the fact that most mortgages are variable rate; the SPE may use swap hedges if the pool contains fixed/variable-rate mortgages. In order to expand the potential investor base on a given deal, the SPE may also swap tranches directly from sterling into dollars or euros. A typical deal is likely to feature up to seven senior short and long-term tranches, along with one or two subordinated sequential pay amortizing tranches. Some UK pass-throughs contain callable features with step-up coupons.³⁰

The UK mortgage market features much lower prepayment rates and more stable prepayment characteristics than the US market (e.g. average of 8 % CPR, maximum of 18 % – half of the US average). This is due primarily to the fact that the UK market has a much larger proportion of variable-rate mortgages, along with larger prepayment penalties on fixed/variable-rate mortgages. LTVs are somewhat higher than in the US market (averaging 90 %³¹), while delinquency and default rates are similar to those in the US.³² Most pools make use of seasoned mortgages. UK default experience suggests the greatest likelihood of mortgage borrower default occurs in the first five years of a loan's life, so pools tend to be weighted with mortgages that have been outstanding for at least several years. Though most mortgages tend to fit certain conforming

²⁸ In the UK, lenders can transfer mortgages to an SPE without notifying the borrower through a process known as equitable assignment; the SPE's interest in the collateral is thus not fully perfected, but the mechanism still meets true sale requirements.

²⁹ In fact, the UK MBS master trust structure functions just as the ABS master trust discussed later in the chapter, using mortgage assets rather than receivables. The advantage of the master trust is that it allows for the issuance of multiple tranches/deals from a single trust umbrella at different times, and permits the creation of bullet securities with short and defined maturities.

³⁰ It is interesting to note that several innovative UK pass-through deals have made use of the credit card securitization techniques we discuss later in the chapter, i.e. issuance via a master trust, use of substitutable assets, and use of soft bullets (which accumulate principal in a funding account for investor interest, make a bullet payment, and then redirect principal to the mortgage seller). The pool generally is set larger than the amount of bonds outstanding in order to achieve these goals.

³¹ Mortgage insurance is required on any loan with an LTV greater than 75 %. In fact, mortgage indemnification of this type provides partial protection for the lender and, by extension, the MBS investor.

³² There is evidence to suggest that despite higher LTVs (e.g. greater leverage), there is no significant difference in delinquency/default, as UK lenders are more willing than their US counterparts to work with borrowers in restructuring troubled loans.

criteria, the marketplace also supports deals constructed with nonconforming loans (e.g. larger balances, longer or shorter tenors, higher LTVs).

Securities generally feature one or more forms of credit enhancement in order to achieve suitable ratings. The earliest transactions in the market carried third party mortgage pool insurance, though this ultimately proved to be too expensive to use on an ongoing basis, and is now relatively uncommon. The market has since evolved to include credit enhancement via reserves, surety bonds, and subordination. For instance, most UK pass-throughs feature prefunded cash reserve accounts designed to protect against slow pay or default. Excess spread between the amount collected by the servicer from the mortgage pool and the amount due and payable to investors may also be used to fund additional reserve accounts. Surety bonds from highly rated financial institutions or monoline insurers may be employed on occasion. These forms of enhancement, however, are less common than the senior/subordinated tranching noted above. Mortgage deals often feature liquidity facilities in order to protect against cash flow shortfalls; importantly, these facilities are not intended to provide any credit protection, so cash flow shortfalls arising purely from defaults do not qualify, and cannot be covered by liquidity facilities. Secondary liquidity tends to be much lower than in the US, as investors purchasing sterling-denominated pass-throughs are largely buy-and-hold institutional accounts.

Denmark

Denmark is a pioneer in the European MBS sector, having started its market at approximately the same time the US market was gaining momentum. The local MBS market features pass-through securities securitized by pools of fixed-rate, level-pay mortgages with final maturities of 10 to 30 years. The mortgages, which are a mix of residential and commercial property loans, are prepayable without penalty and call for quarterly interest payments. Diversified portfolios of loans are grouped together in pools of \$1 b and \$10 b; though this diversification is vital, it is important to remember that the Danish economy is relatively small, and an economic downturn can impact multiple sectors/borrowers simultaneously – there is thus a degree of sector concentration inherent in the securities. Refinancing is fairly common, as borrowers generally are sophisticated and rate-sensitive, and refinancing costs are very low. This means that the pass-throughs are exposed to refinancing-driven prepayment risk, just as in the US market. However, the actual process of allocating prepayments occurs via a “lottery” system, rather than the standard pro rata split amongst tranches. This makes assessment of the impact of prepayment on holdings more complex, as some allowance must be made for the random nature of the lottery allocation.

Germany

The German Pfandbrief, or loan/mortgage bond,³³ market has grown rapidly and constitutes a major portion of the European capital markets. Pfandbrief (meaning letters of pledge), are issued by German banks under special legislation, and are backed by specifically registered and identified asset pools. Pfandbrief are available in two forms: public bonds (offertliche), which are backed by loans to public sector entities, and, of relevance to our discussion, mortgage bonds (hypotheken), which are backed by residential and commercial mortgages with maximum LTVs of 60 %. Though the mortgages underlying a typical mortgage-backed bond may be residential

³³ It is worth noting that although Germany accounts for 90 % of the European pfandbrief market, similar structures also exist in Luxembourg, Spain, Ireland, and France. For instance, France features “obligations foncières” issued by the mortgage bank Soci  t   de Credit Foncier; Spain supports “cedulas hipotecarias,” its own equivalent of the mortgage bond; and so on.

or commercial, we include them as securities in this section for reasons of continuity. Note that although the total pfandbrief market is very large, the portion attributable to mortgage bonds is confined to approximately 15 % of the total.

Hypotheken pfandbrief are structured to give investors recourse to the underlying pool of mortgages should an issuer become insolvent. In fact, investors receive a preferential claim over the assets and need not be involved in insolvency proceedings (this discussion is still theoretical, since there has never been an event of issuer insolvency). However, investors do not have recourse to the issuer's noneligible assets – only those assets that are ring-fenced in support of a specific issue. Approximately 40 institutions are authorized to issue pfandbrief; the group features a mix of private sector mortgage banks and public sector credit institutions.

We have noted that in the conventional MBS structure mortgages are removed from the mortgage originator's balance sheet. The same does not occur with pfandbrief: mortgages serving as collateral remain on the originator's balance sheet. Indeed, the originator/issuer is responsible for making principal and interest payments to investors; cash flows from the mortgages are only used as a second "line of defense" – suggesting that the originator, rather than the investors, bear prepayment risk. Mortgage collateral pools underlying each issue must conform to very strict parameters, and pool characteristics are quite homogeneous.

Though the pfandbrief market has existed for more than 200 years, it was "repopularized" during the mid-1990s, when banks began issuing jumbos³⁴ (e.g. deals of at least €500 m, extending up to €5 b), which helped concentrate liquidity in a core number of benchmark issues; indeed, prior to the 1990s, the market was highly fragmented, with more than 17 000 listed deals with an average size of only €80 m. Secondary trading occurs in lots of €15 m, with at least three market-makers quoting prices (as a spread to German government bonds); securities must be listed on an exchange, but can be traded OTC. Maturities on the securities can extend to ten years, though the five- to seven-year sector is most active and popular.

Italy

Although Italy's MBS market is relatively young compared to the British, Danish, or German markets, growth has been rapid, primarily as banking institutions consolidate and become increasingly sensitive to the need for efficient balance sheet management. Italian residential MBS issuance has increased steadily since changes in securitization laws were effected in 1999.³⁵ Though the country historically has featured a relatively small number of mortgage products (most of them variable rate, based on EURIBOR, the Associazione Bancaria Italiana prime rate, or the Tasso di Sconto official discount rate), fixed- and fixed/variable-rate structures are becoming somewhat more prevalent, primarily as a result of customer demand. As the move towards a greater amount of fixed-rate products ensues, prepayment risks associated with the actual MBS issues are likely to increase.

Mortgage maturities range from 5 to 20 years, and amortization of principal begins after a short grace period. Approximately 10 % of system-wide mortgages feature government subsidies; the balance is comprised of standard, commercial-rate loans. The market does, however, distinguish between two types of mortgage: standard mortgages (*fondinari*), which secure a specific property and can be used for purchase of, or construction on, that property, and general loans (*ipotecari*), which secure a property and can be used for any purpose. Nearly 90 %

³⁴ Traditional pfandbrief are issued via tap, with each series being an individual issue; jumbos are launched via syndicate, using fixed-price reoffer and book-building (e.g. similar to any Eurobond).

³⁵ These changes related primarily to the establishment of "bankruptcy remoteness" for issuers, assignment, ring fencing, setoffs, clawbacks, and usury.

of Italian residential mortgages are of the *fondari* class, while many commercial transactions are of the *ipotecari* class. The typical *fondari* loan, which features a mortgage charge over the property, is likely to have a maximum LTV³⁶ of 80 % and a relatively small average mortgage balance (e.g. €50 000). Most principal and interest collections are made via direct debit on a quarterly or semi-annual basis. The foreclosure process for delinquent/defaulted borrowers is relatively well established, but is quite time consuming as it involves many formal and informal legal steps; in a worst case scenario, proceedings can take up to seven years to complete, though this is starting to decline.

A typical MBS deal features greater credit enhancement levels than other European products, primarily as a result of legal and servicing idiosyncrasies. In most instances, the purchaser of the mortgage pool and the issuer of the securities tend to be the same party (e.g. there is no use of an intermediate SPE or trust), though that is changing as a result of the establishment of strong ring-fencing laws. Structurally, notes generally are issued with super senior soft bullets with a weighted average life of 18 months; cash flows are reinvested in eligible investments or mortgage assets during the revolving period. Deals also feature mezzanine tranches and subordinated mezzanine tranches, all paying out on a sequential basis. In order to avoid a 20 % tax levy, investors do not receive any payments (principal or interest) for the first 18 months of a deal's life.

Not surprisingly, prepayment rates are relatively stable at less than 5 % per year, due primarily to the fact that the underlying mortgages have shorter terms, are largely variable rate, and often carry some type of prepayment penalty. Credit enhancement is driven by credit rating agency requirements, which tend to be a function of expected loss levels (e.g. probability of default and loss severity), and the potential impact on cash flows. The first loss layer is generally the excess spread account; this is often 75–100 bps in size, which is twice as large as the loss layer in the Dutch and Danish markets, but approximately half the size of the UK subprime market. Deals may also include an unfunded reserve account that is kept by the originator (up to 3 % of deal size, which is again twice as large as other European markets); trapping the excess spread finances unfunded reserves. Subordinated tranches generally amount to 7–10 % of deal size.

Australia

Australia's residential pass-through market has grown steadily since the mid-1990s, and now accounts for more than two-thirds of the country's overall securitized fixed income issuance. Banks and non-banks are both involved in the securitization process. Non-bank financial institutions, which account for a large share of local mortgage origination, have come to rely heavily on the securitization process to fund their operations; indeed, non-banks have a limited ability to warehouse mortgage loans, and must therefore turn to the securitization markets frequently.

The local housing market is characterized by a high degree of home ownership and very low default rates (i.e. approximately 25 % of the US rate); more than 80 % of homeowners have their monthly payments automatically deducted from their bank accounts, meaning payment delays are uncommon. Local tax law does not provide for interest deductibility, so borrowers have some incentive to prepay portions of their principal. Residential mortgages contained in a typical pool are 25–30 year amortizing, variable rate loans with a first lien on the property granted to the lender; no prepayment penalties are levied. The market also features hybrid

³⁶ The market still features significant variations between property assessments and valuations, though this is now beginning to change.

mortgages, which carry fixed rates for the first 1–5 years, and then convert to variable rates; prepayment penalties may be levied in a falling-rate environment. Mortgages often have portability, meaning they can be used to finance a new property after an existing one has been sold; this reduces, by a certain amount, the prepayment risk associated with housing turnover. In fact, prepayments historically have been quite stable as a result of a low level of housing turnover (e.g. 20 % + lower than in the US) and refinancing actions that are driven primarily by bank competition rather than borrower demand; these factors equate to a relatively stable prepayment record in the 20–30 % CPR band.

The Australian market is somewhat unique in promoting both domestic and offshore securitizations. Though monthly-pay domestic A\$ deals, targeted at local institutional investors, generated most activity until the new millennium, quarterly-pay offshore deals have since become more popular. Offshore deals floated in A\$, US\$, or euros typically are directed at European and US investors; securities are listed on an international stock exchange, and some are registered with local securities commissions (i.e. the US SEC) so that they can be distributed more broadly. Ratings on domestic and offshore deals are driven by the distribution and quality of the loans in the pool, as well as excess spread, liquidity facilities, or subordination. Pass-throughs arranged by bank intermediaries feature immediate payment of principal and interest to investors (i.e. no delay, even if the bank faces delays in receiving cash flows from the pool) and are enhanced through mortgage insurance policies, excess spread, and subordination. Mortgage insurance is generally applied based on LTVs: loans with LTVs below 80 % can be included in an overall pool policy, while those with LTVs above 80 % must have their own unique policies.³⁷ Bank pass-throughs generally are issued as floating rate securities with a 10 % clean up call that forces redemption once 90 % of the deal has been amortized. Pass-throughs issued by non-banks are also enhanced through subordination, and often feature call dates and step-up coupons. Non-bank pools typically allow for some degree of mortgage substitution within the pool, as long as minimum pool standards are maintained.

Though the countries discussed above feature the largest international pass-through markets, other national systems feature their own securities. For instance, countries such as Spain, and The Netherlands have introduced their own MBS, in many instances using the successful structural technologies developed by the UK, Denmark, and Germany.

Commercial MBS

The commercial mortgage-backed securities (CMBS) market, which is based on pools of loans used to finance significant commercial projects, such as apartment buildings, office complexes, shopping centers, and industrial projects, has grown very rapidly since the mid-1990s. Intermediaries have found the CMBS market is an efficient way of removing large commercial property loans from their balance sheets, while investors have been eager to purchase the assets as a result of good yields, low default risk, and relatively manageable prepayment risk.

Various types of CMBS pools have developed over the past few years, including small loan conduits (pools of 80+ loans, where no single loan accounts for more than 10 % of the total pool), large loan conduits (pools of approximately 30 loans, where several loans may account for more than 10 % of the pool), fusion conduits (a hybrid of the small and large loan conduits, featuring 80+ loans, including at least one that accounts for more than 10 % of the pool),

³⁷ The rating of the mortgage insurer is vitally important; any downgrade can, of course, threaten the rating of the securities being wrapped – particularly the subordinated tranches, since the insurer's rating is an important driver of the initial size of the deal subordination.

floating-rate conduits (comprised solely of adjustable-rate commercial mortgages), and single asset pools (secured by a single flagship property loan).

Though single-asset structures appear in the market periodically, most CMBS are backed by loans extended to a range of borrowers and geographic regions; this helps ensure an appropriate level of diversification. Pool diversity is often measured through an index value that reflects the degree of concentration; an overly concentrated pool may attract penalties from the rating agencies, requiring additional support in subordinated tranches.³⁸ The nature of the properties being financed also factors into pool construction and evaluation; some classes of property are inherently riskier than others, as a result of their cash flows, tenant profiles, location and function, and so forth. In general, the low-risk class includes multi-family residential, anchored retail, industrial, and office, while the higher-risk category includes unanchored retail, hotel/hospitality, and nursing home/health care. Thus, a pool comprised solely or primarily of loans to unanchored retail and hotels would be considered riskier, all other things equal, than one based primarily on multi-family apartment complexes.

Commercial property loans are structurally different from standard residential mortgage loans. For instance, many commercial loans have standard amortization schedules covering a period of 25 to 30 years, allowing for relatively level repayment of principal. Others, however, are structured with 30 year final maturities and 10-year balloons (i.e. modest principal payments for ten years based on a 30-year amortization schedule, followed by a lump sum payment in year ten), which can skew the cash flow profile of the pool. Balloons also add greatly to refinancing risk: the ultimate repayment of a CMBS depends on the ability of the market to support refinancing of the underlying properties as the balloon dates appear. This refinancing capacity is a function of interest rates, property values/conditions, underwriting criteria, general credit conditions, and specific occupancy rates.³⁹ Importantly, many CMBS feature prepayment restrictions,⁴⁰ such as lockouts, yield maintenance, or defeasance requirements,⁴¹ which lead to lower prepayment risks.⁴² This characteristic makes the prepayment component of CMBS easier to model than comparable agency or private label pass-throughs; indeed, while prepayment risk is the primary risk driver in the residential pass-through market, real estate credit risk is the main driver in the CMBS market.

Commercial loans granted on a recourse basis require analysis of both the borrower and the underlying property being financed. Those granted on a nonrecourse basis demand rigorous analysis of property cash flows (i.e. net operating income, net cash flow, debt service coverage ratio, LTVs),⁴³ the creditworthiness and contractual obligations⁴⁴ of the underlying tenants

³⁸ For instance, the rating agencies may begin levying “penalties” when a supposedly diverse pool holds 20 % or more of its loan in a single state or province.

³⁹ Rating agencies and investors examining refinancing risk in a CMBS focus heavily on stressed market conditions, i.e. an inability to refinance at current market rates. This provides some measure of the downside scenario that may ultimately impact value.

⁴⁰ Prepayment restrictions and penalties generally are eliminated approximately six months before any balloon payment comes due, in order to encourage borrowers to refinance the underlying properties.

⁴¹ The typical defeasance process calls for the borrower to pledge securities to the mortgage holder with cash flows sufficient to meet the mortgage principal and interest, and agreeing to compensate the lender for any loss of above-market yields. This means that the borrower will only prepay the loan if it needs to sell the property. CMBS with a significant amount of defeased loans are often upgraded by the rating agencies, as they are essentially collateralized by high grade securities instead of commercial properties.

⁴² Though prepayments through asset turnover or refinancing may not occur as frequently as in residential pass-throughs, CMBS are still subject to prepayments through defaults. While this is a relatively rare occurrence, most analyses assume annual default rates of 1–2 %.

⁴³ For instance, net operating income is rent revenue less vacancies, fees, and operating costs; net cash flow is net operating income less tenant allowances, capital expenditures, and reserves; and the debt service coverage ratio is net operating income divided by the actual mortgage obligation (or cash flows after costs and expenses divided by stressed refinancing obligations). Each one of these measures helps determine whether the underlying properties will be able to supply the necessary cash flows to service the principal and interest on the CMBS.

⁴⁴ A key element of the contractual obligations relates to lease conditions and terms, lease expiry dates, and any concentration of lease rollovers (e.g. a large amount of leases coming due in the same year). Downside scenarios focus on large re-leasing requirements

Table 4.1 Sample CMBS tranching

Tranche/rating	Subordination %	Issue %
Tranche 1/AAA	20	85
Tranche 2/BBB–	7	10
Tranche 3/BB	2	3
Tranche 4/unrated	n/a	2

supplying the cash flows, and any structural enhancements impacting the deal. In general, the lower a deal's debt service coverage ratios (e.g. net cash flows/principal and interest), the greater the credit enhancement requirement. Enhancement levels may also be adjusted upward/downward by examining the diversity of the pool, delinquencies, and seasoning. In fact, nonrecourse CMBS deals typically rely on senior/subordinated tranching to enhance overall credit quality and protect against any cash flow shortfalls. The average level of subordination in AAA-rated CMBS tranches tends to be rather high, in part because the absolute number of loans in a CMBS pool is much lower than in a standard residential pool. The subordination requirement established by the rating agencies is thus a function of the absolute number of loans in the pool and the size of the largest loan. For instance, a small loan conduit, with good diversification across a large number of loans, might require 0–4 % subordination; subordination in a fusion conduit may increase to 3–10 %, while subordination in a large loan conduit might range from 10 % to 30 %.⁴⁵ Subordination requirements can also increase as a result of concentrations related to geography, or loans extended primarily to the “higher-risk” class of commercial projects noted above. Subordination can also increase in the presence of mezzanine funding (i.e. additional borrower funding above the first mortgage that is secured by equity or a second mortgage loan). Large loan conduits or single-borrower transactions must be viewed in the context of a pure corporate securitization, with attendant risks associated with the single or primary borrower. In order to achieve AAA or AA status, LTVs generally must be capped at 45–50 % and additional overcollateralization may be required.

The goal of the senior/subordinated credit enhancement requirement is to ensure that default by a single large borrower does not jeopardize the entire CMBS structure. It is worth noting that overall subordination requirements have declined gradually over the past decade, as a result of better loan underwriting standards, decreased leverage, and a greater focus on less volatile properties.⁴⁶ Table 4.1 illustrates a sample large loan CMBS structure with subordination levels.

Pool payment characteristics in recourse CMBS reflect a heavy reliance on the cash flows from the properties, and only a secondary reliance on the credit quality of the issuer. Actual default experience has been very positive over the years as a result of the multiple levels of security, including cash flow from the underlying properties, credit strength of the issuer, and/or structural enhancements designed to provide additional support. Seasoned CMBS tend to be in good demand among the investor base, particularly in areas where commercial property has appreciated in value.

at below-market rates, which can strain cash flow and create vulnerabilities in the servicing of the CMBS. In general, multi-tenant properties with large and reputable anchor tenants are considered best for mortgage loans being included in a pool.

⁴⁵ To obtain a specific credit enhancement at an A, AA, or AAA level, the probability of default and expected loss within the pool must first be estimated; the result is then geared up to achieve the desired target rating. If the loan/pool varies from the default and expected loss assumptions, then the subordination level is adjusted by an add-on factor.

⁴⁶ Bank lenders generally are unwilling to finance properties that do not have a significant amount of high-quality leasing/preleasing in place. They may also curtail their activities when a market is entering a speculative phase, or when vacancy rates are on the rise (e.g. 12–15 %+).

A master servicer is used in a standard CMBS structure to collect monthly principal and interest payments from the commercial borrowers represented in the pool. These cash flows are forwarded to the trustee, who disburses them, net of fees, to the CMBS investors in proper sequence. If any loan in the pool is in default, a special servicer determines whether to liquidate the underlying property securing the loan, or restructure the loan; the end goal is to preserve as much value in the collateral pool as possible, so that investors can receive the largest percentage of principal and interest due, even in a default situation. In addition to the credit rating tranching noted above, a CMBS issue may feature short- and long-term maturity tranches and, on occasion IO tranches (which we discuss later in the chapter). The long-term tranches effectively are protected from default-related prepayments by the shorter-term securities and IOs.

Seasoned deals benefit from at least two forms of call protection: at the loan level (e.g. with prepayment penalties and lockout periods) and via the sequential pay structure (e.g. the AA cannot be paid down until the AAA is retired). This leads to less contraction risk and lower negative convexity than might be found in standard residential MBS. Secondary trading in benchmark CMBS tends to be relatively active, particularly when the overall commercial property market is robust. Top CMBS tend to feature good value relative to similarly rated corporate bonds. However, the real estate cycle – which can include downturns and troughs that can span many years – means long-term relative value performance is still largely unproven.

International CMBS

The CMBS market is dominated by US mortgage pools. That said, the European CMBS market has been expanding in recent years, and now features activity related to multi-borrower loans (conveyed into an SPE or trust via true sale or a synthetic mechanism)⁴⁷ and single loans from one originator (or in relation to a single line of business, such as healthcare facilities or pubs). Multi-borrower structures rely primarily on cash flows from the underlying property to service obligations. Each loan is analyzed independently, as in the US market, but the entire structure benefits from the number and variety of loans. Significant reliance is placed on LTVs, debt service coverage ratios, leasing schedule (and tenant quality), seasoning, concentrations, property type, cash flow schedule, and servicing. Analysis tends to be market-specific to take account of local requirements. For instance, in Italy, France, and Spain, lease periods are shorter, but foreclosure periods are longer, than in the UK.

The UK features sterling-denominated CMBS deals that are structured in a form similar to those in the US (i.e. multi-borrower conduits or single-asset transactions,⁴⁸ with some amount of senior/subordinated tranching). Commercial leases and loans underlying UK CMBS have 15–25-year terms with five year rolling upward rent adjustments (meaning rents rarely, if ever, decline). Unlike the US market, most commercial mortgages are based on full amortization of principal, rather than balloon payments; this means that the primary analysis is based on the credit quality of the tenant and the value of the property, rather than the refinancing risk arising

⁴⁷ The European marketplace also features a secured loan structure, where the borrower keeps ownership of the assets but investors obtain a lien on the assets. The SPE raises funds and lends to the borrower, who can then pass the funds through to other operating companies. Security thus includes a first lien over the borrower's property and assignment of cash flows (which are channeled through a fixed-charge account).

⁴⁸ Single-asset transactions, often based on large property developments, have proven quite popular, with ratings based on the credit of the underlying borrower or a third party guarantor.

at the end of the balloon period.⁴⁹ Most UK CMBS deals feature liquidity facilities that help ensure timely payment of principal and interest on nondefaulted mortgages; this is necessary as the servicer does not typically commit to advancing cash flows in a timely manner if they are not received as scheduled from the borrower.

Collateralized mortgage obligations

As noted earlier in the chapter, CMOs originally were developed in order to address the specific risk requirements of investors. Though many investors are willing to bear the prepayment risks that are characteristic of MBS, others are not. The CMO mechanism allows prepayment risks to be redirected to various tranches, permitting investors to select a security with a preferred risk/return profile. CMOs have proven popular with investors needing to inject greater certainty into the management of their asset–liability flows. Absent this ability, the contraction and extension risks described earlier can complicate the process. For instance, an asset manager may fund a ten-year MBS with ten-year liabilities; if interest rates fall, shortening the MBS maturity to only seven years, the asset manager will be forced to reinvest for the remaining three years at lower rates. If, in contrast, the MBS can be converted into a ten-year mortgage-related asset with a true ten-year life, the asset manager will be able to run a balanced portfolio for the full ten-year horizon. The CMO achieves this goal by bundling together pools of MBS, and then retranching them into individual securities with different prepayment characteristics. Investors seeking little or no prepayment risk can choose a tranche that is not particularly sensitive to changing rates; by doing so they will, of course, receive a lower return. Those seeking full exposure to prepayment risk can select a tranche with greater exposure, receiving a higher return in exchange. The design of the CMO creates greater cash flow stability over a range of prepayment speeds for investors seeking low-risk securities; this stability creates positive convexity (rather than the negative convexity characteristic of MBS), and allows for more accurate asset–liability management. In addition to the obvious benefits that can be achieved by creating securities with customized prepayment characteristics, the issuer of a CMO may also be motivated to create a transaction if it can purchase the underlying MBS (or other mortgage assets) at a discount to the theoretical value; this creates the possibility of arbitrage profits.

A CMO generally is formed as a trust in order to avoid double taxation issues.⁵⁰ The MBS (or other mortgage assets) used to structure the CMO must be accounted for as an “outright sale of assets” to the trust in order for the securities to be removed from the issuer’s balance sheet. One or more of the rating agencies explicitly rates most CMOs. CMO ratings are a function of the quality and diversification of the underlying pool of MBS and the nature of the trust structure and tranching mechanism; the credit quality of the CMO issuer, which establishes and manages the trust, typically is not a factor. In order for a CMO tranche to gain AAA ratings, the cash flows must be sufficient to service the tranche under a range of prepayment scenarios, including those with a low probability/high impact. CMOs, like other nonagency mortgage products, rely on subordination, overcollateralization, excess servicing

⁴⁹ Commercial tenants are responsible for the cost of repairs/insurance, operations, and taxes, as well as returning the space to its original condition. If the lease is assigned to a third party, the same terms and conditions are conveyed to the new tenant; this makes the underlying cash flows on the property quite stable and certain.

⁵⁰ If the CMO were created as a corporation, it would pay taxes on the mortgage cash inflows, while investors would pay taxes on their inflows from the CMO. It is important to note that in the US, some CMO issuers/originators were unable to pass through all cash flows from the CMO collateral on a tax-free basis to investors when they issued debt securities. This led, in 1986, to the creation of real estate mortgage investment conduits (REMICS), or trusts that could pass through all cash flows to investors without first being taxed. Some CMOs were thus created as REMICS rather than standard CMO trusts, but essentially operated in the same manner.

spreads, and/or cash reserve funds to enhance credit quality of specific tranches (e.g. the AAA tranches).

CMOs can be regarded as structured assets that are created from other structured assets, and are designed to cope with aspects of contraction and extension risk. MBS prepayment risk is reallocated through a securitization tranching mechanism; in fact, the template has proven so popular that it has been used to create other structured assets, including the ABS, collateralized debt obligations, and insurance-linked securities that we consider later in the text. Naturally, in any retranching mechanism, some investor or intermediary ultimately must bear the bulk of the risk that is redirected from other tranches. The accrual bond, also known as a residual or z-bond, typically absorbs this. This point is worth stressing: the CMO does not eliminate prepayment risk; it simply reallocates it to the holder of the accrual bond. If no buyer can be found for the accrual bond, the intermediary arranging/issuing the CMO will be required to retain the asset (or cancel the proposed transaction). The CMO accrual bond represents the excess cash flow from the underlying MBS collateral – over and above what is needed to pay interest, principal, and expenses on other tranches. It can be considered an uncertain “stub” with a value that depends ultimately on rates, the coupons on the underlying collateral, the type of collateral, and the overall structure of the CMO. The main sources of residual cash flow include premium interest (the difference between the coupon on the collateral and the coupon on the highest rated tranche in the CMO), reinvestment income, and bond coupon differentials. The accrual bond is very sensitive to rate changes: when rates fall and prepayments accelerate, CMO tranches with lower coupons pay off and reduce the cash flow of the accrual coming from premium income, bond coupon differentials, and reinvested cash flows; the reverse occurs with rising rates. The accrual bond can thus be used as a form of hedge.

A CMO is comprised generally of pools of pass-through securities, though it may also be created using whole loans and even mortgage strips. Each tranche (or class) may have a unique stated maturity, with principal returned from the underlying pool assets on a predetermined basis. The earliest CMOs were created as sequential-pay securities, using multiple tranches and an accrual bond. Using a two-tranche sequential with an accrual bond as an example, the principal and interest cash flow from the underlying pass-through pools are directed as follows: interest to tranche 1, interest to tranche 2, principal to tranche 1, principal to tranche 2, and principal and interest to the accrual bond. It is easy to see through this simple example that tranche 1 investors receive cash flows before tranche 2 and accrual investors, suggesting less exposure to contraction and extension risks. The accrual investors, in turn, do not receive any cash flows until the other classes have been retired, meaning full exposure to prepayment risk. The senior tranches are de facto protected by the accrual bond. The CMO accrual bond has the longest stated maturity in the structure, as it is the last to receive principal repayment; when the accrual bond is in the accrual phase, it accumulates deferred interest that can be used to accelerate the amortization of shorter-maturity tranches. Accordingly, the larger the size of the accrual bond, the greater the ability to shorten the amortization period on the remaining tranches; this allows significant (or even total) reduction in extension risk on the protected tranches. Figure 4.7 summarizes the generic CMO structure.

As a result of this structural design, the CMO effectively creates MBS with varying maturities and return profiles. Even when prepayments slow, investors holding the top one or two tranches of the structure are invested in relatively short-term assets, as noted in Figure 4.8.

Originators can often generate an arbitrage profit between the collateral and the financing cost, as a result of the cash flow certainty and liquidity of the most stable CMO tranches; though lower-rated/junior tranches carry a high financing cost, the stable, senior-rated tranches

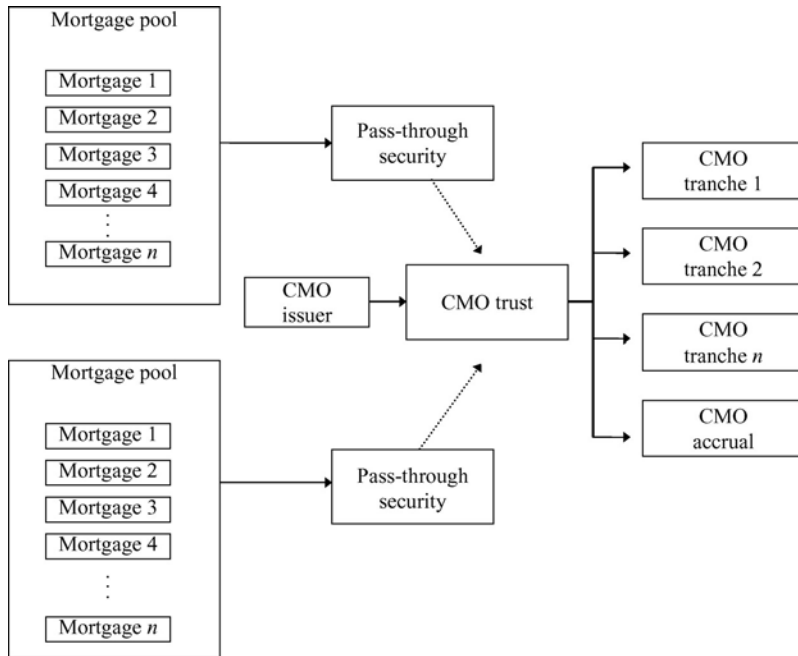


Figure 4.7 Generic CMO structure

carry a far lower cost and generally account for a larger portion of a deal, which leads to profit opportunities. Consider a simple example of a four-tranche CMO based on 8 % CMO collateral. The average yield (cost to the CMO originator) of this structure is 7.50 %; since the yield on the collateral is equal to 8 %, a profit opportunity arises. The tranches, and their relative yields, are depicted in Table 4.2.

Following the initial success of sequential CMOs, new variations appeared, and the sector now supports a wide variety of structures.⁵¹ In addition to sequentials, common CMO structures include floating/inverse floating CMOs, planned amortization class CMOs, targeted amortization class CMOs, and CMO strips.

Floating-tranche CMOs, introduced in 1986, are LIBOR-based securities that feature maximum lifetime rate caps to ensure sufficient cash flows to meet investor obligations. Inverse floating CMOs feature the reverse payoff, and are often issued with floors. In fact, floating and inverse floating tranches are often included in the same CMO structure, in order to ensure that the coupon income from the collateral is sufficient to cover all payments (and eliminate the need for rate caps, which many investors find unappealing); in fact, the two elements must be structured so that they are always equal to the fixed coupon on the collateral.⁵² Floater/inverse floater tranches can be leveraged to provide investors with additional risk/return opportunities. For instance, if a \$100 m CMO issue is launched with an 80/20 split between the floating and

⁵¹ Though not as many, it should be noted, as in the mid-1990s, when the financial creativity of intermediaries led to the development of some particularly esoteric, and risky, tranching structures that ultimately proved unpopular in rapidly rising/falling-rate environments. Some CMOs, launched during the peak of the cycle, featured dozens of very complex and esoteric tranches.

⁵² The sum of a floating CMO tranche and an inverse floating CMO tranche is equal to a fixed-rate bond with a coupon that is set below the average coupon of the collateral pool.




Figure 4.8 Effect of slowing CPR on principal return. (a) CPR base case, (b) CPR slows

Table 4.2 Sample CMO transaction

Tranche	Rating	Percentage of deal	Yield
1	AAA	93.00	7.00
2	AA	5.00	8.00
3	BB	1.50	15.00
4	Unrated	0.50	25.00

inverse floating tranches, a 1 bp increase in LIBOR on the floater leads to a 4 bp decrease in the inverse floater (other combinations are, of course, possible).⁵³ Superfloater tranches are a further variation on the theme, and are structured typically as floating-rate securities with an initial coupon set below current LIBOR, which changes as a multiple of LIBOR – providing a much higher yield than a conventional floater in a rising-rate environment, and a much lower yield in a falling-rate environment.

The planned amortization class (PAC) bond and its associated companion (or support) bond are key instruments of the CMO market. The PAC emerged as a logical evolutionary

⁵³ It is worth noting that an investor purchasing an inverse floating-rate CMO is effectively purchasing a fixed-rate planned amortization class bond and shorting a floating-rate note.

step, creating prepayment protection for certain tranches within the CMO structure. PACs, sometimes known as simultaneous pay bonds, are securities featuring a highly predictable cash flow stream (and hence value), as long as prepayments remain within a defined band. The principal payable on a PAC bond has absolute cash flow priority over the associated companion bonds, which exist in order to absorb prepayment risk above/below the stated prepayment speed bands (i.e. the companions absorb the contraction and extension risks). Companions are extremely sensitive to changes in prepayment speed, and their coupons generally are set well above those on the main PAC tranche.

The PAC redemption schedule is equal to the minimum value produced by the two PSA curves that define the prepayment range; the upper and lower prepayment speed bands of a PAC are known as an initial collar. For instance, a PAC may feature 100 % and 200 % PSA CPR bands, meaning that principal payments from the underlying collateral will be directed to the PAC investors while speeds remain in the 100 %–200 % band. As rates fall, prepayments increase, and the companions are paid off, the degree of protection afforded to the PAC investors changes; the changing upper and lower bands, known as the effective collar, can vary during the life of the transaction, based on rate movements. PACs generally are characterized by a greater degree of contraction, rather than extension, risk, as the lower band is unlikely to be breached for long, while the upper band can be breached for an extended period when rates drop suddenly. If a very high prepayment rate is sustained for a long period of time, the protected range can disappear; this occurrence, which is unusual, causes the PAC to become “busted.” When prepayment speeds are fast (and enduring), the companions amortize quickly; once they mature, the prepayment collar ceases to exist – meaning that any outstanding PACs convert into sequential.

A refinement of the PAC structure groups several PAC tranches within a single CMO, each with its own entry point. This means that the overall PAC collar is subdivided so that individual PAC bonds can be created, i.e. short-term PAC and long-term PAC. A Z-PAC can be created by developing PAC tranches that pay off according to a set sequence (these are often of interest to zero coupon bond investors). In addition, the premium interest from a CMO structure can be stripped to create a separate interest only (IO) tranche, while the bond coupon differentials can be stripped off to create a PAC IO tranche.⁵⁴ Companions, in turn, can be created in floating and inverse floating form; they may also be issued with their own prepayment schedules: so-called PAC II bonds, which are companions with prepayment schedules that are supported by other companions without schedules. Each nested layer of PAC bonds within an overall PAC structure provides for increasingly narrow bands of prepayment protection. Thus, the primary PAC (i.e. PAC I) may have a prepayment collar of 100–300 % PSA, while a second PAC (i.e. PAC II) may provide protection between 150–250 % PSA. Companion bonds accompany the structure and absorb any additional risk.

The targeted amortization class (TAC) bond is a further refinement of the general PAC structure. Indeed, a TAC is identical to a PAC, except that it features asymmetric prepayment protection through the use of a single prepayment speed target, rather than a band of speeds. This means that the TAC protects investors against contraction risk, shifting prepayments to the TAC companion if prepayments rise. As prepayments exceed the speed target, the WAM of the TAC contracts less than that of a standard CMO sequential, since the TAC companion absorbs

⁵⁴ Prior to tax law changes in 1992, a CMO tranche could not be created without a principal component, meaning some deals were issued with a very small principal component and a very large coupon component; these “IOettes” behaved much like standard IOs. Following the tax law changes, a CMO tranche could be issued without a principal element, leading to the gradual replacement of IOettes with standard IOs.

the risk. However, if prepayments slow, the TAC extends and performs in the same manner as a sequential. Reverse TACs, in contrast, do the opposite, protecting against extension risk by diverting the risk to the companions as prepayments decline. TACs and reverse TACs can also be embedded in broader PAC structures as subtranches. Returning to our example of a PAC I tranche covering 100–300 % PSA, and a PAC II tranche covering 150–250 % PSA, a TAC tranche might be added to provide protection at 200 % PSA.

Very accurately defined maturity (VADM) bonds are securities that are paired with accrual bonds; interest accruing on the accrual bond is used to pay principal and interest on the VADM tranche, providing the necessary level of extension risk protection if prepayments slow (i.e. interest on the accrual bond will be sufficient to pay the VADM, hence the accuracy of the maturity). However, if rates fall and prepayments accelerate, the accrual bond pays faster, meaning the VADM shortens. In fact, the VADM is very similar to a reverse TAC, but it provides even greater protection against extension risks.

Stripped mortgage products

The mortgage strip market evolved naturally from the government bond strip market we described in the last chapter. Once intermediaries (and government agencies) discovered the potential value that could be derived from stripping government securities, they applied the same technology to other assets, including mortgages. FNMA was the first US institution to launch stripped mortgages, with an inaugural issue in 1986. The FNMA deal segregated the principal and interest streams from FNMA MBS into separate tranches, allowing investors to purchase the desired PO or IO components. Others soon followed, and the mortgage strip market grew rapidly. The market now includes several different product classes, including conventional MBS IOs and POs and CMO strips. As with other strips, MBS strips are comprised of two different securities: one class that receives all principal flows from the MBS pool, and a second class that receives all interest flows. The two can be recombined to recreate the underlying MBS. In fact, arbitrage forces ensure that prices of the individual components trade at a level that is equal to the price of the underlying MBS; when discrepancies appear, reconstitution activity accelerates until the arbitrage spreads are eliminated. Most MBS strips are issued via specialized trusts that exist explicitly to issue the two classes of securities and manage the attendant cash flows; some, however, are issued as structured IOs and POs through the CMO tranching mechanism described above.

We know from the last chapter that the PO is a deep-discount security, redeemable at par. The mortgage dimension means that effective yield depends primarily on prepayment speeds. Specifically, as rates fall and prepayments accelerate, the effective yield on the PO rises. Let us assume the most extreme scenario, which is instantaneous prepayment occurring immediately after the investor purchases the PO strip at a discount: the investor presents the PO strip for redemption at par, capturing the price differential between par value and the discounted purchase price; this price gain can be converted into an effective yield (which is extremely large). The rising-rate scenario generates the opposite result: if rates rise sharply, prepayments will slow dramatically, and the principal on the mortgages will remain outstanding for many years. The investor's discounted PO position eventually will redeem at par, but perhaps not for 10, 20, or 30 years (depending on the composition of the mortgage pool and the realized prepayment level). Effective yield amortized over the 10-, 20-, or 30-year period will thus be very small.

The IO, in contrast, is purchased without a specific par value; the investor is entitled only to the interest coupons on the outstanding principal balance of the MBS, and therefore hopes

Table 4.3 MBS PO and IO rate/value relationships

Security	Interest rates	Prepayment speeds	Strip value
MBS PO strip	↑ ↓	↓ ↑	↓ ↑
MBS IO strip	↑ ↓	↓ ↑	↓ ↑

for rising rates and slow prepayments. As rates rise, prepayments slow and the pool extends, causing the IO to increase in value as a result of the continued flow of coupons. As rates fall, prepayments rise and the pool contracts, causing the IO to decrease in value as the flow of coupons shrinks. The IO is therefore somewhat unique among fixed-income securities in reacting to rate movements.⁵⁵ While conventional fixed-income securities (and even structured assets such as POs) increase in value as rates fall, and decline in value as rates rise, the IO does precisely the opposite. This means that IOs can be used to hedge certain positions against adverse interest rate movements. It is important to stress that both POs and IOs tend to exhibit greater price volatility than conventional MBS, as a result of their unique structural features. The strips are much more sensitive to prepayment changes than the underlying collateral. The rate/value relationships of IOs and POs are summarized in Table 4.3.

As noted in the preceding section, CMOs can be structured with strip tranches of their own. These CMO strips function much as standard MBS IOs and POs, i.e. coupon or principal cash flows from the underlying CMO collateral are redirected to investors holding the relevant securities. Common CMO strip tranches include conventional CMO IOs and POs (e.g. standard redirection of interest or principal from the CMO collateral), PAC/TAC IOs (e.g. bond coupon differential flows), and Super POs (e.g. leveraged PO strips). As with other strips, CMO POs benefit from falling rates, while IOs benefit from rising rates.

We have already indicated that certain CMBS issues are structured with IO tranches. CMBS IOs receive excess interest after coupons and fees from all other tranches have been paid; the longer the senior and subordinated tranches remain outstanding, the greater the potential cash flow stream to the IO tranche. CMBS IO cash flow uncertainty centers on prepayments, defaults, and refinancing extensions. Prepayments on commercial mortgages are small compared to residential mortgages as a result of yield maintenance, lockouts, and defeasance; this means there is little possibility of negative convexity. Defaults, in contrast, are far more important; if a default occurs, senior tranches receive their principal ahead of schedule, meaning the IO tranche loses its opportunity to earn additional income. In fact, the most severe scenario is based on a widespread real estate downturn that leads to a large number of defaults. Refinancing extensions also play a part in the IO cash flow stream. Since many commercial mortgages are structured with balloon payments, borrowers must arrange for alternate financing prior to reaching the balloon date. If they are unable to do so, the servicer may agree to extend the loan for a certain period, which allows the IO tranche to continue receiving interest. Since CMBS IOs are a relatively recent creation, there is still insufficient data by which to gauge performance. This means investors must make assumptions on prepayments, defaults, loss severity, and default timing in order to assess potential value. Unfortunately, applying such

⁵⁵ IOs also feature negative convexity: as interest rates fall and prepayments rise, the higher speeds overtake the lower discount rates creating negative duration, which continues until prepayments level off. As rates rise and prepayments fall, the slower prepayments outweigh the higher discount rates, causing positive duration.

assumptions to an overall pool may be overly simplistic and insufficiently accurate, while applying them to each individual loan can prove computationally onerous. Accordingly, some “middle ground” may be preferred (e.g. grouping “similar” loans into mini-portfolios and applying relevant assumptions to each one).

4.3.2 Asset-backed securities

ABS are securities backed by cash flows from specific assets, including receivables and certain types of loans. The ABS sector generally is defined to exclude mortgage assets and traditional bond/loan credit assets, which are the domain of the MBS and collateralized debt obligation (CDO) sectors. In contrast to the MBS and CDO markets, ABS pools contain assets with relatively short maturities and some element of credit risk, but relatively modest (if any) prepayment risk. An ABS issue can be structured with fixed- and floating-rate tranches in order to appeal to the widest base of investors. Deals may also be floated in multi-currency form, via public registered or privately placed securities.

Legal and structural issues

A proper legal structure is essential in order to limit liability and bankruptcy risk on an ABS issue. Indeed, the legal and structural issues surrounding ABS deals are similar to those characterizing MBS, including use of bankruptcy-remote entities to isolate the issuer of securities from the originator of the underlying assets. In the credit card ABS market, for instance, receivables are sold by the originator to an SPE or master trust via true sale – this isolates the assets from the originator’s balance sheet and demonstrates that the seller/servicer no longer controls the pool. Assets are sold to a bankruptcy-remote SPE/trust (which acts as trustee) on a nonrecourse basis, with a pledge of cash flows and a perfected security interest granted to investors; in some legal jurisdictions, this involves filing of a lien. Most transactions are accompanied by legal opinions reflecting true sales and security interest status; once complete, the investor is protected and the issuer is entitled to remove the assets from its balance sheet. Note that in some cases assets are originated directly into a master trust, eliminating the need for a legal transfer.

The ABS market generally, and the credit card sector specifically, has evolved considerably over the past three decades from a legal/structural perspective. The earliest transactions (i.e. the mid- to late 1980s), were backed by letters of credit from leading banks; while this was a straightforward approach, it also meant that individual securities were exposed to event risk associated with a downgrading of bank ratings. The next iteration (i.e. early–mid 1990s) was based on the common senior/subordinated tranching we have already described, with lower-rated tranches serving as the first loss buffer for the AAA-rated tranches. The lower-rated tranches, requiring enhancement of their own, were often backed by a cash collateral account or a pool of short-term liquidity from a leading bank. During this stage, issuers began switching from standard single-deal SPEs to more flexible master trusts, which allow issuance of multiple series of securities from a single vehicle over an extended period of time. Master trusts also permit static or dynamic pool structure – static pools preserve the original account structure, while dynamic pools allow new accounts to be added. Master trusts are now the norm in the marketplace as a result of their flexibility. In the next evolutionary stage (i.e. the late 1990s), the senior/subordinated cash collateral account was replaced by BBB/unrated first loss securities (known as collateral invested amount [CIA] certificates), which proved more cost effective; this approach, however,

featured limitations related to trading restrictions and the maximum number of eligible investors.⁵⁶ In 1998, the CIA certificates was replaced by the Class C note – which is essentially a BBB tranche without trading restrictions. The investor and seller tranches carry a beneficial interest in the pool of assets through this structure; this represents shared ownership, or an undivided interest, in the pool. We consider the structure in greater detail in the next section.

Though the servicing function is important in the MBS market, it is essential in the ABS market, since the originator of the asset and the servicer of the asset are virtually always two distinct parties (which may, or may not, be the case for MBS, as loan originators may sell their loans to the pool but retain servicing rights). It is vital, in any ABS structure, to evaluate the strength of the servicer, including its access to liquidity and its ability to deal with crisis situations, where more intensive collection procedures are required; this can often be done by examining treatment of past collections, delinquencies, charge-offs, and so forth. Not surprisingly, the abilities of the servicer are particularly critical when dealing with subprime assets. In addition, the trustee of any ABS deal must develop a backup servicing plan in the event that the original servicer is unable to perform the function.

Credit card securitizations

Credit card securitizations form the single largest component of the ABS sector, and have been integrated successfully into the US and European markets. Credit card deals first appeared in the US in the early 1990s, and have expanded steadily since that time. Much of the securitization business is driven by banks, the largest of which have become significant credit card originators (e.g. the top ten US banks control 75 % of the credit card market).

The standard credit card ABS is comprised of pools of unsecured obligations that individuals and other companies owe bank card issuers (e.g. banks operating under Visa/Mastercard) or retail, travel, or private label card issuers. Each pool generates cash flows through fees, charges, and fixed or floating interest payments; the servicer is responsible for managing all of the cash flows associated with the portfolio. Customers are billed monthly, but need not pay an interest/finance charge if their outstanding balances are paid fully during a stated grace period; if a balance remains, the interest charge is applied to the balance, which becomes due and payable during the next monthly billing cycle.

Each credit card pool can be described by its monthly payment rate (MPR). The MPR is the monthly principal, interest, and fees outstanding as a percentage of an outstanding pool of credit card receivables, and indicates how quickly the receivables base can liquidate the pool balance; MPRs tend to range from 8–20 % monthly, though this depends on the specific market and the nature of customer commitments and economic strength. As MPR declines to low levels, extension risk increases and the probability of insufficient cash flow to repay principal increases.

In addition to MPR, each pool features a pool yield, delinquencies, excess spread, and purchase rate. The pool yield is a function of interest income and annual card fees, less credit losses from defaulting payments. Delinquencies are late payments within a pool; past-due balances are often a sign of potential future defaults, and accounts that reach defined past-due thresholds (e.g. 90+ days) generally are written off by the trustee as credit losses.⁵⁷ Excess

⁵⁶ In order not to create tax disadvantages for the master trust.

⁵⁷ In fact, an analysis of payment patterns in a given pool is important, as better quality credits often pay off more quickly than weaker credits; the remaining assets in a pool may thus be skewed towards poorer credits once the revolving period ceases and the pool seasons. Stress tests and charge-off analysis can provide some insight into this phenomenon.




Figure 4.9 Revolving/amortizing periods in a credit card securitization

spread is the balance remaining after servicing fees, interest expense, and credit loss are subtracted from the pool yield. Repayment patterns, which are embedded in the MPR, are a critical component of the analysis process. For instance, a pool that features a large number of customers that pay off their balances fully each month will feature a higher MPR and greater cash flow than a pool with a majority of customers that only pay the minimum requirement each month. The purchase rate measures the amount and timing of new receivables purchases by the servicer. The servicer's purchases, in turn, are impacted by the start of the accumulation period (i.e. the end of the revolving balance period) and the trust's termination date.

Credit card ABS differ from MBS in their sequencing of cash flows. We have already noted that mortgage pools feature amortizing payments (though certain commercial pools may have a portion of their cash flows allocated to balloon payments). Credit card pools, in contrast, have no principal amortization during an initial lockout period, which can cover from one to ten years. During the lockout period, the trustee accepts payments from the credit card borrowers in the pool, and reinvests in new receivables (at a particular purchase rate) so that the absolute size of the pool remains above some minimum level; the only cash flow payable to investors comes from interest charges and annual fees. The process continues until the end of the lockout period, at which point further "revolving" ceases; the trustee then begins accumulating monthly repayments in a master account in order to repay investors their principal. Once the revolving period has finished and investor payments commence, principal-related cash flows effectively amortize, as in any mortgage-related transaction.⁵⁸ Figure 4.9 illustrates this process. Note that some transactions feature early amortization triggers, where principal is returned sequentially to investors during the lockout period if the excess spread falls to 0. In addition, while certain pools feature a static set of customers with fluctuating balances during the revolving period that correspond with spending and repayment patterns, some master trust pools allow for the addition or removal of customer accounts during the revolving period.

⁵⁸ Note that some amount of the balance (e.g. 5 %) may be set aside specifically to resolve any charge-related disputes that might arise during the payment period.

Though most credit card ABS floated during the 1980s through the early millennium were based on the lockout/amortizing structure noted above, the “soft bullet” structure was introduced in the late 1990s, and has since become more prevalent. The soft bullet structure introduces a new security that is able to absorb amortizing cash flows prior to a bullet date, and generate sufficient cash flows to meet the bullet payment. This essentially allows the creation of new variable notes that are paid off until each bullet security is redeemed.

Individual tranches of securities can be impacted by differences in coupon interest expense and cash flow timing; fixed and floating tranches tend to be launched by the master trust at different times, meaning some tranches have more, and others less, excess spread. Note that although payments from the pool occur every month, coupons on the tranches may only occur every quarter or six months. Interest finance charges from the pool flow to the master trust, which then reallocates them to individual tranches on the basis of each coupon (so-called excess sharing, which is common, or through a pro rata sharing, which is less common). The excess sharing mechanism allocates cash flows based on the proportion contribution of the tranche to total interest; low rate coupons can thus generate flows for higher-rate coupons.⁵⁹ Once this has been completed, principal collections from the pool are allocated to the tranches. This occurs typically in proportion to the outstanding balance, with any excess principal collections used to offset deficiencies for those that are in an amortizing, rather than revolving, period.

In a typical credit card ABS structure the master trust issues a AAA tranche, a subordinated tranche (that is typically rated A), and an equity/seller’s interest; a separate collateral interest stream (itself comprised of a subordinated receivables interest and excess spread) is distributed directly to a separate owner trust,⁶⁰ which then uses the collateral interest to float the Class C note (which is generally rated BBB). All principal and interest on the Class C notes is paid from the collateral interest; any remaining cash flow is considered excess spread and flows to the reserve account and/or the seller. The reserve account is used to protect the Class C notes from collateral deterioration. If the account reaches a particular level, no further funding is required, and any excess cash flow can be redirected to the seller; naturally, if the reserve account is not funded fully, the seller is not entitled to any cash flows. Indeed, the lower the excess spread,⁶¹ the longer it takes to fund the reserve account. Note that the Class C tranche typically features a soft bullet, though it may also amortize or extend under certain scenarios.

Figure 4.10 illustrates the typical flotation structure.

Though tranche sizes are deal-specific, a generic credit card ABS might feature a AAA-tranche equal to 85 % of the deal, a subordinated A-tranche of 7–9 %, and a Class C note of 6–8 %. The AAA and A tranches may be floated on a public or private basis; all Class C notes issued through the early years of the new millennium have been arranged on a private basis. The Class C notes feature maturities of 7–10 years and coupons that are either fixed or floating. The securities tend to be rather illiquid, as a result of their private placement and size characteristics, though they have some crossover appeal to institutional buyers of BBB-rated corporate bonds.⁶²

⁵⁹ Specialized finance charges also exist, where the servicer allocates the interest-related cash flows for all tranches as a single pool, on a weighted-average basis.

⁶⁰ The owner trust is used in order to neutralize any tax effects – since the owner trust is not a publicly traded partnership, it is not subject to taxation.

⁶¹ Comparison of the excess spread across competing ABS deals is important, as losses and yields can vary greatly between transactions.

⁶² Any relative value analysis must focus on pool yield, excess spread, payment rate, and purchase rate, along with credit default scenarios. This can reveal possible yield pickup versus BBB-rated corporates.

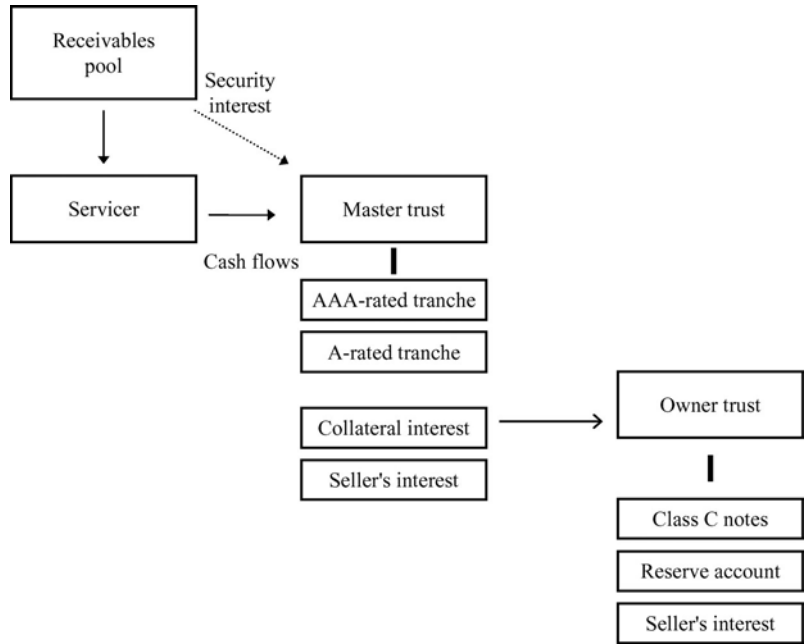


Figure 4.10 Credit card ABS master trust/owner trust structure

A credit card ABS using the structure described above effectively establishes a loss sequence or “waterfall” (which we discuss at greater length in Chapter 6). If a pool is subject to credit losses, the excess spread absorbs the initial amounts, followed by the reserve account, the Class C notes, the subordinated tranches, and the AAA tranche, in that order. The probability of the AAA tranche being exposed to any loss is, of course, extremely small. If losses are especially large, the reserve account may not be sufficient to protect the Class C notes from principal losses. The risk/return profile of the Class C notes and the subordinated tranche depends on the size and pace of assumed credit defaults.⁶³ In general, the rating of Class Cs is a function of historical collateral performance, concentration, legal structure, industry competition, and interest rates. Any excess spread that remains at the conclusion of a transaction may be made available to subordinated investors and the holder of the seller’s interest; the senior investors receive no participation.

Home equity loan securitizations

The home equity loan (HEL) market has expanded in recent years in terms of scope and size. HELs originally were designed exclusively as second mortgages, by which borrowers could access their home equity. In recent years, HELs have grown to include first lien lines of credit for subprime borrowers (often used as a debt consolidation mechanism). Accordingly, when a pool of HELs is securitized, the end result is an asset that can be considered an MBS or ABS. We include our discussion of HEL securitizations in the ABS section in order to preserve continuity with “credit line” securitizations.

⁶³ It is common in any analysis process to hold the pool yield constant while stressing monthly payment rates; this provides an indication of the amount of losses a Class C note can sustain.

The market for HELs is centered primarily in the US, where homeowners frequently borrow against accrued equity in their homes. In fact, the market has grown rapidly since the mid-1990s, and is now approximately equivalent in size to the consumer revolving credit sector. A HEL facility can be used as a revolving line of credit (i.e. drawn down and repaid repeatedly, much as in a standard revolving credit card), a refinancing vehicle (i.e. to take advantage of a decline in rates or improvement in credit quality), or a form of bridge loan (i.e. to purchase another home). The earliest HEL securitizations were arranged by specialty companies that borrowed funds, originated home equity loans, and then sold the loans into a securitization conduit. Deals in the new millennium are often arranged by banks and securities firms, which structure them in the same fashion as other MBS and ABS deals.

HELs can be viewed as lower-rated Alt A loans, and are characterized by similar valuation and prepayment profiles. HELs often feature interest rates that are somewhat lower than comparable revolving credit cards, but well above those characterizing the underlying home mortgages. Pools of HELs are assembled on a geographically varied basis, in order to provide for proper economic and financial diversification; this is particularly critical in the subprime category, where the frequency of credit defaults is much greater than in other types of securitized credit transaction (i.e. averaging 2–5 %, or five times more than agency experience). A typical HEL pool features loans with relatively small balances (i.e. under \$100 000, or half the size of a typical conforming agency mortgage loan) and high rates (i.e. 300+ bps above conforming agency loans); maturities span 15 to 30 years, LTVs average 70–80 %, and cash flows are either level-pay or balloon. Some HELs feature prepayment penalties for the first 3–5 years of their lives; it is typical, for instance, for borrowers to be assessed a penalty if principal repayments exceed 20 % of the original balance. That said, most HEL pools exhibit higher baseline prepayment speeds when compared to agency pools. Pools are more sensitive to refinancing due to credit quality improvements than interest rate declines, and tend to season faster than agency pools. Since subprime borrowers must pay up to 300 bps over conforming loans, they have considerable incentive to refinance as soon as their credit quality begins to improve; it generally takes at least one year of payments to determine whether the credit quality of a borrower merits readjustment. Accordingly, a HEL prepayment model must be calibrated specifically to accommodate credit-driven refinancings. Interest rate prepayment sensitivity is, as noted, far lower, primarily because subprime borrowers have limited opportunities to refinance through other sources, and cannot therefore take advantage of a declining-rate environment to the same extent.

HELs can be structured as open- or closed-end bonds, though the latter tend to be more popular. Closed-end bonds can be viewed as an amortizing mortgage (e.g. fixed maturity with amortizing payments). Securities that are backed by variable rate loans are issued with floating-rate coupons; these are commonly known as HEL floaters. The more sophisticated transactions also make use of the CMO tranching technology described earlier, issuing PAC-style tranches. Secondary market trading in benchmark HELs can be quite active; esoteric structures, however, are quite illiquid.

Other asset-based securitizations

Though the overall ABS market is based primarily on credit card and loan receivables and HELs, certain other structures also exist. The most important of these “second tier” ABS include securities backed by auto loans, student loans, manufactured housing loans, and aircraft leases. We describe each of the mechanisms in brief.

Auto loan ABS

Auto loan ABS, or securitized pools of auto loans, are issued by the finance subsidiaries of auto companies, general finance companies, banks, and auto rental companies operating large fleets. Within the US market, auto loan ABS account for approximately 30 % of the market (third in size after credit card and HEL securitizations); outside the US, the auto loan ABS is less common. Auto ABS are issued via grantor trusts (which function primarily as pass-through conduits) or owner trusts (which allow for greater restructuring/tranching of cash flows). Cash flow structuring tends to operate in a similar manner across deals, with interest paid to all tranches in order of seniority, and then principal paid in order of maturity, until each class is fully retired. If defaults occur, cash flows revert to a pro rata allocation. Note that in some instances, pools are comprised of leases rather than loans (in the US, approximately one-third of all auto purchases are actually leases).⁶⁴

Auto loan ABS generally are considered to be among the safest instruments in the ABS sector, as each security is backed by a pool that is collateralized by automobiles. In addition to senior/subordinated tranching, delinquency triggers, and overcollateralization, deals generally are enhanced by a seller-funded cash reserve account and an excess spread comprised of cash flows remaining after notes have been serviced and fees/credit losses have been paid. External support via bond insurance may also be used, primarily on subprime deals. The auto loan pool itself can be extremely large and diverse, helping mitigate any concentrated risks. The credit scores of borrowers are used to help define pool parameters. Auto ABS are often divided into separate pools that reflect the credit quality of the underlying borrowers; these may be classified as prime, near prime, or subprime, as measured by delinquency rates. Credit problems tend to occur early in the life of an auto loan, so seasoned packages of loans tend to have better credit performance.

Unlike other securitized assets, auto loan ABS begin paying investors principal and interest immediately, and carry much shorter terms than housing-related MBS and HEL structures (e.g. average terms of 3–5 years vs. 10–30 years). Investor cash flows are comprised of scheduled monthly payments of principal and interest, plus any prepayments. Prepayments, which are measured on a monthly basis as a percentage of the original collateral amount, are a function of sales, trade-ins, refinancing, repossession, and destruction. Unlike MBS and HEL ABS, refinancing in the auto sector is not a major driver of prepayments; given the relatively small purchase price of an auto (versus a home) and the short tenor of the loan, the interest savings generated by refinancing are negligible. Accordingly, prepayment behavior tends to be relatively stable.

In addition to standard fixed-rate amortizing tranches, the auto ABS market has also begun to use the revolving structure with soft bullet that is employed commonly in the credit card ABS sector. During the revolving period, new auto loans (or leases) may be added to the pool as existing ones are retired. The trust may thus issue CP and term funding: when the term funding comes due, the trust issues new CP to repay investors in bullet form; the process is then repeated. If the auto loan collateral pays faster than anticipated, the trust can repay CP before the term bullet note is due; if the collateral pays more slowly, the CP will still be outstanding when the bullet notes are due, so new CP can be issued and the two CP issues can be retired

⁶⁴ That said, credit experience for lease-backed deals can be stronger, as the prequalifying credit standards for leases tend to be more stringent than they are for loans. Some distinction must also be made between deals that expose investors to residual value risk, and those that do not. Auto lease bonds use a process where the SPE (often a titling trust) purchases the vehicles and retains a beneficial interest in the leases and vehicles; a special unit of beneficial interest is then securitized and conveyed to investors. Investors in this instance have a right to the cash flows from the lease, but are also exposed to residual value risk.

sequentially. Bullet notes can also be extended if the trust cannot sell CP. There is now also greater issuance of floaters, and more use of lockout periods without principal amortization, in order to extend average deal life (i.e. by up to two years).

Auto ABS launched in Europe are done primarily via SPEs using true sale techniques rather than master trusts. Deals, which feature average lives of three to five years, can include loans or leases from various nations.

Student loan ABS

Student loan securitizations appeared in the US market in 1993, and became popular during the mid-1990s, when the Student Loan Marketing Association (SLMA, or Sallie Mae), a quasi-government agency, started to float a steady supply of issues. In fact, SLMA has been instrumental in promoting activity in student loans by purchasing loans originated by various financial institutions, and selling them into securitization conduits. Various other private sector institutions entered the market with securitizations of their own following SLMA's early successes, and further growth occurred through government-guaranteed loan programs extended by private institutions under the Federal Family Education Loan Program (FFELP).

Student loan ABS are structured with an owner trust that issues individual tranches of securities backed by amortizing pools of student loans. A loan servicer is appointed to collect monthly principal and interest payments from borrowers (commencing six months after a student completes the educational course/degree) and forward the cash flows to the trust. Loans in a pool, which are all of unique size and maturity, may or may not be guaranteed by the federal government (through FFELP/Department of Education). The resulting securities typically are structured as sequential-pay tranches rated from AAA to A, with floating-rate coupons and ten year expected average lives; interest is paid every month or quarter in sequence. The trust relies on excess spread, subordination, and a reserve fund to create the AAA tranches. Student loan ABS tend to feature a higher degree of liquidity risk than other ABS, due to borrower delinquencies and/or delays in receiving payments from the Department of Education or other third party guarantors. Such liquidity risks can be mitigated (though not eliminated entirely) by draw down of a prefunded reserve account within the trust; this helps ensure that monthly principal and interest cash flows to investors remain uninterrupted.

SLMA securitizations, which are based solely on government guaranteed loans, feature relatively modest liquidity, credit and basis risk; accordingly, only a 3.5 % subordination level is required in order to create AAA tranches. Nongovernment securitizations, lacking the same low-risk characteristics, may require up to 15 % subordination in order to support a AAA tranche. Default risk can affect investor payment streams when performance is not government guaranteed. Historical experience suggests that default risk peaks during the first three years of scheduled repayments, and declines steadily from that point on. Unlike other MBS and ABS structures, student loan securitizations are not sensitive to refinancing opportunities, as alternative financing sources are not readily available to most borrowers.

Aircraft lease ABS

Aircraft leases, which are a popular form of financing in the commercial aviation market, are repackaged periodically by intermediaries through aircraft lease ABS or equipment trust certificates. Aircraft lease ABS are backed by a pool of leases from various airlines, while equipment trust certificates are backed by leases from a single airline. Accordingly, the rating

of the aircraft lease ABS is based primarily on the cash flows from the pool, while the rating of the equipment trust certificates reflects primarily the credit quality of the single airline obligor (and/or any third party guarantor that may be involved). Investors appear to favor aircraft lease ABS as a result of their diversification.

Though aircraft feature useful lives of 20+ years, most aircraft leases are operating leases covering 4–5 years. This means that the specialized aircraft leasing companies that are responsible for arranging and servicing the financing transactions must actively re-lease the aircraft. The related securities, however, make no assumptions about the re-leasing process, and are based solely on the terminal maturities of the operating lease contracts. Though the earliest deals in the market relied on aircraft sales to cover principal repayments, most are now based on lease revenue streams, meaning enhancement requirements have declined.

Manufactured housing ABS

The US housing market features strong demand for manufactured houses, which are single- or multiple-section units that are manufactured in a factory and delivered to, and assembled on, the housing site. The key attraction of such manufactured housing relates to cost: units are up to 75 % cheaper than a comparable site-built home. General and specialty lenders extend loans for the purchase of manufactured housing, with financing set to cover only the house (as personal property), or the house and land (as real estate). Personal property financing is arranged through an installment sales contract, while the real estate financing is structured as a standard note and mortgage. Pools of manufactured housing loans are combined using many of the techniques described above, and result in the creation of manufactured housing ABS.

Various other forms of asset-backed securitization also exist, e.g. pooling and securitizing of small business loans, sporting event revenues, music royalties, financing leases, trade receivables, whole business enterprises (e.g. pubs), public sector assets, and so forth. Most of these follow the general principles noted in this chapter, and are driven by many of the same motivations – we shall not, therefore, go into any further detail.

Structured Notes and Loans

5.1 INTRODUCTION

Structured notes and loans represent a vital, and highly creative, portion of the structured asset market. The sector, which combines derivatives and fixed-income instruments, has existed for nearly three decades, and has proven to be an important element of investment, funding, and risk management strategies. The sector, in its broadest terms, encompasses a range of asset classes, including interest rates, currencies, commodities, equities, and credits.¹ Though the volume and growth rates of different segments of the market vary, all play a role in helping intermediaries and end-users achieve specific end goals.

In this chapter we review the development and market drivers and product mechanics/applications of structured notes and loans. Since the market is dominated by fixed- and floating-rate notes, rather than loans, we shall focus our comments accordingly; where relevant, however, we shall introduce analysis of loan products. The market features a variety of basic structures created from combinations of securities and options, swaps, or forwards, that can be applied to create principal-protected, coupon-increased, coupon-decreased, or step-up structures. These can be used across all major asset classes, as noted in Figure 5.1.

5.2 DEVELOPMENT AND MARKET DRIVERS

Structured notes and loans have been issued regularly since the mid-1980s, just as activity and interest in rate, currency, and commodity derivatives began accelerating. Intermediaries discovered that pairing derivatives with funding instruments could provide investors with new investment opportunities, and issuers with better financing levels. Investors with aggressive return targets and significant risk-tolerance thresholds soon incorporated varying amounts of leverage in their structured notes.

Interest-rate-linked bonds appeared originally in the mid-1980s, several years after the introduction of basic interest rate swaps and options. Though the earliest embedded interest rate notes were introduced in the US, they were replicated immediately in the European and Japanese markets. The relative inefficiencies of the swap and option markets in the earliest days of product development helped spur growth, as arbitrage opportunities for issuers were significantly larger and more obvious than in the new millennium; even the most elemental structures allowed issuers to monetize value and lower funding costs, while simultaneously giving investors new opportunities to express views on rates and rate volatility.

Notes and bonds linked to individual currencies also appeared in the mid-1980s, as issuers began embedding derivatives on major currencies in their securities. The most common references during this early period centered on the “majors,” including Japanese yen, British pounds sterling, Deutschmarks, Swiss francs, Italian lira, Australian dollars, and Canadian

¹ Certain insurance risks, including those related to catastrophe and weather, are also embedded in notes/bonds; we consider these unique instruments in Chapter 7.

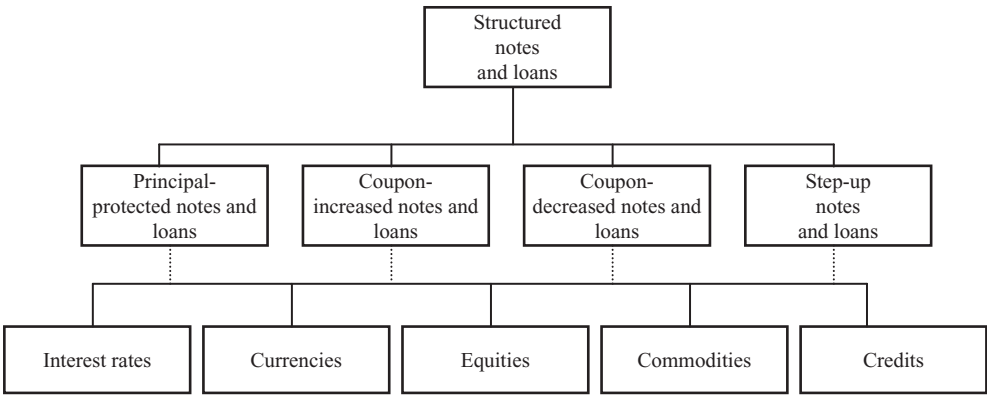


Figure 5.1 Structured notes and loans

dollars; deals related to other reference currencies appeared opportunistically. In more recent years, the market has expanded to include bonds linked to emerging market currencies and, with the advent of the European Union, the euro.

Commodity-linked notes entered the market shortly after currency-linked bonds. Early structures were based on the world’s main commodity references, including oil, precious metals (gold, silver), and certain industrial metals (copper, aluminum). Energy-based issues have proven particularly popular over the years with energy producers (e.g. oil companies) as well as energy consumers (e.g. transportation companies, airlines); producers have been able to use the notes to lock in revenues/receipts, while consumers/end-users have employed them to lock in the purchase prices on the inputs. Common energy references include various forms of crude oil (e.g. Brent, West Texas Intermediate), crude products (e.g. jet fuel/kerosene, heating oil, naphtha), and natural gas. After several years of successful issuance on individual commodities, the market expanded in the mid-1990s to include notes on broader commodity index references.

Equity-linked notes² have proven popular in both bullish and bearish markets. As OTC equity derivatives emerged as an important asset class in the late 1980s, intermediaries began issuing notes referencing the appreciation of major market indexes, such as the S&P 500, Nikkei 225, and FTSE 100; in some instances, these notes allowed investors to participate in market sectors that were otherwise restricted. Strong bull markets in the US, Europe, and Japan throughout the latter part of the 1980s/early 1990s prompted significant issuance. When market conditions became bearish in the early 1990s (particularly in Japan), issuance remained relatively robust as issuers and investors turned to market depreciation structures.³ Banks also began issuing equity-linked deposits, offering investors a choice between instruments with a

² Note that convertibles and other equity-based structures that result in new equity issuance on conversion or exercise are considered separately in Chapter 8; the discussion in this chapter is based on notes and bonds that simply reference equity indicators.

³ For instance, prior to the descent of the Japanese markets, which started in 1990, Japanese investors believed the market would continue to rise, and were thus active investors in high-coupon “bull notes,” which were note securities with embedded short puts, often issued via high-grade European corporate or supranational issuers. Many of these notes incorporated a considerable amount of leverage (e.g. 2 % of lost principal for each 1 % decline in the Nikkei 225 below the participation rate). Such notes were of considerable appeal to zaitech (financial engineering) speculators, who were convinced that the market would remain bullish. Issuers were also eager participants, as monetization of the embedded option led to very low funding levels. US banks arranging the notes purchased the options from the issuers and resold them at higher levels to domestic US retail investors (generally via separately structured notes or warrants listed on one of the stock exchanges). As the Japanese market declined steadily between 1990 and 1993, Japanese investors that originally sold the puts suffered losses, while US investors posted gains.

guaranteed return and a lower equity participation rate, or a higher equity participation rate in lieu of a guaranteed return. Thereafter, variations on the theme began appearing regularly, including sector notes, giving investors an opportunity to participate in the equity movement of individual industries (e.g. pharmaceuticals, financials, cyclicals), and country-based notes, referencing the returns of broad baskets of stocks in a specific country. While much of the equity-linked activity was created for the OTC market (this remains the largest component to the present time), some was directed toward the exchange/listed markets in order to promote retail activity.⁴ Perhaps more than any of the other notes considered in this chapter, equity linked notes have been designed primarily to meet investor needs/requirements regarding equity market exposures.

Credit-linked notes (CLNs), which are structured debt securities that reference the financial performance of credit-risky issuers/borrowers (i.e. a form of funded credit derivatives), have emerged as one of the most prominent parts of the structured note marketplace. Though CLNs originally were developed as investment vehicles for institutions seeking specific credit exposures, they soon came to serve as an important tool for those seeking hedging, risk control, diversification, and liquidity access. The first CLNs appeared in the early 1990s, but true growth came in subsequent phases, including the structured note cycle of the mid-1990s, and the credit derivative expansion cycle of the late 1990s. In fact, the credit markets of the late 1990s and early millennium featured forces that promoted development and activity: deteriorating credit quality due to regional and sectoral slowdowns,⁵ excessive concentrations of credit exposures within the community of intermediaries and originators (and a resulting need to find alternative mechanisms to transfer credit risks and optimize risk capital), and a desire from within the investor base to absorb, directly or synthetically, attractively priced credit risk structures.⁶ The market has expanded rapidly in the new millennium.

Structured notes, regardless of asset class, feature certain common drivers that have been instrumental in fuelling market activity over the past few decades – either cyclically or continuously. The generic form of the structured note can provide investors, issuers, and intermediaries with various benefits; in particular, a structured note:

- gives investors access to asset exposure in unique forms, with asset references, maturities, currencies, leverage, and coupon/principal payments customized to meet individual requirements – this includes providing access to markets that might otherwise be restricted;
- allows investors that are not permitted to deal in derivatives directly to replicate the risk/return profile via a securitized structure;
- creates an efficient mechanism by which to purchase a desired reference asset – this generates administrative savings when purchasing an entire portfolio of assets through a single transaction;
- allows investors to avoid the administrative burden of having to arrange separate derivative transactions to achieve the same results, thus reducing documentation expenses and avoiding use of counterparty risk limits;

⁴ For example, BSC S&P Notes, Market Index Target Term Securities (MITTS), and the Stock Market Annual Reset Term (SMART) notes are listed on various exchanges. MITTS, a popular retail investment vehicle, are structured as five-year notes with principal protection paying 100 % of the appreciation in the target market. Similar index notes have proven popular in the European markets, particularly with the advent of the euro and the introduction of pan-European indexes (e.g. Eurotop, DJ Euro STOXX, and so forth).

⁵ Examples of troubled sectors include emerging market credits in the mid-late 1980s, and again in the mid-late 1990s, US commercial real estate in the early 1990s, Japanese real estate in the mid-1990s, technology/media/telecom in the early millennium, and so forth.

⁶ Interestingly, the CLN market has also been instrumental in “restructuring” other structured credit assets created in the mid- and late 1990s that encountered difficulties as a result of excessive leverage.

- allows issuers to lower funding costs to levels that might not be achievable through normal mechanisms (e.g. a funding arbitrage opportunity) – this is done typically by taking a view on a forward curve;
- permits intermediaries to avoid counterparty risk exposure by transferring derivative-based exposures in funded form;
- makes it possible to create synthetic, tradable references where none exist;⁷
- allows participation in certain markets by synthetically replicating existing assets that are in short supply and high demand;
- transforms otherwise illiquid assets into a more marketable form, thus adding to market liquidity – this ultimately helps lower issuer funding costs, and creates more attractive and secure opportunities for investors;
- offers tax and regulatory advantages that can benefit the investor and/or issuer;
- permits intermediaries or other users to transfer risk exposures that they wish to hedge, diversify, or reduce.

Investors in structured notes/loans include financial institutions, insurers, pension, mutual, and investment funds, and corporates. Though all are active, they often select specific segments of the market in order to express their investment views. For instance, pension funds often participate in credit- and equity-linked structures, but may prefer those with a high level of principal protection; their chief aim is to enhance returns while preserving capital. Corporate treasury operations, investing excess cash on behalf of parent corporations, may favor similar “conservative” risk/return profiles. Hedge funds, in contrast, often assume a more aggressive stance, preferring to establish speculative/high-risk positions; this may involve greater use of leverage and/or limited principal protection (e.g. full capital at risk).

5.3 PRODUCT MECHANICS AND APPLICATIONS

Before examining the specific characteristics of interest rate, currency, equity, commodity, and credit-linked structures, we consider a number of general structural issues that are applicable across all types of linked notes and bonds.

5.3.1 General structural issues

The structured note and loan market involves different types of market reference, but basic structures/product designs are applicable across most references. Indeed, the common initial driver in many structured notes – whether commodity-, currency-, credit-, or equity-linked – is a lowering of an issuer’s funding costs by combining market structure and a view on implied future rates/prices. As noted earlier, structured notes are available in various generic forms, including:

- **Principal-protected.** The investor is only at risk for loss of coupons in the event a defined risk event occurs.
- **Coupon-increased.** The investor’s principal is at risk in unleveraged or leveraged form – generally up to the entire amount of principal.⁸

⁷ It is important to stress that an excessive amount of customization, e.g. creating an asset with an odd coupon payout profile, can drain liquidity from an issue. In addition, securities that are reissued strictly as private placements have a limited investor base (e.g. other qualified institutional buyers).

⁸ In certain instances, principal redemption may not be floored, allowing for a negative redemption scenario, where the investor in the note is at risk of having to repay an additional amount of principal if the market reference moves against it; such a structure is quite unusual, but can be arranged for investors that are interested in highly leveraged and speculative opportunities.

- Coupon-reduced. The investor receives full repayment of principal, and an extra payment in the event a defined risk event occurs; if no event occurs, the coupon is reduced.
- Step-up. The investor receives a larger coupon for each occurrence of a defined risk event.

The payoff profile of all notes can be modified as to degree of leverage, principal protection, and trigger event; principal-protected structures generally only work from an economic perspective on long-dated deals, since the actual risk/return parameters must be embedded in the coupon stream.

Structured notes are, in general, comprised of a host bond (3–7 years in maturity) and one or more derivative contracts.⁹ For instance, a combination of a bond and a forward can be used to create a dual currency bond (FX forward), or an oil-based bond (oil forward). An FRN and a swap can be used to create an inverse floater, or a dual index note (CMT and LIBOR or Cost of Funds Index and LIBOR). A bond or FRN and an option can be used to create a capped or floored note, or a range floater (interest rate option) or an equity-linked note (equity option). In cases where international flows are involved, such as a dollar note based on the movement of a foreign index, the payouts may be protected through embedded currency options (quanto option) in order to neutralize any currency risks.

As indicated, the design of a structured note must take account of various factors, including desired investor participation levels in future market movements, degree of principal protection, payment of interest or dividends, maturity, and optimal derivative use. For instance, if the investor is amenable to the use of an average price, rather than a terminal price, in the determination of principal redemption, then an Asian option can be used. Since the Asian option typically is cheaper than the European or American option, the investor can benefit through an increased participation rate. Alternatively, if the investor is willing to assume a greater amount of risk in exchange for a higher participation level, he can sell the note issuer a higher strike option (e.g. cap or call).

Actual host bonds can be structured in various forms, including coupon-bearing, zero coupon, amortizing, bullet, or balloon. Host issuers, generally AAA or AA corporate or supra-national organizations, allow intermediaries to issue structured liabilities on their behalf, in exchange for a fee or reduced funding cost. Investors purchasing such securities thus face little default risk associated with the host issuer – coupon and principal-based returns are based solely on external references that are incorporated into the relevant payoff formula. While this type of structuring and issuance was popular during the 1990s (and still exists to some degree), activity has shifted in favor of the repackaging structure. Repackaging vehicles, organized as bankruptcy-remote SPEs or trusts, were introduced originally in the late 1980s, but gained true popularity in the mid-1990s during the first cycle of structured note business. Most major financial institutions operate their own repackaging vehicles,¹⁰ which they use to acquire assets, repackage cash flows via swaps, and issue notes to investors. Most vehicles are created for multiple issuance in order to improve cost efficiencies; indeed, once a vehicle is established, the costs of new note issuance are generally very modest, making the process economically viable for both intermediaries and investors. The securities that are ultimately placed by the repackager with investors take the form of notes (SPE) or trust receipts (trust); each series features unique cash flow characteristics that correspond with specific demands of investors. Notes may be arranged on a publicly registered or private placement basis. Those

⁹ Notes with longer maturities, e.g. 7–10 years, tend to carry higher coupons and greater investor participation in appreciation/depreciation of the reference asset or index.

¹⁰ For instance, Merrill Lynch uses the SIREs and STEERS vehicles, Citibank the TIERS vehicle, JP Morgan Chase the CRAVE vehicle, Barclays the ALTS vehicle, and so forth.

that are publicly registered are rated in order to increase marketability and liquidity; ratings are a function of both the underlying bond and the derivatives counterparty (which is often a AAA/AA-rated financial institution). That said, liquidity for structured notes must be considered limited; the very fact that transactions are highly bespoke suggests a relatively small audience for resale and repurchase. The exception occurs with certain standardized notes that are listed on exchanges (and which are often intended for retail investors); liquidity in such cases generally is somewhat greater.

Though construction and valuation of structured notes can be an involved task, the individual components of any given security can be examined separately to determine proper design and pricing. For instance, a bond with an embedded option can be decomposed, and each leg can be priced separately. In an arbitrage-free market environment, the decomposition process helps ensure some degree of price equilibrium. To demonstrate general product mechanics, we present in the sections that follow a sampling of common structures from each of the major asset classes. The products we consider are not exhaustive – many other variations have been, and are, created.

5.3.2 Interest-rate-linked notes

Structured notes with embedded interest rate options are common, appearing frequently when interest rates are low, the yield curve is steep, and interest rate volatility is on the rise; each one of these factors allows investors and issuers to benefit. Key structures in this category include inverse floaters, leveraged floaters, capped and collared floaters, range floaters, step-up bonds, and bonds with debt warrants.

Inverse floaters are structured to give investors a leveraged return to market interest rates, and tend to appear when investors have a view of future rates that diverges from the market implied view. As the name suggests, inverse floaters feature coupons that move opposite the market, i.e. as LIBOR rises, the coupon on a floater declines, and vice versa. For instance, if LIBOR is currently set at 4 %, the structured coupon on the FRN might pay a coupon of 8 % – LIBOR. As LIBOR rises (e.g. to 5 %), a standard FRN investor receives a greater return (5 %), but the inverse floating investor receives a lower coupon (3 %). Riskier versions of the note can also be created; these instruments, known as leveraged inverse floaters,¹¹ can feature risk/return payoffs with many multiples of leverage.¹²

The decomposition of the inverse or leveraged inverse structure suggests the investor holds a package comprised of a standard FRN and a swap, where it receives fixed rates and pays LIBOR on n times notional, where n is the leverage multiple (and $n = 1$ for an unleveraged structure). Thus, if the fixed rate is set at 5 % and the leverage factor is 2, the investor receives ($2 * 5$ %, or 10 %) and pays ($2 * \text{LIBOR}$); it also receives LIBOR on the standard FRN, suggesting a combined total coupon of 10 % – LIBOR. The combined structure, floated via an SPE, is depicted in Figure 5.2.

Financial institutions and corporations seeking to lock in funding levels routinely issue capped and collared FRNs. The capped FRN is a package comprised of an FRN and a series of long issuer caps that establish a maximum borrowing cost. If forward rates are high (e.g. the curve is steep) or short-term rates are especially volatile, the caps are expensive, meaning a higher return for investors; high forward rates increase the “lock-in” value for each reset date,

¹¹ Note that a variation on the leveraged floater theme appeared in the late 1990s – the deleveraged floater, with coupon payments set as a fraction of the actual floating-rate index.

¹² Other variations are also available. For instance, a fixed inverse floater pays a fixed rate for the first one to two years, and then converts into a standard inverse floater; this allows an investor to defer its view on rate movements until a future time period.

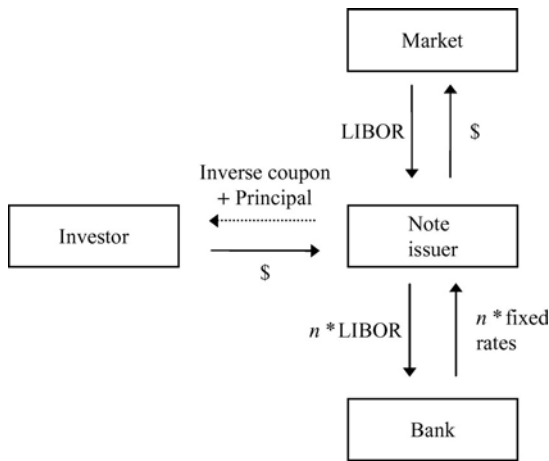


Figure 5.2 Inverse floater with n times leverage

and the high volatility increases the time value for each date. In some instances, the strip of caps sold by investors via the note may be misvalued; this is particularly true when the curve is very steep. In such cases, the issuer can sell an equivalent strip of caps at higher levels in the market, using the crystallized differential to lower its funding costs. Another version of this structure, the collared FRN, incorporates a strip of short issuer floors to give investors a minimum market return while defraying some (or all) of the issuer's cap premium cost. When the curve is steep, forward rates are significantly higher than current rates; this means that the floors the issuer has sold investors may be in-the-money in a current period, but may move out-of-the-money in a future period; the reverse scenario occurs for the caps. If the issuer can sell the floors to investors at a relatively high level, and repurchase an identical strip at a lower cost in the market, it again lowers its funding costs. Cost savings can run to more than 50 bps (though this depends on market circumstances). In certain other situations, the reverse transaction can also be arranged (e.g. the issuer sells caps and buys floors).¹³ Figure 5.3 illustrates the collared FRN.

Note that several other variations seek to capitalize on the same arbitrage opportunities. Collared FRNs with multiple cap/floor strikes can be arranged, with each caplet/floorlet featuring a unique strike. Ratchet-collared FRNs are similar in concept, but contain caplet/floorlet strikes that increase at particular intervals; such structures are most valuable to issuers when the implied forward rates suggest greater moneyness will accrue to the caplets. Capped and collared FRNs with participating caps/floors are also available. These, for example, may allow for increased participation levels as the cap strikes rise for each sequential coupon period, with a payoff equal to $(\text{cap strike} - \text{participation rate} * (\text{strike rate} - \text{LIBOR}))$. Obviously, not all collared or capped structures reference LIBOR; other rates, including constant maturity Treasury (CMT), EURIBOR, bank bills, and so forth, can also be used.

¹³ A breakeven analysis is required to determine the economic viability of a capped or collared FRN, and to examine where the value lies. Let us consider the following example: a benchmark two-year FRN is quoted at $L + 50$ bps, while a two-year capped FRN (6% cap strike) is quoted at $L + 80$ bps. If LIBOR is trading at 4.5%, then we can define the maximum the rate can increase in one year before the capped FRN and the uncapped FRN break even; this analysis involves reinvestment of the first 12-month coupon at the future one-year uncapped rate. Solving the inequalities yields a LIBOR breakeven level of 6.64%. Thus, if LIBOR exceeds 6.64% in one year, the cap feature reduces the FRN's yield by more than enough to cover the 30 bps cost of the cap. A similar analysis can be performed for collared FRNs.

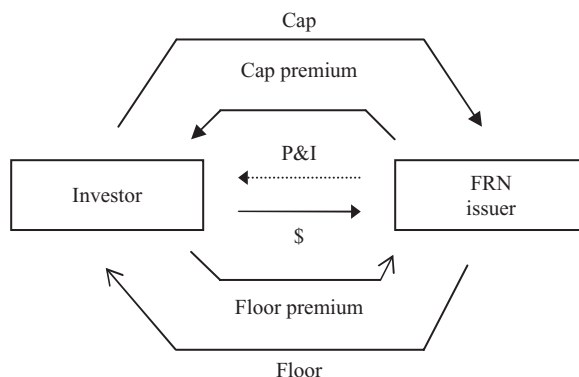


Figure 5.3 The collared FRN

The range floater structure (also known as an accrual note) is similar in structure to the collared FRN, except that the caps and floors are replaced by digital options that provide a payout/coverage of either a fixed amount or zero. The range floater pays interest to note investors only on days when LIBOR (or some other floating-rate index) falls within the boundaries established by the upper and lower strikes of the digital options; for every day within a quarterly or semi-annual period that LIBOR falls outside the range, interest ceases to accrue. If LIBOR falls outside the bands for an entire quarter, investors lose interest for the full quarter. In exchange for potential risk of loss on the coupon, investors receive an enhanced coupon while the benchmark rate remains within the band. This suggests, of course, that investors are selling the issuer a strip of digital options, with the enhanced coupon representing premium. Use of the digital options means that the coupon/protection profiles are discrete rather than continuous, and the value of the note can change rapidly as the underlying rate/price approaches the upper or lower strikes. Ranges can be held constant over the life of the note, or they may be increased at specific intervals (coinciding with a positive yield curve and higher implied forward rates).

Let us examine an example of this structure. Assume that an investor buys a three-year range floater paying an enhanced coupon of three-month LIBOR + 40 bps in a market when LIBOR is equal to 5 %. The note features a step-up range, where year 1 has a range of 5–6 %, and years 2 and 3 have a range of 6–7 %. We can determine the number of days the range floater can trade outside the range and still be equal to a standard FRN (i.e. the breakeven LIBOR level at which the range floater is equivalent to a standard FRN trading at LIBOR flat). In this example, the extra revenue earned over a full year amounts to 40.5 bps; the cost of trading outside of the range for one day is equal to approximately 1.4 bps, meaning that the breakeven position is 29 days. This analysis suggests that if LIBOR trades inside the range for more than 29 days, the investor's position is profitable. Note that a more extreme version of the range floater, the knock-out floater, is comprised of a series of knock-out (barrier) options, rather than digital options, suggesting that if the reference rate trades outside of the range for even a single day, the coupon for the entire period is sacrificed (e.g. it knocks-out as the barrier level is reached); these are particularly risky versions of the accrual structure.

It is important to note that investors in capped/collared FRNs and range floaters face different risks. Investors in capped/collared notes are speculating that the rates implied by the yield curve are too high; if correct, they preserve the enhanced coupon representing the premium. Investors in range floaters, in contrast, are taking a view primarily on interest rate volatility (yield curve

moves are only secondary); specifically, they must believe that the premium received via the enhanced coupon will be sufficient to make up for a potential loss in daily yield driven by volatility.

The step-up note is also quite common. Under this structure, which first appeared in the market in the late 1980s, the investor receives an above-market rate during the note's noncall period (generally one or two quarterly periods), and then receives increasingly higher rates (step-ups) for each period in which the note is not called by the issuer; the step-up rates are likely to be below-market rates for the first few periods. The investor in this instance has sold the issuer a strip of call options, with the initial above-market rate constituting the premium.

A further interest-rate-linked bond meriting mention is the bond with attached debt warrant, giving the investors the option to purchase an incremental amount of new fixed-rate bonds (with terms specified at original issuance). In exchange for the option, the issuer receives a premium payment (generally in the form of a below-market coupon) that lowers funding costs. Since the warrants are attached to the host bond, the issuer also benefits from lower fees for any subsequent issuance. In addition, market convention calls for setting the new issue coupon at 25–50 bps lower than the host bond, creating a built in “cost saving.” Though the issuer can clearly benefit via lower costs, it faces a complication in being unable to estimate, in advance, the precise amount of bond warrants that will be exercised into new debt; much depends ultimately on the level of interest rates and credit spreads. In the extreme, the warrant may expire worthless, requiring the issuer to finance or refinance through alternate sources.

In addition to securities linked to interest rate references, government and private issuers launch notes and bonds that reference inflation levels. Such bonds allow investors to earn real rates of return based on inflation recorded over a particular period of time. In the US, such inflation-linked bonds are tied to the consumer price index (CPI), in the UK, to the retail price index (RPI), and in the Euro-zone, to the harmonized index of consumer prices (HICP).

5.3.3 Currency-linked notes

The currency sector features a number of important structures that allow investors to capitalize on movement in currency rates and/or volatility. Two of the most popular include bonds with embedded currency options (or warrants), and dual-currency bonds.

Bonds with embedded options can be structured to give an issuer exposure to a particular currency. For instance, a company may issue a fixed-rate bond with a long embedded currency option, paying the investor premium via an enhanced coupon. The value of the option impacts redemption: if the currency option remains out-of-the-money, the investor receives par at maturity; if, however, the option moves in-the-money, the principal redemption payable to investors declines on a sliding scale related to the level of moneyness (the further in-the-money, the lower the redemption to a minimum floored level). Principal reduction may be defined by a formula such as:

$$P^* \left(1 - \frac{\text{strike} - \text{spot}}{\text{spot}} \right)$$

where

P is principal

strike is the strike price of the currency option

spot is the spot level at maturity.

Consider, for instance, a structured currency note of \$100 m that pays investors an enhanced coupon and par redemption while the \$ trades weaker against the yen (versus a current spot and strike rate of Y105/\$). As the \$ begins to strengthen, the principal redemption begins to decline, suggesting that the issuer of the note has purchased an embedded \$ call/yen put. Assume that at note maturity, the dollar strengthens to Y110/\$. Under this scenario, investors will receive principal redemption of \$95.5 m (i.e. $\$100 \text{ m} * (1 - (105 - 110)/110)$).

The issuer may sell an identical option in order to monetize value and lock in a reduced funding cost. Increasingly sophisticated variations of this basic structure have appeared in recent years, including bonds that create variable coupons based on currency movements. For instance, a note may feature coupons that increase as a reference currency appreciates (i.e. the base currency depreciates), with additional caps and floors. This asset is equivalent to a note with a put spread and a call spread (e.g. the issuer buys an at-the-money currency put and sells a lower strike put, and vice versa with the currency calls). Ultimately, the issuer obtains a hedge on foreign currency revenues; if the reference currency strengthens, then its revenues are higher in base terms, which offset the increased interest cost payable on the note. The reverse scenario occurs if the reference currency weakens. Multivariate notes are a further extension of the basic structure, linking coupons and/or principal to the level of a reference currency and interest rates. Notes can also be constructed with embedded currency baskets (e.g. emerging market currencies), or exotic currency options, including barriers, digitals, lookbacks, and ratchets; these, not surprisingly, are somewhat less common.

In a similar light, bonds with embedded currency warrants have featured since the late 1980s, and are still issued in the new millennium – though in cycles corresponding with significant currency events or trends. Funding reduction is, again, the primary driver in this structure, which involves the sale of an embedded call or put warrant on a target currency via the note. In order for this to occur, the currency warrant sold through the note must be more expensive than the mirror option the issuer purchases from an intermediary. If this can be achieved, the issuer crystallizes a gain that is used to lower funding costs.¹⁴ Investors receive enhanced principal redemption if the reference currency warrant moves in-the-money prior to maturity; in most instances, principal is protected, with minimum redemption value set at par. Such currency-linked warrant bonds can be issued on emerging market currencies, giving institutional (and retail) investors access to risky emerging market currency movements in a convenient form.

Dual-currency bonds were among the earliest, and most successful, of the currency-linked fixed-income securities, and remain prevalent in the market to the present time. The securities were introduced originally on a private placement basis in the late 1980s, with Japanese insurers acting as capital providers to US and European issuers. The structure gave insurers protection on their yen principal, but provided the potential for gains via foreign currency coupons, such as A\$. The “opposite” form of the structure, the reverse dual-currency bond, was structured with yen coupon and A\$ principal flows, for those seeking to express a different view on the currency.

The dual-currency bond consists of a relatively complex series of spot currency transactions and swaps, but these are transparent to investors, which receive the benefit of the structured coupon/principal flows. In a typical dual-currency bond, such as the yen/A\$ structure, an investor might purchase a yen-denominated structured private placement that pays coupons in an alternate currency, such as A\$. The issuer, in turn, might seek to convert the yen proceeds

¹⁴ During the earliest phase of this period, many structured currency notes were sold into the retail sector, where it was possible to sell the options at a relatively high premium; the repurchase of back-to-back options from the institutional market at lower cost locked in a funding cost saving.

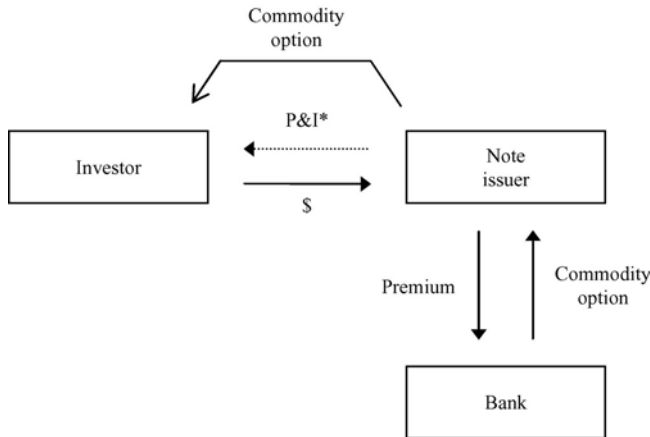


Figure 5.4 Commodity-linked note. * Principal redemption increased if option moves in-the-money; interest coupon decreased as a form of premium payment

directly into dollars via a swap. The package in this instance is comprised of yen notes, the first leg of a yen/\$ currency swap, spot currency trades converting yen into A\$ for the purchase of an A\$ bond investment to generate A\$ coupons, and a final leg of the yen currency swap (e.g. the investor receives yen principal at maturity). From the bank intermediary's perspective, enough yen must be invested at the start of the transaction via a zero coupon yen/\$ LIBOR swap to provide sufficient proceeds at maturity. The bank swaps a portion of the borrower's yen proceeds for dollars; the balance is used to buy a portfolio of A\$ bonds, which generates the periodic A\$ investor coupons.

5.3.4 Commodity-linked notes and loans

Commodity-based structures are created on common commodity references, including oil, gold, silver, and aluminum, as well as broader commodity indexes. In some instances, the notes/loans are issued by commodity producers, which can use the instruments to hedge core commodity risks. In other cases, they are issued by intermediaries or other institutions that seek to use the references to monetize a market view (e.g. the forward commodity curve); this monetization can help lower funding costs or increase investment returns.¹⁵

Common structures involve embedding of commodity options, forwards, or swaps in fixed- or floating-rate notes. As above, issuers selling embedded commodity options lower their funding costs by offering investors a lower yield (i.e. one that is equivalent to the theoretical premium on the options, amortized over the maturity of the note). Investors receive the upside or downside of the commodity if it moves in-the-money. An issuer selling a commodity option that it does not require for its own operations typically repurchases an identical option from an intermediary in order to neutralize the risk; even after paying the premium, the issuer's all-in funding costs can be lower than those of a straight issue, though this depends on the relative strikes of the two options. Figure 5.4 illustrates this general structure.

¹⁵ For instance, Swedish Export Credit was one of the pioneering institutions in the market. It has issued commodity-linked structures regularly in order to lower its funding costs, and has been active across a range of commodity references, including silver and oil. The firm has also issued dual-reference notes (oil and gold), with principal repayment dependent on the price of both commodities.

Let us consider another example to illustrate the economics of a typical transaction. In this instance, an issuer floats a three-year fixed-rate bond with a coupon of 9 % and embedded gold call options (or warrants); each option gives the investor the right to buy 1 oz of gold at \$400. If the bond alone is a par-valued security, and the present value of the entire structure is 109.45, then the value of the call embedded in each bond is \$9.45. In order for the funding arbitrage to work properly, the premium on the gold option the issuer repurchases must be less than \$9.45. For example, if the issuer purchases an offsetting call at \$8.75, the yield on the bond may decline to 8.45 %, a 55 bp saving. If the issuer opts not to hedge the short call exposure, it has a speculative position that negatively impacts its cost of funds as the price of gold rises.

We can also examine an alternative, but similar, structure, where the note contains an embedded prepaid forward. In this instance, the issuer, a gold producer, can sell gold forward as a hedge via the note; the note effectively is a security collateralized by future gold sales. Assume that spot gold is now \$405, the 12-month forward price is \$420, the risk-free rate is 5 %, and the annual gold lease rate is 1.3 %. Each bond raises \$405 today and pays off 1 oz gold in one year, plus interest based on the gold lease rate (e.g. \$5.25/bond, $\$405 * 1.3 \%$). Accordingly, the present value of the payoff is simply the coupon and the prepaid forward price of gold – which, in an arbitrage-free market, amounts to \$405. The producer is thus indifferent to what happens to the price of gold: if it rises, the payout on the note is greater (e.g. 1 oz gold at a higher price), but the value of its gold rises in tandem; if gold falls, the payout on the note is lower, which offsets the lower value of its gold inventory. Note that while an institution can use the forward alone to accomplish the same hedging task, the structured note provides it with a hedge plus financing. To value this type of structured note we must compute the coupon stream and the value of the prepaid forward. The prepaid forward is equal to paying for an asset today and receiving it on an agreed future date. This, of course, is different from an outright purchase, because the asset is not received until maturity, and it is different from a standard forward, because payment is made today, rather than at maturity.

Commodity-linked loans also appear periodically, primarily in support of a commodity producer's funding program. For instance, a copper producer may borrow funds via a loan, with a cost that declines through the sale of embedded copper calls; since the producer owns the underlying copper, it is effectively operating a covered call writing program, selling options on a portion of the production, and using the premium to lower its funding costs. If the price of copper rises, the calls move in-the-money and the lenders/loan investors holding the liabilities will receive an enhanced coupon or principal; the copper producer is indifferent to the increased payout, however, as it earns more on its core copper stock.

5.3.5 Equity-linked notes

Equity-linked securities are floated in a variety of forms, including FRNs and bonds (primarily via MTN and EMTN programs in order to reduce costs/time)¹⁶ with embedded derivatives referencing individual equities, baskets/sectors, or indexes. Though the embedded derivatives can take the form of equity forwards, swaps, and options, the latter tend to be most popular. This is because options are unilateral contracts, so cash flows can be controlled more precisely; in addition, the growing range of exotic equity options (e.g. barriers, digitals) means that highly bespoke risk profiles can be created for investors. In fact, notes can be created to give investors bullish or bearish views on market direction or market volatility.

¹⁶ Public deals issued via shelf registration need only adhere to the prospectus supplement requirement, rather than a full prospectus.

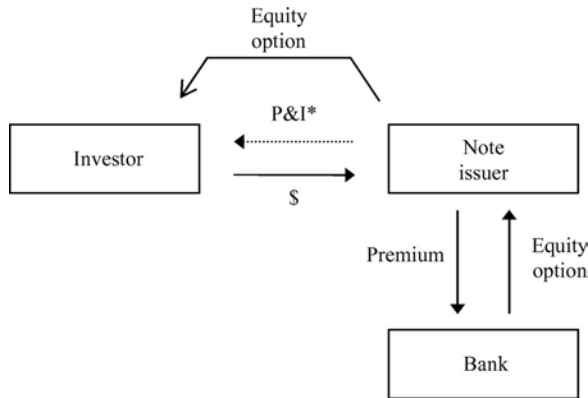


Figure 5.5 Equity-linked note. * Principal redemption increased if option moves in-the-money; interest coupon decreased as a form of premium payment

An investor in a basic equity note purchases a security that pays no coupon (or a below-market coupon), in exchange for appreciation (e.g. a call) or depreciation (e.g. a put) in the equity reference. The redemption of the note typically is set at par (i.e. a minimum guaranteed fixed payment) plus a percentage of the appreciation or depreciation in the reference equity index.¹⁷ The issuer of the note, such as a bank or company, benefits from a lower cost of funding through the receipt of the premium payment (e.g. the zero coupon or below-market coupon rate on the note). As with the other notes mentioned above, the issuer may purchase a similar/identical option in order to monetize value. Figure 5.5 summarizes this structure.

Creating this type of equity-linked security follows a multi-step process that begins with computation of the issuer's funding level. Once the target level is known, it is converted into a fixed rate covering the term of the note; this generates the value of the implied zero coupon bond. For example, if the funding level suggests a 75 % zero coupon threshold is required, \$75 of each \$100 of note proceeds is used to purchase a zero coupon, which grows to \$100 by maturity. The balance of the note proceeds is used to acquire an option (e.g. a call). With a market value on the option of \$20, the maximum amount of call coverage is \$25/\$20, or \$125 per \$100 of note issuance. In practice, some amount of the remaining \$25 balance must be used to cover issuance fees and/or to pay an interim coupon (on nonzero coupon deals). It is easy to see that if principal protection is reduced (e.g. the \$100 is not guaranteed at maturity), larger coupons and/or larger participation can be introduced into the process. In fact, the coupon/principal value of the note is a function of interest rates, maturity, option moneyness, option volatility, and desired level of equity participation. A note with an out-of-the-money equity option commands less premium (and thus more coupon), but also provides less upside potential; one that is at- or in-the-money commands a greater premium (and thus a lower coupon), but supplies a larger upside potential. In general, as moneyness increases, the participation rate increases; the same is true if volatility increases. For example, a note might feature a five-year maturity and a 2 % semi-annual coupon plus 90 % of the appreciation in the S&P 500 above a strike price of 650, with a minimum redemption of par (i.e. no principal at risk). The investor retains \$0.90 of every dollar after the S&P exceeds the 650 strike. In the

¹⁷ In some cases, the appreciation/depreciation is paid via an enhanced coupon as it is accumulated, rather than via principal redemption.

opposite form of this structure, investors write the issuer equity calls or puts in exchange for a higher coupon. For instance, under a call-based note, the investor receives a higher ongoing coupon, but smaller principal redemption, as the equity reference rallies above the strike:

$$P * \left(\frac{Index_{mat} - Index_0}{Index_0} \right)$$

where

- P* is principal
- Index_{mat}* is the index level at maturity
- Index₀* is the index level at trade date.

The issuer of the note, which owns the underlying call option can, of course, hedge itself and monetize value by selling a similar or identical call at a higher price.

Put-based structures tend to be less popular than call-based structures, as investors often prefer to participate when markets are in an uptrend. This suggests that put-based structures may be cheaper, allowing a greater level of embedded investor participation. Other types of equity notes can be created, including those with synthetic collars, where an investor receives both a minimum and a maximum return; the investor is thus long a put and short a call via the note (meaning, of course, that the issuer is short a put and long a call). The note can be “tranching” by allowing investors to select from a series of preferred ranges (e.g. minimum 5 %–maximum 10 %, minimum 0 %–maximum 18 %, and so forth); these tranches represent collars with different strikes. Figure 5.6 illustrates the collared equity note.

In a similar fashion, coupons may be embedded with call spreads that set minimum and maximum boundaries for a particular quarter or semi-annual period (e.g. minimum appreciation of 1 % and maximum appreciation of 5 % in any quarter); the short call on the upside (e.g. sacrificing gains above 5 %) increases the investors’ participation rate. Other notes may feature embedded cliquet/ratchet options that lock in unrealized value, even if markets subsequently retrace; more extreme versions feature lookback options, supplying investors with the maximum gain achieved during the life of the transaction. Any option structure that features such “enhanced” profit opportunities is likely to be more expensive than one with a vanilla option.

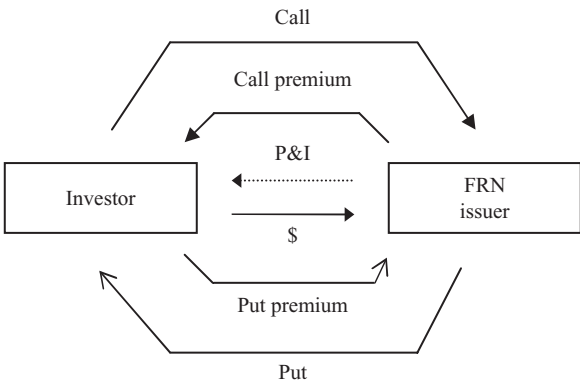


Figure 5.6 Equity note with embedded collar

Another form of equity-linked structure is the synthetic equity warrant, also known as a covered equity warrant (this warrant is distinct from the bond/equity warrant package we describe in Chapter 8, which is an instrument that leads to the flotation of new shares). The synthetic warrant is akin to a long-dated third party call (or put) option that can be floated independently, or embedded in a host note. The original synthetic warrant structure was developed in the Euromarket of the 1980s, when Swiss investors sought a greater amount of participation in the Japanese markets without assuming any currency risk. Accordingly, US investment banks issued Swiss franc denominated synthetic warrants to investors, and hedged their exposures through a combination of the actual dollar or yen warrants and foreign exchange contracts. Since that time, equity warrants have expanded to include a variety of options, including knock-outs, lookbacks, and spreads.

Equity-linked structures generally feature a modest level of secondary liquidity, primarily from the underwriters/dealers responsible for structuring the deals. In general, however, the securities are intended as “buy and hold” investment strategies. Publicly issued notes are listed on one or more exchanges; those that are listed must adhere to specific exchange requirements related to minimum distribution size, financial standards, and so forth.

5.3.6 Credit-linked notes

The broad class of CLNs includes basic CLNs, repackaged (or synthetic) bonds, and credit portfolio securitizations; we focus our comments on basic CLNs and repackaged bonds, leaving the topic of credit portfolio securitizations (collateralized debt obligations) for the next chapter.

CLNs, which are combinations of credit-risky assets and derivatives, or risk-free bonds and credit derivatives, can be issued through highly rated host issuers or dedicated repackaging vehicles. We distinguish between standard CLNs, which function as pass-throughs by forwarding relevant cash flows from a risky asset to investors, and repackaged bonds, which alter or restructure cash flows before forwarding them to investors; the latter is similar, in concept, to the asset swap structure we discuss in Chapter 10.¹⁸

The basic CLN is a package consisting of a low-risk bond or FRN and a simple or structured credit derivative, and provides investors with the returns of a credit-risky investment without the need to own the underlying credit asset. In certain instances, the host fixed-income vehicle may be a loan, rather than a bond or FRN, though this structure is less common.¹⁹ The “vanilla” CLN generally is issued at par with a payout calibrated to a defined credit event. Since payout is dependent on the credit event, the investor typically serves as the credit protection seller. As with other structured notes, the issuer often acts as a “pass-through,” selling the credit derivative it purchases from investors to the bank arranging the transaction, in order to reduce funding costs. The vanilla structure can settle in cash or physical: if a credit event occurs, the note terminates and the investor receives as principal repayment the cash difference between the

¹⁸ CLNs generally are synchronized to credit events defined through the ISDA Credit Derivative definitions (e.g. 2003 with supplements as they are negotiated). The buyer of credit protection (i.e. the issuer) generally is the “notifying party” that triggers the credit event under a CLN or credit derivative. However, in some instances, both parties (buyer and seller) can act as notifying parties. Apart from market-based credit spread movements associated with nondefaulting situations, the definitions are designed to cover bankruptcy, failure to pay, repudiation, moratorium, and restructuring.

¹⁹ Examination of the CLN structure suggests that risks must be considered on two levels: the performance obligation of the note issuer/asset repackager, and market/credit default elements of the reference obligor. The two dimensions can be considered separately, though, in practice, ratings agencies and sophisticated investors tend to view them jointly on an expected loss basis. Note issuers/asset repackagers generally are sponsored by highly rated financial institutions (e.g. AAA–A), so the risk of nonperformance is regarded as quite small; that said, legal risks may exist that could lead to performance failures, so caution is necessary.

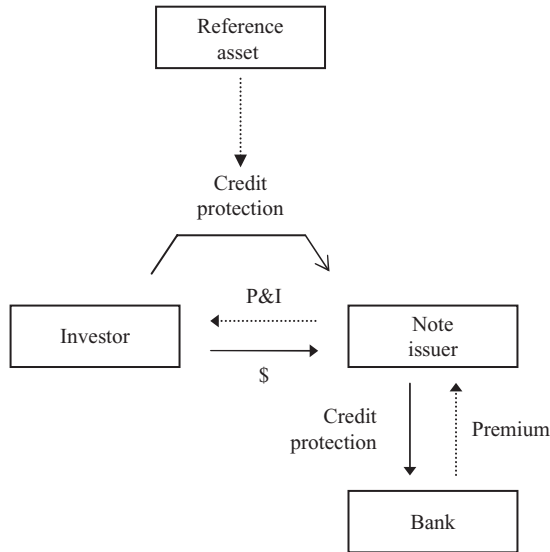


Figure 5.7 Basic CLN

asset’s value pre- and post-default, or it delivers the reference asset and receives the post-default market price. This basic structure is illustrated in Figure 5.7.

CLNs can be created using total return swaps (TRSs), credit spread options, credit spread forwards, or default swaps, which we discuss in greater detail in Chapter 10. For purposes of this chapter, we note that TRSs are contracts that provide the total return receiver with the economics (coupons plus capital appreciation) of a reference credit asset, in exchange for payment of depreciation plus a LIBOR spread. A default swap provides the buyer with protection against one or more credit reference defaults in exchange for a premium payment; similarly, the credit spread option provides the buyer with a payment if the spread on a reference credit widens or tightens. The credit forward involves the forward purchase/sale of a credit spread reference between two parties.

The TRS-linked CLN replicates the economic returns of a credit-risky asset by combining the TRS with a fixed- or floating-rate host bond. For instance, an investor can purchase a CLN that pays a LIBOR spread based on a bank’s credit; the spread is greater than might be obtained by purchasing the bank’s liabilities directly, because the bank uses the CLN vehicle to receive the total return of an underlying asset, paying the spread in exchange; this structure is illustrated in Figure 5.8. Alternatively, the investor can receive the total return of a risky asset through the CLN vehicle. The bank pays the vehicle the total return, which passes it on to the investor, creating for the investor an exposure to the economics of the risky reference asset. TRS-linked CLNs can also be based on an index or basket of credits, rather than a single reference, providing for a greater level of diversification. Banks often favor this structure, as it allows for hedging of single or multiple credit references through creation of a synthetic, rather than actual, short position.

Any of the structures noted above can be leveraged to provide investors with the opportunity of increasing returns. Setting the notional value of the TRS greater than the face value of the reference bond creates leverage. The note coupon is thus equal to (leverage factor * TRS

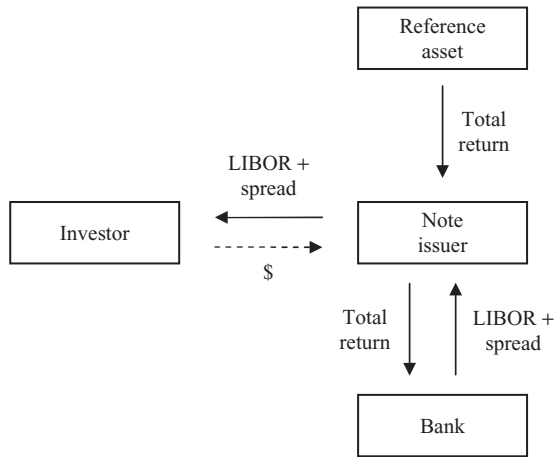


Figure 5.8 TRS-based CLN

margin) + LIBOR. If the leverage factor is large enough, a cap generally is included in order to limit losses on the face value of the note (e.g. to avoid negative redemption value); the maximum an investor can lose is thus principal invested.²⁰

The credit spread note, which embeds a credit spread call or put option (unilateral), or forward (bilateral), in a host security, creates an exposure to the spread movement of a reference asset. An investor purchasing the note monetizes a view on expected versus implied forward credit spreads. Thus, if the note is created using a call option, the investor anticipates a tightening of spreads; if a put option is used, then a widening of spreads is expected. In exchange for purchasing the put or call, the investor pays a premium through a lower-coupon yield. A note created using a forward, rather than an option, features bilateral flows. Thus, if the investor expects a tightening (long the forward) and the reference spread tightens, principal redemption increases; if the reference spread widens, the principal redemption decreases. The credit spread in a note is always defined as the yield to maturity of the reference security, less the yield of the risk-free benchmark; this allows the interest rate risk of the position to be eliminated. As with the TRS-based note above, leverage can be added by increasing the face value of the structure; alternatively, creating an off-market spread can enhance yield. The note can also be based on the average spread of a basket of credits, rather than a single reference, and can also include maturity mismatching (as discussed below).

The credit default CLN, which packages a host bond with a credit default swap (CDS), is similar in form and function to the credit spread note, but generates a payoff that relates only to the default of the reference asset (while CLNs referencing the default of a single reference asset are most common, they can also be created using basket swaps, or first-to-default basket swaps). The default note allows an end-user to invest in default risk (through the embedded sale of the CDS) and an issuer, such as a financial intermediary, to hedge default risk (through the embedded purchase of the CDS). Notes issued by corporates, rather than financial institutions, generally include back-to-back swaps between the corporate issuer and

²⁰ In fact, such an “extreme” risk profile is relatively unusual; most structures have some degree of principal protection (e.g. 50–90 %).

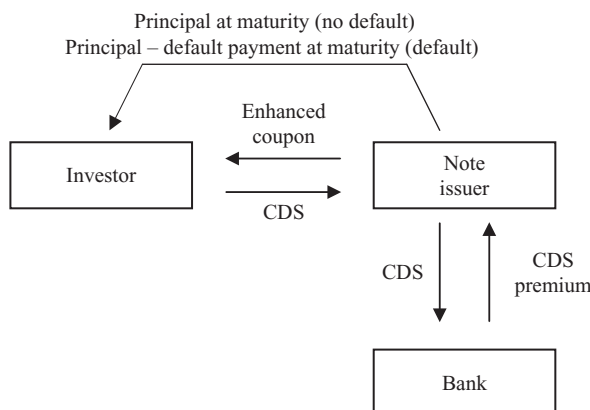


Figure 5.9 Default-based CLN

the bank, in order to transfer the CDS. Under the standard structure, the note provides the investor with an enhanced yield based on the value of the CDS being sold, along with par redemption at maturity if no default occurs, or par redemption less a default payment if an event occurs. The default payment can be set as a percentage of recovery, or the change in the reference asset price between issue date and the default date. The default event can be defined specifically or generally, though the industry's move towards standardized definitions typically means that bankruptcy, restructuring, cross default, failure to pay, repudiation, and moratorium are included. Figure 5.9 summarizes the default-based CLN structure.

The CLN market features a number of variations on the basic structures noted above. For instance, a CLN may include a knock-out feature, which pays the investor a lower return if the reference asset trades above or below a particular spread. This is simply a package of a bond, a credit forward (providing payout within a range), and two digital options (creating the upper and lower range boundaries). Thus, a note might redeem principal at 120 % of face value if the reference spread trades between 75 and 100 bps, but only 100 % if it trades below 75 or above 100 bps. The enhanced return that is generated when the spread is in the range is created through the investors' sale of the two digital options. Assuming that the bond is issued through a company, rather than an SPE or asset repackager, the note issuer, which buys the digital options from the investor, sells them to a dealer in order to lower its all-in cost of funds.

The callable credit default note is an extension on the standard default CLN. This structure, which is a package of an FRN, CDS, and issuer call option, functions just like a default note, but gives the issuer the right to call the note back at regular intervals (generally every quarter or semi-annual interest payment date). If the bond is not called, the coupon steps up to a higher predefined margin over LIBOR. Through this structure the issuer determines, on each call date, whether or not to preserve the default protection; if the note is not called, it may indicate that the reference credit is deteriorating, meaning that the cost of protection via the higher coupon is justifiable. A further variation is the principal-guaranteed credit default note, which allows the investor to participate in the credit risk exposure without placing principal at risk; this can be structured as a risk-free bond and a call option on the credit reference. Another extension is a default CLN linked to the credit performance of the asset and continued convertibility in a local marketplace. This structure, which is often applied to local currency emerging market

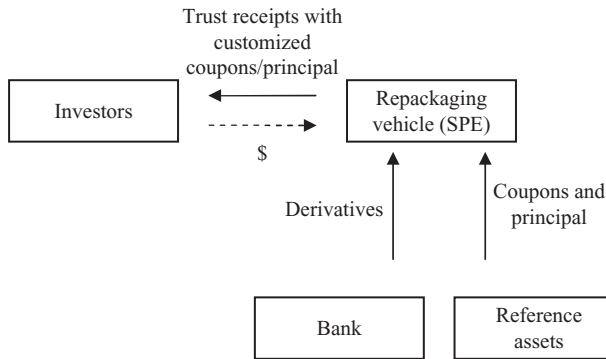


Figure 5.10 Repackaged bond flows

bond investments, where capital controls present an additional dimension of risk, features a coupon/principal payout that is a function of default and restricted convertibility. If either event occurs, principal redemption is adjusted.

Repackaged bonds comprise the second main sector of the overall CLN market, and have proven popular because of their payout flexibility. The repackaged (or synthetic) bond essentially is a securitized form of the asset swap, packaging credit-risky securities with derivatives into a synthetic bond.²¹ Issuance via an SPE, rather than a corporate or supranational conduit, creates greater flexibility and cost savings (e.g. no third party issuer compensation is required). In order to structure a repackaged bond, a bank determines investor demands and then acquires relevant credit-risky assets in the primary or secondary markets that it believes are trading cheap to theoretical value (e.g. assets that are perceived to be out-of-favor as a result of maturity or coupon). The bank then sells the assets to the SPE/trust, which enters into one or more derivatives with the bank to reshape the cash flows to a profile required by investors. The new cash flows are repackaged into notes/trust receipts, which are placed with investors; the proceeds from the note issuance are used to fund the initial purchase of assets. The resulting synthetic notes give investors access to customized credit-linked asset returns. Repackaging can be based on new and seasoned assets: primary market transactions involve repackaging of newly issued bonds, while secondary market transactions center on repackaging outstanding bonds, or reverse engineering structured notes/CLNs back into “vanilla” form. Figure 5.10 illustrates the basic repackaging process.

Repackaging trades are often used to create relative value opportunities for investors. For instance, this approach has been used in emerging market countries that have floated bonds under the Brady program.²² The repackaging vehicle purchases Brady bonds of a target country, attaches a currency swap related to the target country’s currency, and passes the combined flows to investors via a note; after accounting for the value of the Brady principal and interest

²¹ The use of a single bond, rather than a bond and swap package, generally is regarded by issuers as convenient and efficient. Such convenience generally commands a premium, with the SPE charging the investor a slight premium for mechanically arranging the package.

²² The Brady program is a debt restructuring framework, where sovereign nations with rescheduled bank loans and their creditor banks agreed to swap outstanding troubled loans for a new series of dollar-based bonds backed by rolling coupon guarantees and a US Treasury zero coupon bond covering principal redemption, all in order to create a liquid obligation; banks exchanging their loans accepted a discount for a higher-yield bond, or preserved par, but accept a lower coupon. The Brady programs have been extended to countries such as Mexico, Brazil, Argentina, Poland, Venezuela, and others, and have proven very successful.

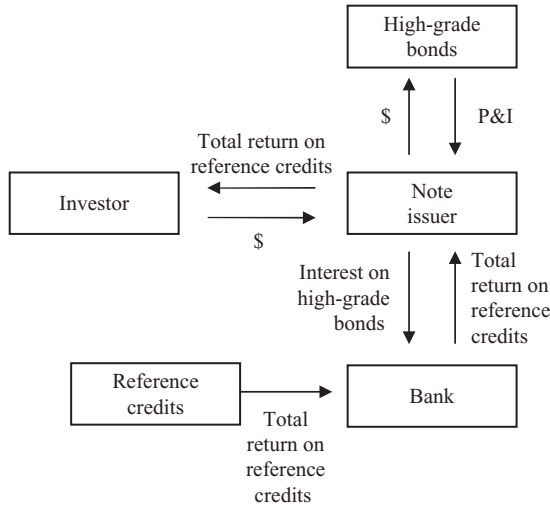


Figure 5.11 Leveraged TRS bond on portfolio of reference credits

protection, the yield on the package provides a relative value pickup versus straight local currency sovereign issuance.

Repackaging vehicles are also used for maturity shortening trades, a very popular strategy that allows investors to create an optimal investment horizon by decoupling the maturity of the reference asset from a desired maturity horizon. Under this process, the repackaging vehicle purchases a medium- or long-dated credit asset, and then arranges a credit derivative transaction (e.g. TRS, CDS) with a bank, synchronized to the investor’s preferred horizon. For instance, in a TRS version of the trade, the investor receives periodic coupons from the security until trade maturity, after which the asset, which may still have several years remaining until maturity, is sold in the secondary market. If the asset has gained in value, the investor receives principal plus the gain; if it has depreciated, the repackaging vehicle withholds the differential from the principal redemption. A further structure provides investors with a return based on a leveraged portfolio of reference credits: the repackaging vehicle issues notes to investors and uses the proceeds to purchase high-grade assets, and then arranges a TRS with a bank based on the desired portfolio of reference credits, receiving the total return in exchange for paying the coupon flow on the asset pool. The vehicle forwards the total return on the reference assets received from the bank through a structured coupon; in order to leverage the returns, the notional of the TRS is made larger than the face value of the bonds. At maturity, the notes pay off a principal amount that is linked to the value of the reference assets in the TRS; any appreciation is paid in the form of enhanced redemption, while any depreciation results in discounted redemption. This repackaged leveraged TRS note is depicted in Figure 5.11.

Repackaged notes, like CLNs, can be used to manage convertibility risk. For instance, an SPE can purchase a sovereign nation’s local currency obligations and issue dollar-denominated notes to investors in exchange (which pay dollar coupons and principal through an FX swap arranged between the SPE and a bank). The SPE then arranges a separate note with a bank, where final redemption is based on the onset of a convertibility event; if an event occurs, the bank receives a lower redemption amount (i.e. it bears the convertibility risk) and the investor receives the full redemption amount (i.e. it receives the benefit of convertibility protection).

Note that the investor continues to bear the sovereign nation's default risk, as payout depends only on the convertibility event.

Synthetic bonds are a further extension of the CLN family, and can be used to replicate an issuer's existing bonds. Though synthetic bonds are similar to asset repackagings arranged through an SPE, they feature several key differences: the structures are arranged by banks that are seeking to hedge risks, rather than investors seeking return opportunities (meaning the deals are often much larger in size); and, the transactions generally are launched in the public, rather than private placement, markets, allowing a certain level of liquidity to build. For instance, a bank may issue a \$100 m synthetic bond on reference credit ABC through a repackaging SPE or trust; the investor purchases the note, and the SPE uses the proceeds to buy a high-grade asset, such as US Treasuries. The bank, through a TRS with the SPE, pays the investor the total return on the ABC bond and receives the return on the underlying US Treasuries; the investor thus receives principal and interest on the \$100 m ABC credit as long as no default occurs. In the event of a default by ABC, the bank pays investors a discounted redemption price (via the SPE), based on market quotes. Investors face various dimensions of risk through the synthetic bond: default risk on ABC's credit, counterparty risk on the bank supplying the TRS flow through the SPE, and default risk on the underlying notes representing the SPE's obligations.

Pension and investment funds, hedge funds, insurance companies, and other financial institutions are all active buyers of CLNs. Some investors acquire a variety of notes, including those that are not principal-protected or which are of lower credit quality; others face investment restrictions and are only permitted to purchase principal-protected instruments or unleveraged structures.

Collateralized Debt Obligations

6.1 INTRODUCTION

The collateralized debt obligation (CDO), a mechanism that pools, securitizes, and redistributes corporate credit obligations, is a relatively new addition to the world of structured assets, with the first transaction having been introduced in the late 1980s. Despite this relatively short history, the CDO has been instrumental in transforming intermediary and investor perceptions of credit risk investment and risk management, and growth, through the early years of the new millennium, has been rapid. Since the generic CDO functions as a form of securitization, it can be considered an extension of the general class of ABS discussed previously. That said, the product features certain unique characteristics that differentiate it from its ABS peers.

We have seen that pooling is an essential characteristic of the ABS structure, providing scalability, flexibility, and operational/investment efficiency. By aggregating what might be a fractionalized and illiquid group of assets, the CDO is able to transform the pool into a scalable asset that can then be restructured into multiple securities with specific risk/return profiles. Many of the earliest CDOs were driven by regulatory capital pressures that led banks to transfer portions of their physical balance sheet assets, including loans and bonds, into securitization conduits. However, the market soon expanded into opportunistic, arbitrage-based deals based on a combination of physical assets and/or derivatives.

In the section below we consider the main forms of CDO, which can be viewed along multiple axes, including: product axis – collateralized loan obligations (CLOs) and collateralized bond obligations (CBOs); motivation axis – balance-sheet CDOs (risk transfer) and arbitrage CDOs (profit); cash flow axis – cash flow CDOs and market value CDOs; and structure axis – cash CDOs (funded, asset-based) and structured/synthetic CDOs (unfunded, derivative-based). Figure 6.1 summarizes these general CDO classes.

6.2 DEVELOPMENT AND MARKET DRIVERS

CDOs, as indicated, are relatively new to the global capital markets. The first CLO transaction was launched by NatWest (now part of RBS), which sought to shift a portion of its loan portfolio, and attendant credit risks, off balance sheet. The market for CDOs grew steadily from that point on, accelerating as more institutions came to realize the regulatory capital benefits that could be gained by transferring credit risk exposure (and as commercial analytics for valuation became more widely available).

CBOs, which were introduced several years after the first CLOs had entered the market, remained relatively inactive until the mid-1990s, when intermediaries and institutional investors began designing mechanisms to optimize the capital treatment of bond portfolios. Specifically, active high-yield investors, including insurance companies and pension funds, which faced regulatory capital challenges on their bond portfolios, drove the earliest CBOs. To circumvent capital problems, they used the CBO structure to repackage their holdings: the highest-rated tranches were retained within regulated entities, while the lower-rated tranches and residuals

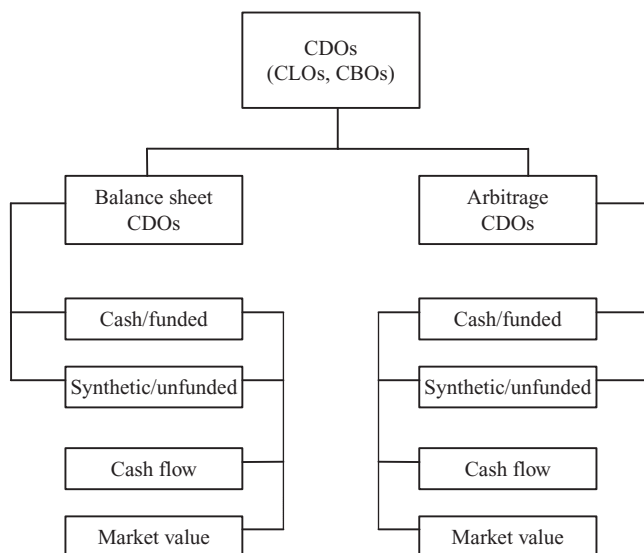


Figure 6.1 General CDO classes

were transferred to unregulated entities. This, of course, is similar to the driving force that caused many banks to transform their loan portfolios into CLOs. After loans and bonds were securitized successfully, intermediaries began repackaging structured financial assets using the same technologies; though the first deal entered the market in 1995, meaningful growth did not really commence until the end of the millennium. Structured finance CDOs backed by pools of MBS, CMOs, real estate investment trusts, and other asset-backed structures are now created and traded actively in the market. These instruments represent the pinnacle of structured assets, as they involve at least two tranches of separately structured cash flows.

The market for CDOs is based primarily on US and European asset pools and investors. The most significant transactions originated through the early part of the new millennium have come via large international financial institutions, with the resulting securities placed through their global institutional investor networks. While US and UK banks were pioneers in the original balance sheet CDO structure, they were ultimately joined by Swiss, German, Dutch, and French banking and asset management institutions. CDO structures are not yet prevalent in other countries, primarily as a result of inadequate securitization laws, and/or insufficient pools of local credit risk (i.e. too small or concentrated to be used in the development of a CDO).

CDOs are arranged for specific risk/financial management reasons, or to capitalize on market opportunities. For instance, a CDO:

- provides investors with professionally managed credit risk portfolio securities with a desired risk/return profile;
- allows credit originators to shift the physical and/or risk dimensions of their credit portfolios – the balance sheet CDO, for instance, allows the sponsoring institution to reduce its credit exposure to a series of borrowers simultaneously;
- optimizes capital and balance sheet usage – by arranging a CDO, the sponsor gains financial flexibility; the transfer of credit-risky assets to end-investors allows it to lower its

credit reserves and internal/regulatory capital allocations, and decrease its balance sheet footings;¹

- permits new credit originations – thus, if a sponsor is comfortable with its level of reserves and capital allocations, it can generate new business to replace transferred business;
- allows access to new sources of capital (i.e. capital markets investors, rather than traditional wholesale or retail depositors) and crystallization of term funding;
- gives intermediaries a chance to take advantage of market opportunities. The arbitrage CDO, for instance, is driven by a sponsor's desire to capitalize on perceived discrepancies between the fair and theoretical value of credit-risky assets, and to lock in a profit margin by acquiring the assets and funding them via the issuance of securities.

While risk management and arbitrage motivations have been key market drivers, arbitrage-related opportunities have gained momentum in recent years. Indeed, as financial institutions have become adept at using credit derivatives to manage their exposure levels, some of the original risk/capital motivations have begun to subside. It is important to note that, despite these powerful drivers, CDOs cannot always be arranged at will. An environment featuring very tight credit spreads is not conducive to risk transfer or arbitrage transactions, as the spread available to investors may be insufficient (unless the collateral pools are comprised solely of very high-yielding, and thus risky, assets).²

Investors may participate in the market if they are seeking to invest in repackaged credit risk managed by professional portfolio managers – in a specific form that matches their desired risk/return profiles. This appears to be particularly relevant for investments based on higher-yielding assets, where a diversified, professionally managed portfolio of assets may be considered preferable to investment in single-obligor exposures. Senior CDO tranches are often placed with high-grade asset buyers (e.g. funds restricted to AAA/AA risks that are primarily interested in capital preservation); subordinated tranches, in turn, are sold to sophisticated institutional investors (e.g. insurers and hedge funds that are eager to assume subordination risk in exchange for a yield pick-up of 10–50 bps); residual/equity tranches generally are retained by the sponsor and/or portfolio manager, and may also be placed with hedge funds (since the residual/equity tranche is the riskiest component of any CDO, it has the greatest exposure to the true skills of the portfolio manager). Investors have discovered that CDOs offer good value relative to similarly rated corporate bonds; this appears to hold true across a variety of market cycles. CDOs generally offer yield pick-ups of 50–60 bps over comparable corporate bonds, and feature less default risk as a result of the diversified nature of the collateral portfolios; they are, however, less liquid.

6.3 PRODUCT MECHANICS AND APPLICATIONS

Before introducing the specific CDO structures that can be formed along the axes described above, it is useful to review a series of general structural issues that are applicable to all types of CDOs.

¹ Though capital treatment varies, a bank transferring credit-risky assets into a CDO may only face an 8 % capital charge on the unleveraged amount of the transaction, or 100 % on any new liability it retains (e.g. the small residual tranche); by some estimates, these can represent up to a 50 % saving in regulatory capital.

² In such instances, there may be greater activity in investor-driven, single-tranche synthetic CDOs. Since the deals are entirely synthetic, no cash purchases are required, and since only one tranche is placed with one institutional investor, the degree of acceptable enhancement can be established in advance. Indeed, the investor generally has significant input into the selection of the portfolio manager and the initial portfolio.

6.3.1 General structural issues

CDOs can be structured in various forms, each designed to achieve particular investor/intermediary goals, or take advantage of specific market opportunities. The standard CDO, which has an average life ranging from 5 to 15 years, involves a sponsoring institution (generally a large bank), portfolio manager, end-investors, and asset issuers (bonds) or borrowers (loans). The sponsor establishes an SPE or trust in a tax-friendly jurisdiction³ that purchases the relevant credit assets through the issuance of tranches of securities, each with its own risk and return characteristics; in practice, the credit assets are warehoused by the sponsor and are sold to the SPE/trust once the liabilities have been issued and good funds are available. The SPE/trust also engages a trustee, custodian, paying/settlement agent, and portfolio manager to manage the operational and asset requirements of the deal.

A typical CDO pool may include from several dozen to several thousand obligors/issuers from different industries. An experienced portfolio manager is appointed to manage the portfolio⁴ under the supervision of the trustee; the portfolio manager must adhere to the investment parameters set forth in the indenture and prospectus, as well as any additional guidelines imposed by the rating agencies. The risk of the collateral pool depends primarily on credit quality and diversification: better credit quality leads to lower default risk, but lower returns; greater diversification reduces the variability of losses, but again generates lower returns. Note that the portfolio manager often retains a significant portion of the residual/equity first loss tranche,⁵ and thus has incentives that are aligned properly with a goal of maximizing earnings while taking prudent risks.

The sponsor and portfolio manager work together to construct the portfolio. The process passes through various distinct stages: in the ramp-up phase, or the accumulation period prior to deal-closing date, the sponsor and portfolio manager assemble or purchase a portfolio of assets; in the reinvestment phase, principal from maturing or amortizing assets is used to acquire new credit assets (note, however, that not all CDOs feature a reinvestment phase); in the wind-down phase, principal from maturing assets is returned to investors, creating a note redemption effect.

A typical CDO, which is often created in private placement form, may feature one or more senior tranches (AAA/AA-rated), one or more subordinated tranches (BBB- to B-rated), and a residual or equity tranche (unrated),⁶ all backed by the pool of credit assets.⁷ The basic structure

³ The tax haven locale is essential in order to reduce friction costs and maximize returns; the downside may be relative lack of transparency on certain reporting/regulatory issues.

⁴ The portfolio manager performs a crucial role in constructing and managing the collateral pool. Accordingly, it is essential for a manager's experience and past performance record to be vetted thoroughly. A proper alignment of economic interests can be a considerable help in the process, ensuring that the risks contained within the fund also drive the portfolio manager's own economic result.

⁵ This share may rise up to a maximum of 49 % in order not to breach nonconsolidation accounting rules; any remaining balance may be sold to hedge funds and other institutional investors.

⁶ The market value of the equity tranche can be viewed as the mark-to-market value of the asset portfolio less all outstanding liabilities/obligations, including senior/subordinated notes issued, other borrowings, and service fees.

⁷ Note that some transactions also feature participating coupon notes (PCNs), which are subordinated debt/equity hybrids, paying coupons and a percentage of any equity returns. Though the PCN structure dates back to the mid-1990s, deals did not actively feature such tranches until the early years of the new millennium. The increased use is due to the fact that many equity CDO investors have allocated their capital in support of existing transactions, making it difficult for sponsors to sell the equity risk to traditional buyers. The PCN has thus surfaced as a mechanism to broaden the investor base. The earliest iterations combined CDO equity and a zero coupon bond in a trust, which was rated explicitly with respect to the zero coupon principal guarantee. After encountering certain tax/regulatory difficulties, the next version combined notes and equity from the same deal through internal structuring (and not via a trust). Investors in the PCN receive an investment-grade tranche with a below-market coupon, and an additional equity-related cash flow. Some PCNs are rated on the basis of principal only, or principal and below-market coupon. If an investor requires an A-rated security for principal and coupon, then the PCN can be tailor-made to accommodate that requirement (e.g. perhaps 90 % debt, 10 % equity component). If it needs the A rating only for the principal, the percentage split can be shifted considerably (e.g. 60 % debt, 40 %

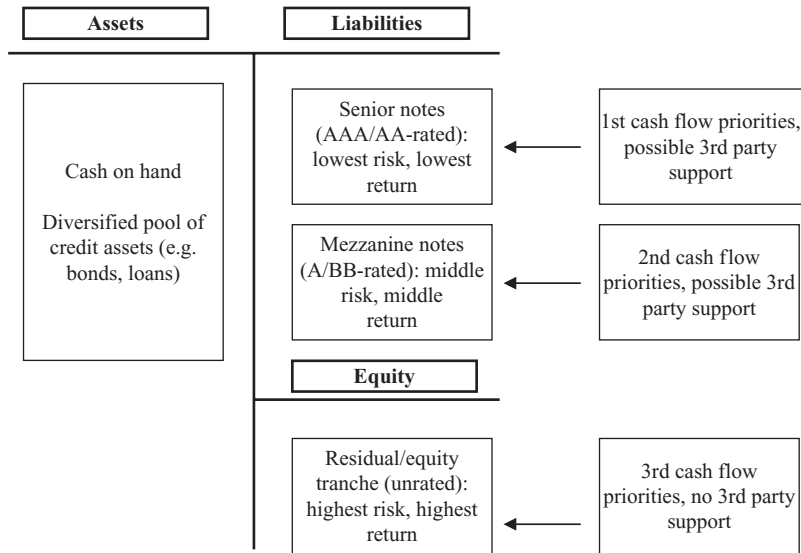


Figure 6.2 The general CDO balance sheet

can be viewed in standard balance sheet form, as noted in Figure 6.2.⁸ The more rated debt tranches a deal has, the less residual/equity it requires; this can lead to lower returns but a more secure structure. For instance, a high-grade CDO, with lower default probabilities in the asset pool, requires a smaller residual/equity tranche than a high-yield CDO (3–5 % versus 8–10 %); by extension, the larger the equity tranche, the smaller the resulting AAA/AA-rated tranches. CDOs generally feature a predominance of high-rated tranches in order to take advantage of cheaper funding levels. This is particularly critical for arbitrage structures.

Tranching and cash flows

The actual tranching structure of a CDO, and the amount of subordination and equity, depends ultimately on the size of the profit spread; this is true of all CDOs, not simply the arbitrage CDOs considered below. The portfolio of credit assets, financed by a combination of debt tranches, must yield enough to service the debt obligations and provide some return to the equity investors. The driving factor is whether the available returns are sufficient to attract the right combination of investors. If the returns appear to match investor expectations, then

equity). If default rates are lower than expected, the PCN outperforms the conventional rated tranche, as a result of the “equity kicker.” Other transactions include a “PIK” feature, allowing for payment-in-kind if current coupons cannot be serviced with cash. Some CDOs permit the trustee to halt current coupons to investors under certain circumstances (though sometimes these can be interpreted very broadly, e.g. the current coupon may be halted even if a sufficient amount of cash exists to service the notes); in exchange, the par value of the tranche is increased by an amount equal to the suspended coupon. This is precisely equal to the PIK structure that is used commonly in the high-yield markets. The PIK CDO allows the portfolio manager to conserve cash or redirect it to pay down the principal on senior tranches in order to avoid any violation of the overcollateralization or interest coverage tests (any breach of which would create a deleveraging). In fact, the more stringent the tests, the more likely it is that cash flows from the deal will be diverted, and the greater the likelihood that subordinated tranches will become PIK-able. Investors need to approach PIK CDOs with caution, as they may be unaware of the fact that tranches that they purchase are “deferrable” in this manner; indeed, the rating agencies may not award an A rating to a PIK-able tranche.

⁸ It is worth noting that, although the equity tranche serves the same purpose as common equity (e.g. first loss capital), it is actually issued as a form of debt security in the CDO structure.

the CDO's structural features can be refined and stress-tested to determine performance under a range of plausible and implausible scenarios.⁹ If the returns appear to fall short under plausible scenarios, the CDO generally must be abandoned. The modeling process must also incorporate the effect of fees (both upfront – which are generally 1–3 % – and annual – 0.50–1 %), funding costs (based on a range of forward curve stresses),¹⁰ delayed portfolio ramp-up, and default and recovery scenarios; for CDOs created with structured finance assets, such as MBS, prepayment effects must also be evaluated. Only after completing these exercises will it become clear whether the CDO can be structured profitably, and whether the cash flows generated by the securitized pool will be sufficient, under extreme scenarios, to service noteholder obligations. When the modeling scenarios suggest that debt obligations can be serviced successfully and the equity tranche will generate a return of at least 15–20 %, then a successful transaction may be at hand. The sponsor can then adjust the details of the tranches, leverage, and asset pool to optimize risk/returns. Naturally, a CDO that appears to “work” based on overly generous assumptions, may actually fail to provide the intended returns.

Cash flows are distributed in sequence based on seniority; the deal prospectus contains information on the structural “tests” that are performed to ensure proper cash flow coverage, as discussed below. A CDO in compliance with its tests typically pays cash flows as follows: fees/expenses of the CDO SPE (typically 25 bps upfront, and 25 bps after satisfying the structural tests below¹¹), interest to senior tranches, interest to subordinated tranches, interest to the equity tranche, principal to senior and subordinated tranches until retired, and then principal to the equity tranche. However, if tests reveal a cash flow shortfall, the payment stream is redirected in favor of the senior-most tranches.

Returns on subordinated tranches can be examined on a relative value basis versus corporate bonds, though they encompass a different set of risks. Default risk differences, in particular, must be considered closely (e.g. a BBB-rated pool of default risk versus a BBB-rated single-obligor default risk); a better relative value comparison can be made against a corporate bond index comprised of BBB obligors. Historical and estimated default experience between the two asset pools can be analyzed to determine whether the CDO's subordinated tranche appears cheap or rich to the BBB corporate bond index. The residual/equity tranche, which protects the senior and subordinated tranches that rank above it, has the potential of providing attractive returns if the pool performs well; the equity can thus be viewed as a leveraged position in the collateral pool's default performance.¹² Importantly, it occupies the first loss position and is thus at considerable risk; violation of any of the structural tests can seriously impair the value

⁹ In practice, a hypothetical portfolio of assets is run through the sponsor's model to compute potential returns. Different model variables can then be added (e.g. changing portfolio credit quality, diversification, concentrations, interest rates, and so forth). These will impact the structure of the deal and reveal whether more or less returns need to be generated. For instance, increasing the credit quality of the portfolio means reducing the specter of default risk, which means lowering overcollateralization and reducing the equity tranche needed to support the overall structure. Since equity returns are greater (though more variable) than senior/subordinated returns, the cost to the sponsor declines. The opposite is also true. Alternatively, the model can be run with more subordinated tranching and less equity, thereby increasing deal leverage; while greater leverage can lead to greater returns, it can also magnify default-induced losses and must be subjected to tighter overcollateralization and interest coverage tests (any violation of which will cause a more rapid deleveraging).

¹⁰ Funding costs are vital in determining the amount of subordination; the higher the interest costs, the lower the level of subordination, as junior tranches are more expensive than AAA tranches. That said, too little subordination can impact AAA participation: AAA investors tend to focus heavily on the amount of subordination that exists in a deal as their first loss protection; if they consider that it is too small, they may be reluctant to participate, or may demand a greater risk premium.

¹¹ Split fees are increasingly common and are used to align the portfolio manager's interests with those of other investors. By subordinating half of the fees, the manager has incentives to manage the portfolio very closely.

¹² The CDO equity can also benefit from credit spread tightening versus LIBOR if the interest rate risk is hedged versus floating coupons. The equity also benefits from a high-spread environment, as investors will receive excess coupons over the amount needed to service the senior/subordinated tranches.

of the instrument. Equity investors only receive cash flows after the CDO's fees have been paid and the senior/subordinated tranches have been serviced. If the tests fail, it is very difficult for the equity to later generate significant returns, as the structure will already have deleveraged by a certain amount. Note that equity investors tend to receive a disproportionate amount of their cash flows in the early part of a deal (e.g. before any credits in the pool have defaulted), meaning that seasoned equity tends to sell at a discount to the flotation price.¹³ The potential value of the equity tranche can be analyzed through standard internal rate of return (IRR) scenarios, including scenarios where defaults shift IRRs from positive to negative. For instance, with 0 % defaults (an unrealistic scenario), the IRR of the equity may exceed 25 %; with 2 % defaults (an unlikely, but possible, scenario), the IRR may drop to 0 %. In practice, a well-managed CDO, with default experience near historical averages for a pool of a given set of characteristics, can generate an IRR of 15–20 %. More sophisticated analyses, including those performed via simulations, allow potential equity investors to examine the distribution of equity returns, and establish confidence levels across a range of default and recovery experience.

The portfolio manager generally is required to sell defaulted assets in the pool within one year, and may also have to mark the position at the lower end of the post-default market price or some fixed price (e.g. 30 %). The underlying assets comprising the pool generally have final maturities that are less than, or equal to, the maturities of the securities. In practice, the actual average life of pool assets tends to be shorter than the maximum life, meaning that a given deal may have a higher yield to maturity than is otherwise apparent.¹⁴ As noted in the ratings section below, the sponsoring institution or portfolio manager attempts to structure and maintain a properly diversified asset pool, avoiding undue concentrations by obligor, industry, country, or ratings class. CDOs actively use swaps, caps, floors, and collars to transform cash flows from fixed into floating rate, or to hedge any interest rate or basis risk between the asset pool flows and the securities. Since hedges are rarely perfect, most CDOs feature a certain amount of outright overexposure or underexposure to market factors. Transactions generally require about six months of lead-time to structure and arrange (depending on portfolio complexity, asset liquidity, and market conditions). Figure 6.3 illustrates the general CDO structure.

As noted above, a CDO must undergo certain tests to confirm the investor payment sequence (e.g. the “waterfall”). Consider a CDO comprised of three tranches, Class 1 (AAA-rated), Class 2 (BBB-rated), and Class 3 (BB-rated). In order for principal and interest to flow from one class to the next, overcollateralization/interest coverage tests must be passed successfully.¹⁵ The process begins with an interest income flow, which, after being used to cover fees, is used to pay Class 1 interest. The Class 1 test is then performed: if the test results are positive (e.g. the asset/interest coverage ratios are adequate), Class 2 interest is paid; if the test results are negative, Class 1 principal, rather than Class 2 interest, is paid – by using excess interest to pay Class 1 principal, the senior-most investors are protected. The paydown of senior investors is equal to early amortization of the structure, which creates a deleveraging effect. Assuming the test is passed, however, the waterfall structure then moves to a Class 2 test: if the test is passed, Class 3 interest is paid, and any remaining cash flow is allocated to the residual; if the test fails, additional Class 1 principal is paid, and then Class 2 principal is paid. The structure

¹³ The value of the CDO equity that has experienced no defaults also declines over time as cash flows are returned as interest or amortized principal.

¹⁴ Note that although the collateral pool may contain callable securities, in practice, these are unlikely to cause a deal's average life to vary by more than 12 months; the effects of this optionality are thus limited.

¹⁵ Note that it is increasingly common for CDOs to feature separate overcollateralization and interest coverage tests for each class of notes.

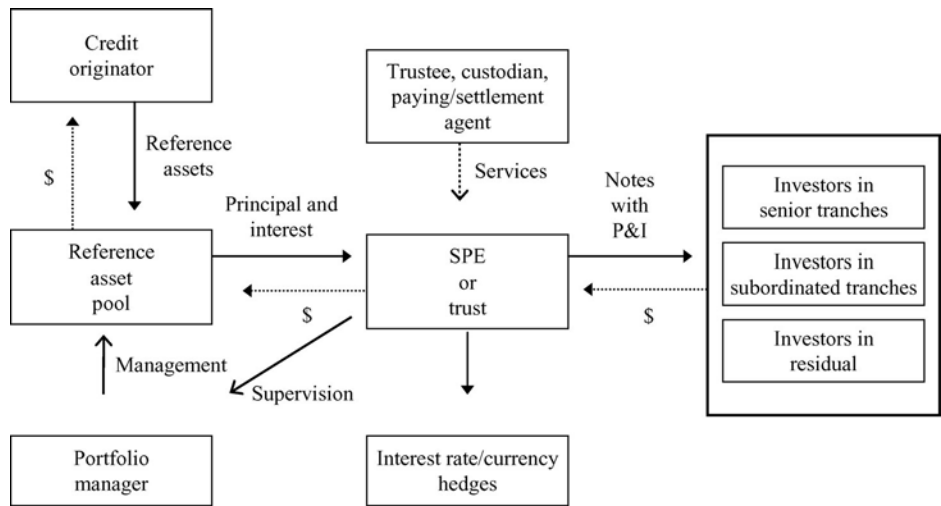


Figure 6.3 General CDO structure

thus continues for each payment period. The payment sequence of interim and final cash flows on a standard CDO that remains current on its tests is summarized in Figure 6.4.

Certain standard tests are used to evaluate the allocation of principal and interest via the waterfall structure; these include overcollateralization, interest coverage, and quality tests. The overcollateralization test is given as follows:

Principal value of collateral portfolio

(Principal for a tranche + Principal for all tranches ranking senior)

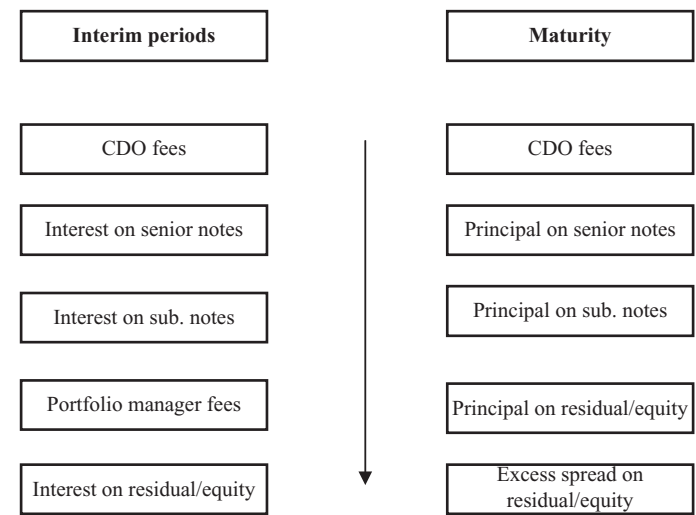


Figure 6.4 CDO interim and final cash flows

The results of the overcollateralization are compared to overcollateralization triggers. If the computed figure is greater than the trigger, the test is passed successfully and the waterfall structure remains intact; if it is less than the trigger, the flows are redirected, in the manner described above. Not surprisingly, a series of defaults within the pool, which subtracts from the principal value of the collateral pool, may cause a test to fail. Typical overcollateralization ratios at a BBB-rated level might range from 102–105 % for high-grade asset pools, and 105–115 % for high-yield asset pools. It is worth noting that CDOs that use a greater amount of collateral selling at a discount to par have a greater level of implied overcollateralization, allowing for a lower equity requirement. Thus, deals created during severe market downturns or corrections (e.g. 1997, 1998, 2000) can feature significant overcollateralization buffers. The reverse is also true: a widespread, systemic spread widening event that occurs after a CDO has been created, can lower the principal value of the collateral pool by enough to create test failure – even though no defaults occur.

The second major test, interest coverage, is computed as:

$$\frac{\text{Interest on collateral portfolio}}{(\text{Interest for a tranche} + \text{Interest for all tranches ranking senior})}$$

Again, the results of the computation are compared to guidelines and explicit coverage triggers; a successful result leads to a continuation of the waterfall, an unsuccessful result leads to cash flow redirection. Typical interest coverage at the BBB-rated level can range from 100–105 % for high-grade asset pools, and 110–120 % for high-yield asset pools.

As noted, if the overcollateralization and/or interest coverage tests fail, principal and interest receipts from existing pool assets are redirected to the senior-most tranches and the structure begins to delever. The equity investors absorb the majority of this redirection. The process of deleveraging and redirecting cash flows continues until the proper coverage levels are regained, at which point the CDO can resume its normal operations. It is worth emphasizing that the smaller the difference between the actual and trigger levels, the greater the structural leverage contained in the CDO, and the greater the risks and returns to the equity investors and subordinated tranche investors. CDOs constructed with an average asset pool rating of BBB tend to feature smaller differentials between the actual and trigger levels on the two tests than those comprised of BB or B assets, since the volatility of default rates is smaller.

Quality tests, which relate primarily to portfolio construction, diversification, concentration, expected default, and loss given default (e.g. 1 less the recovery rate), are also used. Portfolio assets are decomposed into industry groups to determine the actual level of diversification in the pool, given an assumed correlation between industry groups. Cross-industry diversification is regarded generally as a favorable characteristic of the asset pool. Pools with more diversification require less equity and can support more leverage; those with less diversification require more equity. Maximum allowable concentrations per obligor are also computed; in most cases, these cannot exceed more than 2 % of a given pool. Expected defaults based on estimated default probabilities are applied to the entire portfolio based on historical data (adjusted for possible anomalies). Defaults are value-weighted across the target portfolios to create a time-dependent cumulative probability of default. Recovery rates are taken generally as a historic constant based on seniority. Expected loss within the portfolio is thus a summation of (loss given default * probability of default) for each asset. For instance, if asset 1 in the pool (\$10 m notional) has a 5 % probability of default and a 40 % recovery rate, the expected loss is 3 %, or \$300 000. The same computation is applied to the next asset, and so forth. Default correlation features in the portfolio calculation as well. A portfolio with low default correlation has a

high probability of experiencing only a few losses, but virtually no probability of sustaining zero losses. A portfolio featuring high default correlation, in contrast, functions more as a single credit obligor: either many credits default, or no credits default. As default correlation increases, the probability of large losses grows, and the probability that the CDO's senior tranches will be exposed to default-related losses increases in tandem; this means that even AAA/AA investors will demand a greater return.

CDOs may be structured to allow some asset substitution if particular references in the collateral pool show signs of deterioration. Such substitutions, which may be capped at five or ten per year, are designed to avoid an event of default (which would permanently damage CDO value), rather than to promote opportunistic trading.¹⁶ If an asset is removed from the portfolio, it must be replaced by another asset in order not to deleverage the structure. If a capital loss is sustained in liquidating a troubled asset, the losses are reflected through the tranche coupon. Certain CDOs are structured so that they can be called at par by the sponsor after a blackout period; however, a "prepayment penalty" may be levied to discourage exercise of the call.¹⁷

A secondary market in CDOs has developed over the past decade, primarily for senior and some subordinated tranches; equity tranches have very limited resalability. However, since the market is still relatively new and each CDO has somewhat unique structural features, liquidity is not deep, and pricing inefficiencies exist. Hedge funds have emerged as aggressive players in the CDO market, and some have capitalized successfully on price discrepancies. Their ability to model CDOs and derive certain theoretical values that may be different from market values allows them to purchase "cheap" CDOs in the secondary market, and/or to sell "rich" ones. Most price support and market-making activity comes from underwriters initially responsible for issuing CDOs on behalf of the sponsor; this is logical, as an underwriter (generally one of the major investment banks) is most familiar with the structure and cash flows of a given deal. It is important to know that valuation becomes more transparent once a CDO exits the primary offering stage, as the portfolio backing a particular deal becomes widely known;¹⁸ this allows potential secondary investors to analyze the deal on its merits. Secondary investors buying seasoned CDOs must pay particular attention to an issue's overcollateralization and interest coverage tests, and whether the pool contains any concentrations of "borderline" securities that may trigger defaults at a future time horizon.

CLOs: high grade and high yield

While the generic CDO template is relatively straightforward, and is quite comparable to the securitization technologies used in the MBS and ABS market, specific details related to the construction of CLOs are rather more involved. Indeed, the legal and structural issues

¹⁶ There are, however, instances when substitution is allowed if certain assets have increased in value by a significant amount, in order to help crystallize reserves that can be used to protect investors.

¹⁷ A transaction may be called if it is performing well or poorly: if it is performing well, the sponsor generally is required to pay a call premium equal to some percentage of the annual coupon, and if it is performing poorly, the sponsor may not have to pay any such premium (in which case, the call functions as a deal "clean-up" mechanism). Some CDOs grant the equity holders the right to call the deal after the lockout period under certain predefined circumstances (generally performance-related). CDOs may also feature refinancing options that allow the portfolio manager to refinance the structure while leaving the asset pool intact. This occurs when the value of the liabilities has risen and it is unwise to sell the collateral (e.g. it is illiquid, or worth more if held to maturity); such refinancings are, however, relatively uncommon. Early termination is also possible, though only if the CDO's covenants are breached, the senior tranches fail to receive cash flows, the issuer of the CDO declares bankruptcy, or the portfolio manager/team departs.

¹⁸ The offering memorandum and trustee reports reflect the structure and current holdings of a CDO; this information also includes overcollateralization and interest coverage test levels.

surrounding CLOs can be particularly complex, primarily because loans are not always transferred easily from the original borrower into a collateral pool. A CLO can include various types of loans, including direct loans, participations/subparticipations in drawn or undrawn commitments, revolvers, leveraged leases, project finance loans, and distressed loans.¹⁹ A CLO may be arranged using loan participation agreements (sale without the knowledge/agreement of the borrower), meaning if the seller becomes insolvent, the SPE or trust may not have direct recourse to the borrower. Alternatively, it may be structured through the assignment process, which provides for a full assignment of the seller's rights, thus linking the SPE or trust with the borrower; this, however, is a more time-consuming method, as it requires prior notification and/or approval by each borrower. The difficulties associated with transferring loans easily are one reason why banks have begun using structured and synthetic CDOs more actively.

CLOs may be based on high-grade or high-yield loans; though high-grade conduits are easier to structure and sell, they offer investors lower returns. High-yield CLOs, which feature more attractive returns to compensate for the added default and liquidity risks, are more difficult to structure, as sourcing a sufficient amount of suitable high-yield loans can be challenging; indeed, deals originally intended as pure CLOs may become broader high-yield CDOs through the inclusion of high-yield bonds. In addition, it may be difficult identifying a portfolio manager with suitable high-yield experience, and investors interested in acquiring the residual tranche of the deal.

CBOs: high grade, high yield, emerging market, structured finance

CBOs can be collateralized by pools of high-grade, high-yield, emerging market, convertible, distressed, or structured finance bonds; on rare occasions, asset pools may contain a mix of bonds from each sector, though this demands considerable portfolio management expertise and very specific investor needs. While most assets in CBO pools tend to feature fixed coupons, floating-rate notes may also be used; this requires the SPE to arrange additional hedges to reduce or eliminate interest rate risk. Since bonds are traded/transferred among investors easily (unless they are structured as private placements), a CBO generally can be assembled more rapidly than a CLO. Creating a security interest in the pool of bonds for the holders of the securities is a well-established process that has been used widely in the MBS and ABS markets for several decades.

CBOs often use inefficiently priced illiquid securities that can be held in the collateral pool until maturity; high-yield bonds, emerging market bonds, private placements, and other structured notes/bonds tend to feature prominently in CBOs. Emerging market CBOs have become a popular component of the market; though emerging market borrowers have a significant amount of loans outstanding, bonds tend to be the asset of choice for collateral pools, as default experience is often much better.²⁰ Structured finance CDOs, comprised of other structured assets, tend to feature better credit quality than high-yield CDOs, but also feature

¹⁹ That said, a CLO comprised of an excessive amount of unfunded commitments or revolvers may not work, as cash flow uncertainties may simply be too great.

²⁰ For instance, emerging market borrowers have a long history of defaulting on, or renegotiating, their external loans, as well as their domestic (local currency) bonds; many, however, have avoided similar actions on their external bonds, as they prefer not to threaten their ability to access the international capital markets at future points in time (renegotiating terms with investors holding bearer bonds is also largely impractical). Consider that at the turn of the millennium, emerging market borrowers featured a combined default rate of nearly 12 %, but only 2.5 % involved external bonds. Default/restructuring actions can, in fact, be selective, meaning certain debt obligations may be damaged while others remain current.

Table 6.1 CLO and CBO collateral assets

CLO assets	CBO assets
Senior, secured bank loans	High-grade bonds
Senior, unsecured bank loans	High-yield bonds
Mezzanine loans	Emerging market bonds
Funded/unfunded revolving credit lines	Convertible bonds
Leveraged leases	Structured finance bonds (ABS, MBS, CMO)
Project finance loans	Distressed bonds
High-yield loans	
Distressed loans	

longer maturity tails and, on occasion, a degree of prepayment risk.²¹ Some portfolio managers include callable bonds in their collateral pools; since these securities generally are callable at a premium equal to a certain percentage of par or the current coupon, returns to CBO investors can improve by a commensurate amount.

Table 6.1 summarizes assets that are used commonly as collateral in CBOs and CLOs. Many of these assets also serve as references for structured and synthetic CLOs and CBOs.

Ratings

Rating agencies are contacted once legal structure and investment strategies have been developed, in order to obtain detailed guidance on portfolio construction, overcollateralization, and interest coverage. CDOs are rated typically through an analysis of various factors, including credit risk, market risk, legal risk, cash flow timing, expected credit losses, credit enhancements, portfolio manager experience, operational complexity, and use of third party support (e.g. monoline insurer guarantees, bank letters of credit).²² Each one of these factors may be subjected to certain stress tests to determine how the structure will perform under adverse conditions. The rating agencies often require CDOs to adhere to certain standards, including minimum pool diversification, maximum asset maturities, minimum weighted average pool asset ratings, minimum overcollateralization (e.g. asset principal greater than securities issued), and minimum interest coverage; they assign specific scores to many of these attributes, which influence ratings of various tranches within a deal. Rating agencies generally also require deals with a greater amount of leverage to use higher-quality asset pools and adhere to stricter overcollateralization and interest coverage tests. Conversely, those with less leverage can comfortably accept lower-quality assets.²³

Diversification is an especially critical factor, particularly for structures that may be subjected to industry- or country-wide difficulties. For instance, the rating agencies tend to follow a conservative approach in considering diversification attributes of emerging market CDOs; they often group emerging market assets into broad groups and give only partial benefit to

²¹ Since most structured finance CDOs feature less than one-third of their pools in residential MBS, negative convexity and prepayment risks are much smaller than in standard MBS or CMO structures.

²² Consider, for instance, the process used by one major rating agency in evaluating a CDO: the parameters of the hypothetical portfolio are evaluated by examining aggregate and individual industry scores; the weighted average probability of default is then determined by computing the expected loss of each credit in the portfolio; a stress factor is then applied to the weighted average probability of default on the entire portfolio; weighted average recovery rates are then applied, and a cash flow model is then used to determine investor losses under particular scenarios. The resulting expected loss of a senior or subordinated note is then mapped to a default table to obtain an equivalency rating.

²³ In certain unusual cases, the equity tranche (which generally is unrated), may be rated investment-grade if deal leverage is low (e.g. sub-8x) and the credit quality of the pool is high (e.g. average BBB or better).

cross-country diversification, in order to avoid damage that might arise from contagion effects. Borrowers within a single emerging market are given no diversification benefit.²⁴

Most CDOs rely on internal/external enhancement in order to achieve a particular minimum rating level for the highest-rated tranches within the overall structure. As we have noted, it is common to create standard senior/subordinated tranching in order to redirect cash flows to the senior tranches; alternatively, the structure may feature fast-pay/slow-pay tranches (with accelerated cash flows directed to the highest-rated securities). A CDO can also be enhanced using excess spread that accumulates in a reserve account, or a third party guarantee from a highly rated monoline insurer or financial institution. CDOs that feature lower diversification scores (e.g. emerging market or high-yield CDOs) must be created using higher structural standards in order to achieve the same ratings level as a well-diversified, high-grade transaction; this typically means more overcollateralization, greater subordination, and/or a greater amount of equity (e.g. the equity tranche of an emerging market deal may be as high as 20 %, while that of a high-yield deal may be 8–10 %, and that of a high-grade deal only 2–4 %).

While most early CDOs had limited/no right of asset substitution, that has changed with market evolution and introduction of new CDO structures. Right of substitution has expanded, making it more difficult for the rating agencies to monitor and enforce minimum standards. This is likely to remain a challenge as greater “flexibility” is built into new transactions. The rating agencies generally require the trustee to compile and report regularly on the target CDO’s portfolio, diversification, defaults, overcollateralization and interest coverage tests, excess spread account, asset turnover, and portfolio manager performance, in order to ensure adherence to required parameters and thresholds.

6.3.2 Cash flow CDOs and market value CDOs

A CDO can be created as a cash flow structure or a market value structure. The cash flow CDO relies primarily on principal and interest cash flows from the securitized portfolio to service outstanding notes, while the market value CDO depends on both principal and interest streams and active management (e.g. maturing assets, properly timed sales creating capital gains) to generate debt service. Most of the balance sheet CDOs described below are structured in cash flow form, while arbitrage CDOs may take either cash flow or market value form.

Under the cash flow structure, the sponsor and portfolio manager design the collateral pool of loans or bonds so that the assets generate sufficient principal and interest to pay investors periodic coupons, along with principal at maturity. Pool cash flows thus drive investor cash flows. Under the market value structure, the pool is subject to a daily or weekly mark-to-market evaluation and overcollateralization testing. If the market value of the portfolio times the advance rate²⁵ falls below outstanding debt, a portion of the pool²⁶ is sold in the

²⁴ For instance, borrowers in a non-Latin American region are assumed to have a default correlation of approximately 30 %, those in Latin America, a default correlation of 60 %, and individual borrowers within a country, a default correlation of 100 %.

²⁵ The advance rate reflects the discount, or “haircut,” on the value of the asset pool required by the rating agencies; thus, an 80 % advance rate implies a 20 % discount, a 70 % advance rate a 30 % discount, and so forth. Haircuts on high-yield bonds are approximately 5 %, distressed bank loans 10 %+, and so forth. Advance rates are critical to the success of a market value CDO; if they are set too high, investors may not be protected adequately, and if they are set too low, the economics of the structure may not function properly. The advance rate is a direct function of price volatility and liquidity of the asset – the greater the volatility and the lower the liquidity, the larger the discount and the smaller the resulting advance rate. Each asset in the pool features an advance rate that is related to the target rating, based on the structure of the deal and the composition of the portfolio. The advance rate also increases with portfolio diversification, at a level that depends on the correlation between pool assets (correlation may be measured under normal or stressed conditions, depending on the specifics of the collateral pool); thus, the lower the correlation, the greater the diversification benefits, and the greater the advance rate.

²⁶ In practice, the amount of collateral that needs to be liquidated is computed as deficit/(1 – advance rate).

market and the proceeds are used to redeem investor notes; this preserves the appropriate level of overcollateralization. If the market value stays above the trigger level, the notes remain outstanding. Market value structures also allow for tests of the equity/residual tranche; if a minimum level of equity is not maintained (i.e. an amount sufficient to protect the senior and subordinated tranches), the senior noteholders can elect to accumulate their payments.

While most CBOs and CLOs are structured using the cash flow method, certain CBOs based on distressed securities use the market value approach, suggesting that rebalancing within the collateral pool may be more active. The market value structure is more flexible; the portfolio manager may enjoy considerable discretion in buying/selling securities within the portfolio (e.g. up to 20 % of the portfolio per year), though she must still adhere to operating parameters related to diversification, overcollateralization, and credit quality. In fact, portfolio managers do not always have complete flexibility to trade in and out of positions in the portfolio. The goal of the market value structure is to maximize total returns while minimizing price volatility; sharp and sudden declines in asset value can be dangerous for such structures, even triggering note redemption. Repayment of liabilities under the market value structure can occur when assets are sold or when they mature. As such, they are not well suited for collateral pools that feature unpredictable cash flows, or those with assets maturing well after the final maturity of the CDO; both of these can expose the structure to significant losses if assets need to be liquidated early.

6.3.3 Balance sheet and arbitrage CDOs

Balance sheet CDOs were developed originally by sponsoring banks attempting to optimize their balance sheets, capital, and risk exposures by transferring portions of their loan (and eventually bond) portfolios to third party investors. Like other securitizations, balance sheet CDOs involve the creation of an SPE (for a single issue of securities) or a master trust (for multiple issues), which buys a pool of credit-risky assets from the transaction sponsor or credit originator. The asset purchase is funded through the issuance of multiple tranches of securities, which are placed with investors; the sponsoring institution may retain the equity. A typical balance sheet deal is likely to feature senior notes comprising 90–95 % of total deal size, subordinated notes (both strong and weak investment grade) of up to 8 %, and an equity tranche of approximately 2 %.

Arbitrage CDOs have grown more rapidly than balance sheet CDOs over the past few years, as sponsors and investors seek to capitalize on market opportunities.²⁷ If a bank identifies credit assets in the secondary market that it perceives to be cheap to theoretical value, it can purchase them, and then repackage them in an SPE or trust to lock in value. By doing so, it crystallizes the spread differential between the cash outflows (funding) and cash inflows (principal and interest), and is guaranteed an attractive return as long as default performance within the portfolio does not outpace default-related assumptions used in constructing the CDO. Arbitrage CDOs, which are often leveraged from 8–12 times, may include high-grade bonds, high-yield bonds, and emerging market bonds; sponsors increasingly use loans in their arbitrage structures as well; this has been made possible by growing secondary liquidity in certain types of high-grade bank loans (it does not apply uniformly, however, as the transfer of many loans must still occur on a case-by-case basis between the borrower and originator/seller, as noted above).

²⁷ While a cash flow arbitrage CDO functions in a manner similar to a balance sheet CDO (except for the goal of profit generation rather than risk transfer), a market value arbitrage CDO appears very similar to an actively traded investment or hedge fund.

The economics of the arbitrage CDO are driven ultimately by the returns payable on the equity tranche. The net flows on the structure can be separated into inflows and outflows: interest inflows include interest from collateral assets and any swaps that might be used to convert tranches into floating rates; interest outflows include interest payable to the senior and subordinated tranches and any swaps used in the structure; outflows also include fees payable to the trustee and portfolio manager. The net balance remaining is the amount payable (if any) to the equity investor. The amount payable varies based on defaults, asset callability, spread performance on assets that are sold during the period, swap hedges, and fees. If the arbitrage CDO has been structured correctly, the equity may generate returns of up to 30 %.

6.3.4 Structured and synthetic CDOs

Cash-funded CLOs and CBOs, which rely on true sale transfers, represent the core of the CDO market, and remain popular to the present time. The structured and synthetic CDO sector developed several years after the first funded CDOs were introduced as a logical extension of the asset class. Structured and synthetic CDOs rely on derivative replication transfer, rather than true sale transfer, to achieve the risk transfer and investment goals outlined above. This is, of course, similar to the concepts described in Chapter 4.

For the purposes of this discussion, we define a structured CDO as any CDO that relies on credit-linked notes (CLNs) for risk transfer, and a synthetic CDO as one that uses only credit derivatives to achieve the same results. As we noted in Chapter 5, a CLN is simply a combination of a bond and a credit derivative, and is thus one step “removed” from a pure synthetic derivative structure.²⁸ Structured and synthetic CDOs ultimately convey the risks/returns of the collateral pool without transferring legal ownership into an SPE or trust.

In the structured CDO, the sponsor, via an SPE, issues a CLN that references individual loans in the target portfolio; the CLN transfers the risk of the loans into the SPE, but not the physical assets, which remain in the sponsor’s possession. Investors in structured CDOs purchase tranches of rated securities that are backed by a pro rata share in CLNs (meaning they have no recourse back to the originating bank). Principal and interest from the CLNs are paid to investors; the CLN sources its principal and interest flows from the pool of reference loans or bonds. If an obligor from the reference portfolio defaults, the SPE bears the initial loss of principal and interest, which is transferred to the investors. In common with other CLNs, the defaulting note effectively is redeemed at a fixed recovery rate percentage, or at a dealer quoted post-default price.

Standard diversification techniques are employed in the structuring process, so that the portfolio of CLNs references a pool that has maximum exposures by obligor, industry, country, and rating.

The final maturity of any CLN reference asset must always be shorter than the maximum maturity on any CLN or funding tranche. Minimum levels of overcollateralization are typically used, just as in a conventional CDO; thus, a pool of \$105 m of loans or bonds might be needed to secure \$100 m of CLNs. Since there is no separation between the CLN reference obligations and any other obligations held by the bank in its overall portfolio, the structure depends on the bank’s performance; accordingly, the highest rating of the best tranche in the CDO

²⁸ A CDS embedded in a CLN that is used for a structured or synthetic CDO allows for payments to the protection buyer should a credit event occur; such credit events, defined under ISDA, generally include bankruptcy, restructuring, moratorium, repudiation, failure to pay, and/or acceleration of obligations. The ISDA restructuring definition was revised in 2003 to encompass instances where multiple bank lenders exist, and two-thirds of the lenders agree to the restructuring.

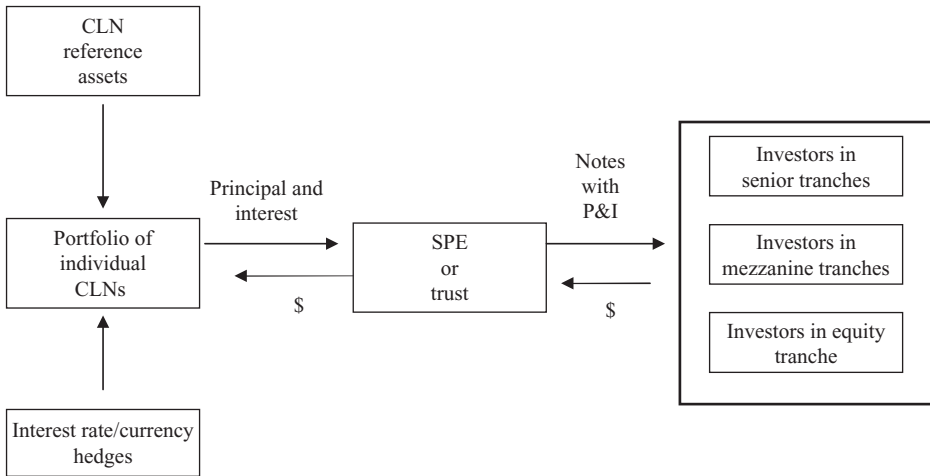


Figure 6.5 The structured CDO

cannot exceed the sponsoring bank's rating. Some structures incorporate early amortization triggers to protect investors in case reference asset loss experience accelerates; this results in an unwinding of the entire structure. In some instances, the structured CDO also allows for dynamic replacement (e.g. pool assets can be retired or new credits extended), indicating the CLNs are redeemable and reissuable on a quarterly or semi-annual basis. Each CLN represents a single pool obligation, facilitating the management process.

Structured CDOs avoid the security interest and transfer issues that characterize standard balance sheet CDOs, and simplify the hedging process by allowing currency and interest rate hedges to be embedded within each individual note. One notable disadvantage, however, is that the sponsor retains the reference assets on its balance sheet, which inflates total footings. Figure 6.5 highlights the basic structured CDO.

The synthetic CDO serves as an extension of the structured CDO. JP Morgan's inaugural transaction of 1997 (BISTRO), demonstrated the viability of using CDSs to transfer credit risk in the CDO structure, reducing the need for funding via the SPE/trust.²⁹ Under the standard synthetic CDO, which has since become widely used,³⁰ a sponsoring bank issues CLNs to investors to finance a pool of high-quality collateral assets, which are used to secure a CDS between the sponsoring bank and the issuing SPE or trust; the CDS references the specific credits in the sponsoring bank's portfolio requiring protection. The SPE receives a fee for providing the default protection; this fee, combined with the principal and interest cash flows

²⁹ Through the pioneering BISTRO deal, which transferred the credit default risk on \$9.7 b of loans, JP Morgan retained the residual (\$32 m), subordinated investors accepted the next layer of risk (\$237 m), senior investors the third layer of risk (\$470 m), and JP Morgan the remaining excess layer of risk.

³⁰ For instance, after the BISTRO deal, UBS (via Warburg) created Eisberg Finance, and CSFB created Triangle II, to transfer portions of their own credit portfolios. Citibank applied the same technologies to the European market by structuring the €4 b C-Star structure in 1999: the bank retained the 1 % first loss tranche, arranged a credit default swap with C-Star, which issued credit-linked notes to collateralize the transaction; Citibank then arranged uncollateralized credit default swaps with high-grade banks directly. The end result was a shift in the credit risk to a variety of investors and counterparties. It is worth noting that the same synthetic securitization technologies have been applied to repackaging of mortgage loan risk, letters of credit, and underwriting commitments. For instance, rather than using standard private label pass-through technologies, certain banks have used the synthetic CDO mechanism to shift the risk of portions of their mortgage loan portfolios. Each of these transactions must, of course, be viewed as pure risk transfer rather than funding.

from the high-quality collateral pool, is paid to investors funding the CLNs (again, via multiple tranches). The amount of cash funding is based on the extent to which a synthetic CDO uses credit derivatives; in the extreme, a 100 % unfunded synthetic CDO relies on complete use of credit derivatives to transfer risk to investors.

Synthetic CDOs, like conventional and structured transactions, feature senior/subordinated tranching that allocates losses in sequence; the cash flow stream is preserved while coverage tests remain in order. The sponsor often retains the first loss equity tranche in a synthetic CDO. In order for the sponsor to gain capital relief, it must limit the amount of default exposure it retains through the first loss tranche; generally this means that the portfolio is structured with high-grade, rather than high-yield, credit assets. It is also important to note that in order for a synthetic structure to work properly, the derivatives counterparty on the CDSs must have payment priority over investors in the CDO's tranches.

Synthetic CDOs can be created with static or dynamic reference pools; though static pools still account for more than two-thirds of the synthetic market, investor interest in dynamic pools has been on the rise (creating additional challenges for the rating agencies, as noted earlier). Investors in static transactions have full knowledge of the underlying pool, while those in dynamic pools are aware only of the general ratings parameters, rather than specific obligors, of the pool. Since static pools feature less "active management" than dynamic structures, they command lower fees. Synthetic arbitrage CDOs remain the fastest growing sector of the CDO market, far outpacing activity in traditional balance sheet structures; more than 70 % of new deal flow occurring since the millennium has been in the form of synthetic arbitrage transactions. The appeal lies in the fact that deals can be assembled much more rapidly, ramp up periods are shorter,³¹ and structures can be customized more precisely.

Partially funded synthetic CDOs, where the SPE issues only 5–15 % of the notional amount of the credit risk being transferred via the default swaps, have become quite popular as they are funding- and capital-efficient. Such deals function just as any fully funded cash, structured, or synthetic CDO, but the SPE or trust issues less notes and, thus, holds less high-grade collateral.³² The senior-most risk position of the deal, which exists above the senior-funded position, remains unfunded; this so-called "super senior" tranche is characterized by only a small probability that losses will exceed the funded component of the transaction. In some cases, the risk of the super senior tranche is transferred via a separate CDS. A typical partially funded CDO might be comprised of 80–90 % unfunded super senior tranches, 5–10 % funded AAA tranches, and 5–10 % mezzanine and equity tranches. The fact that the super senior tranche remains unfunded surfaces as a significant cost saving for the sponsor. For instance, if a bank funds at LIBOR + 50 bps, a super senior CDS costs 10 bps, and reinvestment in high-quality assets yields LIBOR flat, then the sponsor saves 40 bps running. In addition, the super senior tranche generally receives very favorable capital treatment (e.g. 20 %, rather than 100 %, risk weighting).

Synthetic CDOs, like structured CDOs, preserve assets on the sponsor's balance sheet, though credit risk clearly is transferred through the CDS. Though synthetics do not generate funding, neither do they require the full amount of securities to be placed with investors; for instance, a fully funded balance sheet CDO involving \$5 b of credit risk assets requires \$5 b of

³¹ A synthetic transaction can ramp-up and close in a matter of days, rather than the weeks or months required for a conventional CDO. Even CDOs that make use of warehouse facilities, where up to 75 % of assets are purchased prior to settlement and the residual investor bears the spread/default risk, can only shorten ramp-up periods by a few weeks.

³² It should be noted, however, that the SPE must retain enough liquidity to fund the physical delivery or cash settlement of any underlying assets that default.

Table 6.2 CDO characteristics

Feature	Cash CDO	Structured CDO	Synthetic CDO
Transfer mechanism	Assignment, sale, participation	CLN	Credit default swaps
Funding	Yes (via sales)	Yes (via CLN)	No/partial
Capital impact	Risk capital reduced by amount of transfer	Risk capital reduced if CLN acts as hedge	Risk capital reduced if swap acts as hedge
Balance sheet impact	Transfer assets off balance sheet	No transfer of assets off balance sheet	No transfer of assets off balance sheet

tranches to be placed in the market – a considerable distribution task. However, the synthetic CDO may only require placement of \$400 m of securities, which is sufficient for the sponsor to acquire the collateral needed to secure the CDS covering the \$5 b credit risk portfolio. Synthetic CDOs also give the portfolio manager greater flexibility; the asset portfolio may ultimately include a combination of physical assets and CDSs, while the liability portion may include securities and CDSs. The ability to short credits synthetically also exists, allowing certain CDOs to operate as de facto hedge funds.

Synthetic CDOs have proven appealing to investors, who gain exposure to high-grade credit risk portfolios, at an incremental yield pick-up over standard bond portfolios, and face none of the prepayment or extension risk characteristics of other ABS. That said, some price opacity exists, since synthetics are based on OTC derivatives that are impacted by liquidity, model, and hedging risks; the “true value” of a synthetic CDO trading on a secondary basis may be challenging to determine.

Table 6.2 summarizes certain key characteristics between a standard cash CDO, a structured CDO, and a synthetic CDO.

Variation 1: single-tranche CDOs

The single-tranche CDO, created in both physical and synthetic form (though the latter has become dominant), emerged at the turn of the millennium as an effective way of using CDO technology to customize individual tranches for investors. Rather than featuring a full capital structure (from senior to equity), the single-tranche CDO contains only a single portion of the loss distribution, and is tailored to an investor’s specific requirements. The investor specifies the underlying reference portfolio (which may include dozens to hundreds of reference credits), as well as the attachment point and the cap (which define the subordination and the size of the tranche, e.g. 3–5 %); the sponsoring bank retains the balance of the capital structure. The investor may also specify whether the portfolio is to be static or dynamic (though even when dynamic it is only “lightly” managed). Through the single-tranche structure, the bank sponsor essentially buys a tranche of credit protection on the bid side of the market on each credit in the portfolio, and can use the position to hedge portions of its own portfolio or monetize value by selling protection on the offer side of the market. If a reference defaults, the sponsor is protected, though the revenue it earns declines once the defaulted credit is removed.

For example, a bank might create, on behalf of its client, a \$30 m mezzanine tranche based on a \$1 b notional portfolio of 100 reference CDSs (\$10 m each). The tranche might attach at a loss level of 5 % and detach (or cap out) at 8 %, meaning that the investor bears losses of up

Table 6.3 Index tranches: Dow Jones CDX North America

Tranche	Default “loss width”
Equity	First 3 % of index default losses
Mezzanine	3–7 % of index default losses
Senior	7–10 % of index default losses
Senior	10–15 % of index default losses
Senior	15–30 % of index default losses

to \$30 m once the first \$50 m of losses has occurred. If we assume for simplicity that recovery on defaulted assets is 0, the investor will not begin to suffer losses until the sixth credit in the reference portfolio defaults. Thereafter, it will suffer losses until the cap level is reached. In exchange for accepting credit risk at the mezzanine level, the investor might earn a coupon of LIBOR + 200 bps.

Variation 2: CDS index tranches

Standardized tradable credit default swap indexes have been developed over the past few years. For instance, the Dow Jones indexes, comprised of up to 125 of the most actively traded CDSs in the market, are available for regions (North America, Europe, Japan, Emerging Markets), credit type (e.g. high-grade, high-yield), and industry sectors. The promulgation of standardized indexes such as these has boosted CDS liquidity considerably, and allowed the development of additional products, such as the CDS index tranche – a hybrid of the CDS index swap and the single-tranche CDO. A CDS index tranche is a derivative contract covering a specific segment (or “loss width”) of the CDS index default loss distribution. A CDS index tranche is similar to the single-tranche synthetic CDO, with two key differences: the CDS index tranche references a standard pool of credits that is always the same, while the CDO tranche references a customized, deal-specific pool; also, the index tranche features a standard “loss width” (e.g. the equity tranche is always the first 3 % of losses), while the CDO tranche “loss width” can vary by deal. These standard features have helped the index tranche product build a significant mass of liquidity, leading to attractive pricing.

In practice, the investor in the equity tranche (who is synthetically long the index of credits and is thus the protection seller) pays its counterparty an amount equal to the losses from the default (e.g. par – post-default price) up to a maximum of 3 % of the index. The mezzanine, senior, and super senior tranches likewise bear loss levels of their own, as noted in Tables 6.3 and 6.4, which depict the Dow Jones CDX North America and iTraxx Europe/Japan indexes.

Table 6.4 Index tranches: Dow Jones iTraxx Europe and iTraxx Japan

Tranche	Default “loss width”
Equity	First 3 % of index default losses
Mezzanine	3–6 % of index default losses
Senior	6–9 % of index default losses
Senior	9–12 % of index default losses
Senior	12–22 % of index default losses

In exchange for bearing the risk of loss, investors (as protection sellers) receive a quarterly premium payment from the protection buyers equal to the basis point premium times the outstanding notional of the tranche; the outstanding notional, in turn, is equal to the original notional less any losses impacting that tranche. Mezzanine and senior tranche investors receive only a quarterly premium spread. Equity tranche investors receive a quarterly premium spread plus an upfront payment equal to some percentage of the original notional amount of the contract; this additional payment is a reflection of the higher degree of risk the equity investors bear.

Variation 3: CDOs-squared

The CDO-squared, as the name suggests, is a CDO of a CDO, or a repackaging of other single-tranche CDOs (in fact, it can be likened to the CMO discussed in Chapter 4). A CDO-squared is characterized by a “parent” tranche and a series of “children” tranches; the latter reference pools of credits, some of which may contain some of the same references. Like other single-tranche CDOs, each child tranche has an attachment and detachment point (e.g. 5–10 %), which defines the loss level and magnitude. Similarly, the parent tranche features its own attachment/detachments (e.g. 10–20 %). Thus, a reference pool of \$1 b (100 credits, \$10 m each) might be used for eight child tranches of 5–10 % (\$50 m loss layer), suggesting the total notional is \$400 m. The parent, in turn, would feature a loss layer of \$40 m (10–20 % of \$400 m). In some instances, arrangers do not sell the child tranches to external investors; they may simply use them as components of internal deals.

Insurance-linked Securities and Contingent Capital

7.1 INTRODUCTION

The market for insurable risks has emerged as a new frontier in the structured asset marketplace. Though insurers and reinsurers have a long history of accepting and managing a range of insurance risks, they have been motivated to access additional risk capacity via the global capital markets. This has become increasingly important in an era when insurable risks are rising, and demand for risk transfer via the insurance mechanism is expanding. In order for the insurance community to tap into the global capital base, it must be able to convert insurance risks into securitized form; this ultimately allows creation of assets that can be bought and sold by traditional capital market investors in a convenient form.

In this chapter we consider several classes of the most prominent insurance-based capital market instruments, including insurance-linked securities and contingent capital issues.¹ The general class of insurance-linked securities includes catastrophe bonds (including hurricane bonds, earthquake bonds, windstorm bonds, and multiple-peril bonds) and noncatastrophe insurance bonds (weather bonds, residual value bonds, trade credit securitizations, and life insurance bonds); the class of contingent capital, in turn, can be separated into contingent debt (committed facilities, contingent surplus notes, and contingency loans) and contingent equity (loss equity puts and put-protected equity). Figure 7.1 summarizes these categories.

7.2 DEVELOPMENT AND MARKET DRIVERS

Financial intermediaries active in securitization of various types of assets and risk began applying the same techniques to the insurance market in the 1990s, creating notes and bonds based on insurance-related events. The concept behind insurance-based capital markets issues, or insurance-linked securities (ILS, we use the terms interchangeably), is similar to the securitizations we discussed in Chapters 4 and 6: namely, issuing securities referencing insurance risks in order to transfer exposures to investors, thus reducing/hedging risk or creating incremental risk capacity. Investors, for their part, receive the benefit of attractive returns that are generally uncorrelated with traditional financial asset classes.

The first ILS issue did not appear until the mid-1990s, though the groundwork started in 1992. Hurricane-based catastrophe (cat) bonds were first introduced to the market via USAA's pioneering issue of 1995, and have proven to be enduring, with steady to increasing annual issuance since that time. Following this successful "proof of concept," the same analytic technologies were applied to the development of earthquake bonds, windstorm bonds, and multiple-peril bonds. Other noncatastrophic insurance-related risk securitizations have appeared in recent

¹ Portions of the material in this chapter draw on Banks (2004) and Banks (2003); readers interested in a more detailed review of catastrophic risk and/or alternative risk transfer mechanisms may wish to consult these two texts, listed in the reference section of this book.

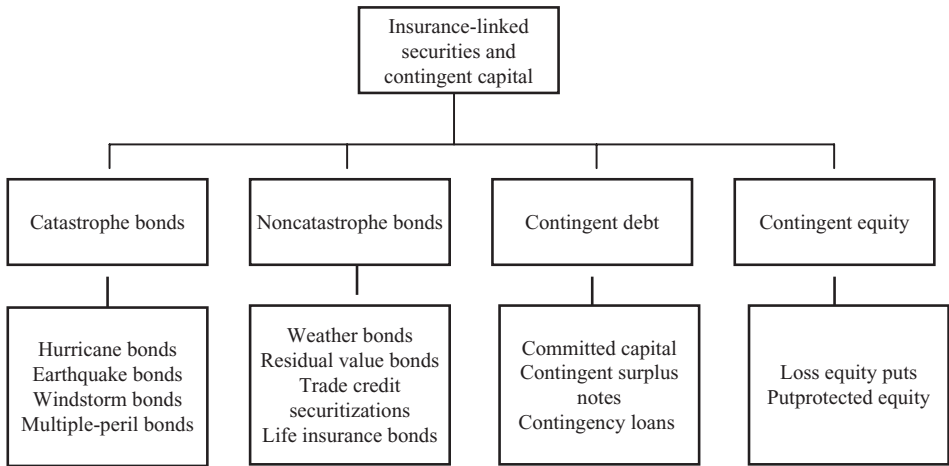


Figure 7.1 Insurance-linked securities and contingent capital

years, including those referencing temperature levels, residual value, life insurance policy acquisitions costs, auto insurance, workers’ compensation, and so forth; these sectors remain small compared to the catastrophe market.

Contingent capital issues, which provide for post-loss financing through structures arranged in advance of any loss, were also designed in the mid-1990s, following the reinsurance capital shortages that appeared in the aftermath of two major US catastrophes (Hurricane Andrew, 1992, and the Northridge earthquake, 1994); the catastrophe equity put, for instance, was launched formally by RLI in 1996. These structures were intended originally to provide insurers and reinsurers with alternative avenues by which to raise debt or equity funding in the aftermath of a disaster, and in instances where reinsurance prices were simply too high to be a cost-competitive risk tool. Though overall activity is still relatively modest, the likelihood of greater activity appears strong given the growing dollar-value of disasters.

One of the most appealing factors in the development of ILS and contingent capital structures has been the ability for institutions to link the insurance and capital markets – permitting the insurance sector to tap in to the tremendous supply of capital held by investors. Not surprisingly, most ILS and contingent capital issuers are insurers and reinsurers seeking alternate tools by which to manage their risk portfolios. Direct corporate issuance has been very small, with only a handful of issues appearing in recent years; in fact, most companies with catastrophic exposures find it simpler and more efficient to use standard insurance products to cover risks to hurricanes, earthquakes, and so on.

Securitization of insurance risks benefits various parties, including ceding companies, investors, and intermediaries. For instance, the ceding company (generally an insurer, as noted) can make use of another loss-financing mechanism to manage risk. During a hard reinsurance market, when supplies are tight and prices are high, this might be an attractive alternative.² It

² Since ILSs are a substitute (though not permanent replacement) for insurance/reinsurance, the price differential between reinsurance and capital markets issues has an influence on overall activity. When a hard market develops, ILS issuance can accelerate (though it remains within a relatively tight boundary, i.e. there is no evidence of a large spike in issuance). Since creating the ILS structure can be relatively expensive – based on costs associated with forming SPEs, preparing documentation, engaging investment banks to underwrite the issue, and so on – it is only justifiable in the cost/benefit framework when other loss-financing alternatives are more expensive. While an insurer/reinsurer’s decision to proceed with an ILS will depend on price, it must also take account of other issues, such as the amount of overall risk it wishes to retain and reinsure, the amount of credit exposure it wants outstanding to various reinsurers, and so forth.

also reduces its credit exposure to individual reinsurers; since the risk is repackaged into notes and sold to investors via the SPE, the ceding insurer no longer needs to be concerned about specific performance of the reinsurer. In addition, since the marketplace is highly bespoke, the insurer can design its preferred note structure: assuming greater basis risk but eliminating moral hazard, bearing the incremental cost of moral hazard but reducing basis risk, issuing single-year or multi-year cover, protecting against single or multiple perils, and so forth. Investors also gain, by purchasing securities that are likely to have little, or no, correlation with other risk assets in their portfolios. This is very appealing for investment managers, who are eager to find opportunities to earn extra yield without compounding the risk effects of the portfolio (e.g. a hurricane or earthquake event is not correlated with the movement of bond yields or the stock market, meaning diversification possibilities exist). Investors are also able to capture good returns. Most deals of the late 1990s and early part of the new millennium featured a “novelty premium” of 50–100 basis points in excess of what could be earned on similarly rated corporate bonds; though margins have compressed as investors have grown more familiar with potential risks, they remain attractive.

Activity in ILSs and contingent capital structures is driven by several key factors. Specifically, the instruments:

- allow insurers, reinsurers, and corporations to transfer exposures to catastrophic or non-catastrophic risks in an efficient and cost-effective manner;
- provide insurers, reinsurers, and corporations with an additional risk management tool, thus reducing their reliance on traditional insurance/reinsurance arrangements (which is especially important when the insurance markets are in tight supply);
- create additional risk capacity, allowing insurers and reinsurers to continue offering cedants coverage for their core insurance risks, without having to rely solely on the reinsurance markets;
- develop a formal linkage between the insurance risk market and the capital markets, allowing a transfer of capital as needed;
- permit investors to access a new asset class (i.e. catastrophic and noncatastrophic property and casualty risk), that is uncorrelated with other financial assets and offers excess returns.

Product innovation and expansion continue to accelerate as a result of these drivers. Though the overall ILS and contingent capital market is still very small compared to the other structured asset sectors we have discussed, its potential is extremely large – especially since insurable risks continue to escalate rapidly.

Investors in ILSs and other contingent capital structures are institutional parties that seek primarily to diversify their portfolios. In fact, investment funds, pension funds, and hedge funds that are willing to assume insurance risks are particularly active. Many are attracted by the spread premiums that can be earned from participating in rather novel, and quite illiquid, markets.

7.3 PRODUCT MECHANICS AND APPLICATIONS

Since the ILS and contingent capital markets merge insurance and financial technologies, their design characteristics are somewhat different from those we have considered in the ABS, MBS, and CDO markets. In this section we consider the specific mechanics of the asset class to demonstrate their unique construction.

7.3.1 Insurance-linked securities

ILSs are tradable instruments that transfer the financial risks of an insurance event from the ceding insurer or reinsurer to end investors. Though ILS structures have been refined and customized in recent years, the basic architecture has remained relatively unchanged: an insurance or reinsurance company issues securities through a special conduit, and bases repayment of interest and/or principal on losses arising from defined insurance events. If losses exceed a predetermined threshold, the insurer/reinsurer is no longer required to pay investors interest; if structured with a nonprincipal protected tranche, all, or a portion, of the principal can be deferred or eliminated as well. Through this elemental structure, new risk supply is created: the issuer passes a defined exposure to capital market investors, lowering its risk profile; this provides capital and reserve relief and allows new business to be written.

Legal and structural issues

The standard ILS structure is similar to other securitized capital market structures, except that a special purpose reinsurer (SPR), rather than an SPE or trust, acts as the issuance vehicle. A pure securitization of risk does not allow an insurer to meet its statutory capital surplus requirements; thus, some amount of risk must be reinsured to the SPR. This permits the risk to first be reinsured, and then securitized, which allows for the necessary capital relief. In a standard structure, the bankruptcy-remote SPR is responsible for writing a reinsurance contract to the cedant in exchange for premium. Since the protection provided to the cedant is in the form of a reinsurance contract, rather than a derivative, the SPR must be established as a licensed reinsurance company. Naturally, in order for the insurer to receive the benefit of ceded exposure, risk must be transferred, meaning the ceding insurer cannot own the SPR directly. In fact, charitable foundations sponsor most SPRs in order to fulfill this “independence” requirement. The SPR issues notes to investors, channels proceeds of the premium to the trustee for further investment, and arranges any swaps that might be necessary to fix coupon payments to investors. Proceeds from the issuance of the notes are used to purchase low-risk securities, which are held in a collateral pool; the return on the securities and the payment of premium from the ceding company are used to pay investor coupons. The collateral in the trust account is used to repay principal at maturity, unless an insurance event triggers a reduced payout; if this occurs, investors may not receive interest and/or principal on a timely basis, if at all (in some cases they will only receive recompense after all claims and contingent liabilities arising from the insurable event have been paid). By reducing, or eliminating, the payout to investors in the event an insurable event occurs, the ceding company mitigates its exposure to that event.

The general SPR structure is summarized in Figure 7.2.

In some cases, a second reinsurer is interposed between the ceding company and the SPR, meaning that the contract becomes one of retrocession (i.e. reinsurance of reinsurance), rather than reinsurance; this structure, depicted in Figure 7.3, permits the reinsurer to accept indemnity risk and hedge with index contracts, so that the ceding company does not have to bear basis risk.

Triggers

ILSs feature one or more triggers that determine the conditions under which the ceding company can suspend interest and/or principal payments (either temporarily or permanently), or

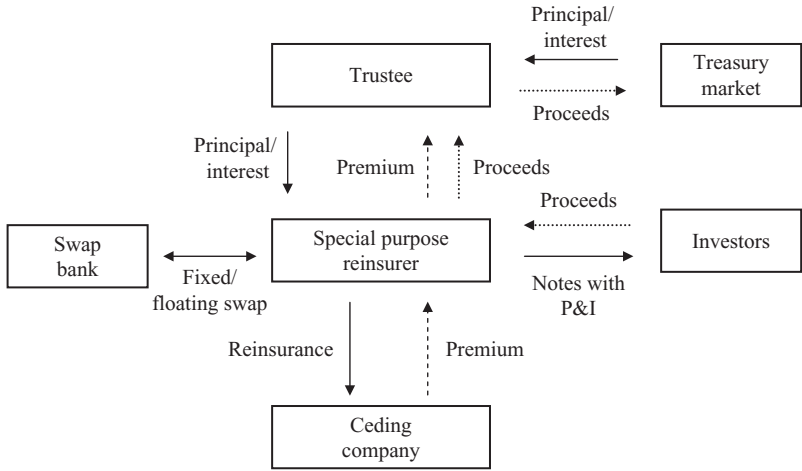


Figure 7.2 ILS with SPR as issuance vehicle

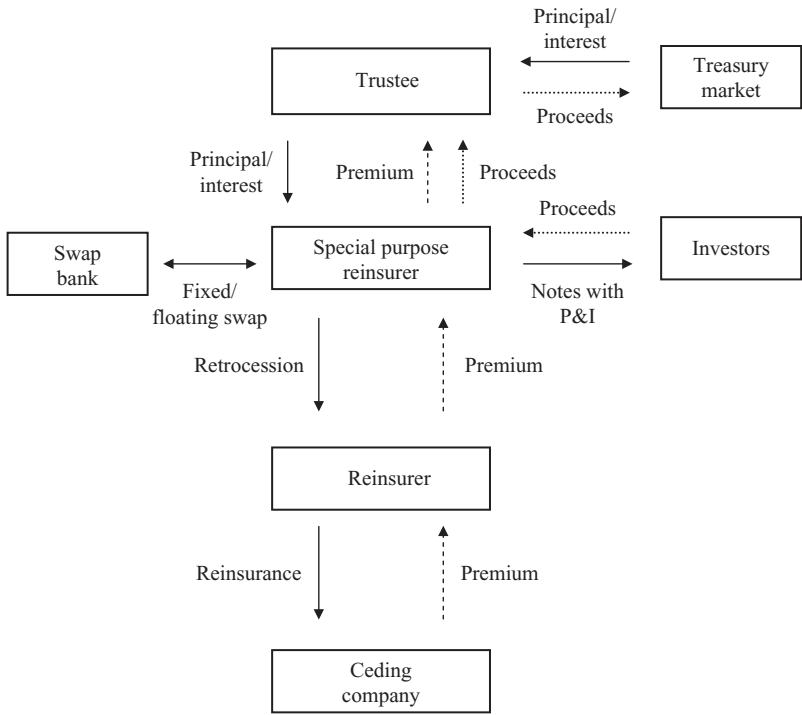


Figure 7.3 ILS with intermediate reinsurer

receive a capital inflow. A trigger can be based on single or multiple events (occurrences) and becomes effective after a cedant's losses exceed a particular amount (e.g. a de facto deductible). Referring back to our financial building blocks in Chapter 2, this is precisely equal to an option contract that moves in-the-money and becomes exercisable when a loss-based strike has been breached. This option, when packaged with an existing bond (or with a future financing commitment in the case of contingent capital), leads to the creation of a structured asset with insurance risk characteristics and dependencies.

Triggers can take one of the following forms:

- Indemnity trigger. The suspension of interest and/or principal occurs when actual losses sustained by the issuer in a predefined segment of business reach a certain level (e.g. an actual book of business).
- Index trigger. The suspension of interest and/or principal occurs when the value of a recognized third party index reaches a certain threshold.
- Parametric trigger. The suspension of interest and/or principal occurs when a specific loss metric reaches a certain value (e.g. a disaster of a particular magnitude/location).

Since indemnity bonds are based on a ceding insurer's actual book of business, they give rise to moral hazard risks: the cedant knows that the ILS trigger (and, therefore, restitution) is based on actual loss experience, and might not therefore be as diligent in assuming risks or enforcing loss control behaviors. However, since the loss experience is perfectly matched, basis risk is eliminated. In practice, indemnity deals require the cedant to reveal the nature of risk exposures and/or underwriting standards, and share in a portion of the losses; this helps obviate some of the moral hazard that might otherwise exist.

Index and parametric bonds remove the specter of moral hazard, since suspension of principal and interest is based on external events or values that are tabulated by third parties. The tradeoff, naturally, is an increase in basis risk for the issuer, as it is very unlikely that actual exposure matches the trigger; indeed, the ceding company must determine whether there is an index or parametric gauge that is correlated sufficiently to actual exposures to make the transaction viable. When a firm uses a parametric or index trigger, it faces lower costs (because it is assuming more basis risk), and is not required to divulge the actual details of the business being secured.³ Index and parametric securities may also be somewhat more liquid and tradable, as they are based on transparent metrics that all investors can evaluate (that said, the securities are still considered illiquid, certainly as compared with similarly rated corporate securities).⁴

Tranching

In common with other structured assets featuring subordination, ILSs are issued in multiple tranches that allow investors to select a desired level of risk and return. For instance, hedge funds may purchase mezzanine/first loss tranches, while investment funds and bank/insurance company investment accounts may purchase senior tranches.

³ In fact, investors are often indifferent as to whether the cedant has particular exposures – they can simply review the analytics and index construction rather than the cedant's portfolio.

⁴ Most early transactions in the market were based on indemnity and index triggers, and only a small number on parametric triggers. In recent years, the market has shifted from a majority of indemnity deals (e.g. 70 %+ in the late 1990s) to a majority of index transactions (e.g. 70 %+ in the new millennium by both dollar and volume). This is consistent with investor preferences; many investors favor index transactions because they add transparency and do not require a full evaluation of the cedant's underlying risk portfolio.

Table 7.1 Sample ILS tranches

Tranche class	Risk	Credit rating
1	Credit enhanced; no loss of interest payments or principal repayments	AAA
2	Loss of interest payments	AAA–A
3	Loss of interest payments	A–BBB
	Delay in principal repayments	
4	Loss of interest payments	A–BBB
	Partial loss of principal repayments	
5	Loss of interest payments	BB
	Loss of principal repayments	
Equity	Residual equity risk	Unrated
	Loss of interest and principal payments	

Tranches can be structured to reflect different levels of interest and/or principal delay or forfeiture; though every unsecured tranche is at risk (i.e. there are no principal and interest-protected securities unless specifically enhanced by a third party), potential loss can range from modest to extreme. As we have noted, some tranches might be credit enhanced by a highly rated guarantor in order to boost the credit rating and broaden distribution; since some investors cannot purchase sub-investment-grade securities, there are occasions when distribution requires the issuer to bear the cost of such an external “credit wrap.” Table 7.1 illustrates a sample of tranches that might be encountered on a typical ILS. Tranche 1, credit enhanced through a credit guarantee from an insurer or a letter of credit from a bank, might be rated AAA; Tranche 2, featuring possible loss of interest payments, may be rated from A to AAA, depending on structure; Tranche 5, with the potential for complete loss of principal and interest, is akin to a BB-rated security. As in other securitizations we have discussed, the rating agencies play an active role in evaluating the risks associated with individual deals and tranches (either directly or through shadow ratings), though they take different approaches: for instance, Moody’s rates issues and tranches on the basis of expected loss estimates, while Standard and Poor’s bases its assessment on the probability of first dollar loss.

Tranches that feature a delay mechanism might return principal as scheduled, and the balance over a period of time through a funded zero coupon position. Though individual tranches carry stated final maturities, actual maturity can be lengthened after an insurable event occurs, because claims may be slow to develop; thus, stated and actual maturity may differ. In practice, cedants like long loss-development periods as they permit the accumulation of a greater amount of claims, which can help reduce principal/interest repayments. Investors, not surprisingly, prefer shorter periods, so that they can receive and reinvest their principal/interest.⁵

Catastrophe bonds

As we have noted, catastrophe (cat) bonds comprise the bulk of the ILS market. Though growth in cat bonds has been relatively consistent over the past decade, the scope of coverage

⁵ Note that while many ILS transactions are multi-tranche and multi-year (given the desire to take advantage of the one-time expenses needed to establish the program) they are still governed by caps that can be breached before the final maturity. Thus, if a five, year Japanese earthquake bond has a \$250 m cap, and a \$300 m loss event occurs in year 2, the issue effectively is extinguished with three years remaining until final maturity (the loss-development period may cover six months to one year).

Table 7.2 Catastrophic risk classes for ILSs

Risk class	Territory
Earthquake	California
	New Madrid (US Midwest)
	Japan
	France and Monaco
	Taiwan
Hurricane/Typhoon	US Northeast/Atlantic
	US Gulf
	Puerto Rico
	Hawaii
	Taiwan
Windstorm	Japan
Hailstorm	Europe
Event interruption*	Europe
	Selective

* including terrorism-induced interruptions.

has expanded steadily – more types of peril and regions/territories are now referenced in specific deals, as noted in Table 7.2.

Hurricane bonds

Hurricane destruction can be devastating, so it is no surprise that firms seek appropriate coverage in the insurance market, and insurers seek cover of their own through reinsurance or hurricane bonds.⁶ Hurricane bonds were the pioneering issues of the cat bond market, and have proven to be enduring, with steady to increasing annual issuance since the mid-1990’s. Most hurricane peril covers are written on the US Northeast/Atlantic, Gulf Coast, and Hawaii, as well as typhoon equivalents in Japan and Taiwan. These areas feature important commercial and residential developments, many in the centers of hurricane trajectories, so adequate risk coverage is essential. Deals are structured with indemnity, parametric, and index triggers. Various analytic simulations suggest that a hurricane producing \$75–\$100 b of insurable losses could occur in the future; any such occurrence would impact reinsurance capital seriously and constrict risk capacity severely, making alternatives such as hurricane bonds an essential tool.⁷

Earthquake bonds

Earthquakes, like hurricanes, are another source of concern for companies and insurers attempting to manage property and casualty (P&C) risks, meaning the earthquake-based ILS has emerged as an important corporate risk management tool. Experience has shown that the financial damage wrought by earthquakes, and associated tsunami events, can be considerable (e.g. Northridge 1994, Kobe 1995, Southeast Asia 2004), and the possibility of even greater

⁶ The financial losses accumulated over the years from hurricane destruction have been considerable (e.g. the aggregate losses of Hurricanes Hugo (1989), Andrew (1992), Amber and Iniki (1992) were so large that they left 15 P&C insurers insolvent).

⁷ Hurricane Katrina, which struck New Orleans and other Gulf Coast areas in 2005, created extensive direct and indirect damage, estimated at well over \$100 b; actual insurable losses were approximately \$60 b.

losses exists. Some analytic projections estimate that an 8.5 Richter earthquake on the US New Madrid fault line could lead to direct and indirect losses of \$115 b (from P&C, business interruption, and so on). It is no surprise, therefore, that issuance of earthquake ILS has been reasonably active over the past few years. Deals are now structured routinely to cover earthquakes in California, Midwestern US, Taiwan, and Japan, with indemnity, index, and parametric triggers.

Let us consider the mechanics of a typical earthquake bond. A reinsurer might decide to shift a portion of its California earthquake exposure through a cat bond with an index trigger. Accordingly, it contracts with an investment bank to create and distribute a \$500 m bond with four distinct tranches. Under the terms of the bond, which has a two-year maturity (and a further one-year loss-development period, during which time additional claims can be made for any damages sustained), the tranches will pay reduced principal based on the Property Claims Service (PCS) index. Tranche 1, \$200 m of fixed-rate notes, and tranche 2, \$200 m of floating-rate notes, feature 50 % principal at risk. Tranche 3, \$70 m of sub-investment-grade fixed-rate notes, has 100 % principal at risk. Tranche 4, unrated, represents a \$30 m first loss piece that sustains 100 % loss if the PCS index losses for a California event exceed \$3 b. Through this structure, the reinsurance company transfers a portion of its California earthquake portfolio to the capital markets, based on loss events measured by the PCS index. In doing so, it retains basis risk equal to the difference between losses on its actual California portfolio and those that are estimated via the PCS index; it does, however, reduce the specter of moral hazard, and thus reduces its overall cost. To match its exposure more precisely, it could substitute the PCS index trigger with a parametric trigger based on locational and magnitude events (e.g. a 7.5 Richter Los Angeles county epicenter event); this could create a more accurate match, but would result in more expensive coverage. In the most “customized” version of the structure, it could embed an indemnity trigger into the bond, precisely matching any loss experience from a California earthquake with the exposure it has in its portfolio. As expected, this solution is the most expensive. The investors in the notes will, of course, receive a fixed or floating coupon that depends on the level of risk cover being provided. The senior tranche might earn a return of 100–200 bps over a benchmark, while the subordinated tranche might generate a return of 300+ bps.

Windstorm bonds

Windstorm risk is a third class of catastrophic exposure, and relates specifically to the P&C damage arising from very strong winds and rain. Windstorm-based ILSs have grown more popular in recent years and are now issued at fairly regular intervals, primarily on references in the UK, Continental Europe, and Florida.

Multiple-peril bonds

In some cases, ILSs are structured to handle multiple cat perils, such as losses from earthquakes and hurricanes occurring in different regions of the world. Such multiple-peril ILSs are intended to give the ceding company maximum flexibility and efficiency by eliminating the need to launch separate transactions for each named peril. Multiple-peril ILSs can be issued with indemnity, index, or parametric triggers, and with single or multiple tranches (related to overall interest/principal protection rather than peril exposure, as all named perils are covered in each tranche). The number of multiple cat peril ILSs appears to be on the rise: starting with the

earliest multiple-peril issues in 1999, issuance has grown in terms of size and creativity.⁸ Consider the example of a reinsurer that wishes to combine three perils into a single bond, receiving protection if any of the three events occurs. It structures a multiple-peril ILS that results in a reduced payout to investors if any one of the following three events are triggered: a minimum € 500 m loss from a French windstorm based on a parametric trigger; a minimum € 500 m loss from a Monaco earthquake based on a parametric trigger; or a minimum € 500 m loss from a UK windstorm based on a parametric trigger.

Multiple-peril ILSs are distinct from multiple-tranche bonds that cover distinct perils. An investor in the former buys a single security whose value can be affected by any one of several perils. An investor in the latter purchases a security that references a specific peril, with securities issued under an “umbrella” that allows for multiple-tranche issuance (i.e. peril by tranche); there is thus no commingling of risks, and investors need not deal with the valuation and risk complexities characteristic of a multiple-peril security. For instance, an insurer might issue a bond with three different tranches: tranche 1 referencing Japanese earthquakes with a parametric trigger, tranche 2 referencing New Madrid earthquakes with an index trigger, and tranche 3 referencing North Atlantic hurricanes with an indemnity trigger. The fact that the insurer issues all three tranches under a single program creates deal efficiencies and reduces costs.⁹

Bond/derivative variations

Variations on the standard cat ILS structure can be designed to meet specific issuer or investor goals. In some instances, insurers float bonds that generate a specific payoff related to an index; this can be compared to a bond with embedded derivatives, as discussed in Chapter 5. For instance, a bond may be structured to provide investors with an enhanced coupon if the index value of a named peril (e.g. a Japanese earthquake) remains below some threshold, or no coupon if it rises above it. This is simply equal to a bond with a strip of embedded short investor digital options that provide an enhanced coupon (or zero) at each evaluation date.¹⁰

The synthetic cat bond, which is simply an option on a cat bond, can be created to cap future reinsurance costs arising from a hardening of the reinsurance market. The optionable bond grants protection should the pricing cycle for standard reinsurance cover become unfavorable, but does not obligate the insurer/reinsurer to commit to a specific ILS transaction until (or if) the cycle turns. For instance, an insurer may purchase an option on a hurricane bond from another reinsurer or financial intermediary, giving it the right to issue a three-year hurricane bond with principal and interest payments tied to a parametric Florida hurricane once a particular loss threshold or reinsurance pricing index is triggered. If the initial loss threshold is triggered, the

⁸ For instance, Swiss Re launched the SR Wind bond, covering French windstorms and Florida and Puerto Rico hurricanes via two separate, but contingently linked, notes: if one peril attached, then the limit on the other could be transferred to cover losses from the peril already triggered. In the same year, French insurer AGF launched the Med Re bond, covering both European windstorms and French earthquakes. Securities were floated in dollars, but covered losses in euros for both events (with 65 % quota share reinsurance) and the first wind event (35 % quota share reinsurance).

⁹ For instance, in 2002, Swiss Re created the \$2 b Pioneer “catsec” program, allowing for issuance of specific tranches of securities covering P&C risks attributable to North Atlantic windstorms, European windstorms, California, and Midwestern US earthquakes and Japanese earthquakes. In 2003, Swiss Re issued the three-tranche Phoenix Quake bond for Zenkyoren (the Japanese National Mutual Insurance Federation of Agricultural Cooperatives). The \$470 m issue covered Zenkyoren’s exposure to earthquake and typhoon, and investors were able to select from among the three parametric bonds (Quake Ltd, Quake Wind, and Quake Wind II), each with its own peril, trigger, payout, and coupon.

¹⁰ As an example, Swiss insurer Winterthur has issued subordinated convertibles with hail-based catastrophe coupons. Investors in the deal received a coupon that was one-third greater than standard convertibles as long as the number of auto claims from hailstorm damage remained below 6 000, and no coupon if claims exceeded 6 000; the deal was thus a package of a subordinated bond, an equity option, and a strip of digital options referencing hail damage.

insurer exercises the option, placing the three-year bond with the intermediary (which might then distribute the securities to its own end-investors). If, during the three-year bond period, a Florida hurricane triggers the parametric index specified under the terms of the transaction, the insurer will pay the intermediary/investors a reduced principal/interest payout. Naturally, if the initial loss threshold is never breached, the insurer will not exercise the option to issue the bond. It is worth noting that the “forward” nature of the commitment entails a dimension of credit risk, i.e. being certain that the option seller (intermediary/investor) is financially willing and able to provide funds when (if) exercise occurs.

Noncatastrophe bonds

Following the success of cat-based securitizations, firms started applying the same techniques to securitize portions of their noncatastrophic insurance risk exposures. The market for these structured assets is still very nascent, but its prospects appear strong, particularly as the base of investors comes to understand the risk/return characteristics of individual categories of perils, and as companies are induced to produce a steady supply of bonds with enough uniformity to build a critical investment mass. Examples include bonds that transfer risks associated with noncatastrophic weather, residual value, trade credit, and life insurance mortality risk and acquisition costs, among others.

Weather bonds

Weather bonds, including those referencing temperature and precipitation levels in select cities or regions, represent one of the most promising sectors of the market. The OTC and exchange-traded weather derivative markets have proven quite successful since the turn of the millennium, and this has led to interest in securitized structures. The concept of the temperature-linked bond, in particular, has drawn a degree of interest. Such a structured asset, with principal/interest redemption that is tied to the level of cumulative temperatures in a particular city, group of cities, or region, has been mooted for several years, but overall issuance is still modest.¹¹ Activity is expected to accelerate, however, as the OTC/exchange markets continue to grow, and a greater number of companies start to manage their exposure to weather risk. A similar expansion is expected in precipitation-linked bonds (e.g. rainfall, snowfall).

Residual value bonds

Residual value bonds are ILs designed to protect firms against the residual value risks embedded in asset leases by shifting exposure to capital market investors. Companies with assets such as airplanes and automobile fleets often retain risk if they have provided a value guarantee. The risk contained in such deals centers on the difference between the original estimate of terminal value and the fair market value of the asset at the end of the lease (i.e. the former being greater than the latter). By transferring the future value (and/or credit) of the leased assets

¹¹ The inaugural issue in the sector was launched by trading firm Koch Industries in 1999. The company issued a \$54 m, three-year bond to create risk capacity by paying reduced principal in the event temperatures in 19 US cities breached predefined levels; if they remained within the trading band, investors earned an enhanced coupon. The Kelvin issue had various structural complexities, including two separate tranches governed by “events”: under terms of the transaction, the second tranche (event) was activated only if the first tranche (event) had been previously attached (though not necessarily exhausted); the second tranche would not attach, however, until the first tranche was exhausted (in addition, coverage was confined to predefined days).

to the financial markets via a structured bond, the issuer reduces the risk associated with the differential (i.e. the lessor de facto sells the risk of the lease receivables to investors). This provides protection against future residual value at the end of the lease term. Residual value ILSs can be arranged for leasing companies, financial institutions with lease portfolios, and original equipment manufacturers. Insurance companies writing residual value insurance (and financial risk insurance, which includes risk to residual value and the credit of the lessee) can securitize their risks through the same mechanism.¹² Though residual value ILSs are a niche sector of the market, they are gradually becoming more popular with firms facing lease value claims.

Trade credit securitizations

Trade credit securitizations, structures which provide for the transfer of trade credit insurance to the securities markets, are arranged periodically by insurance companies that guarantee third party trade credit facilities. Companies implicitly or explicitly extend trade credit to customers or buyers when they deliver services or products prior to receiving payment. The receivables that are generated through this activity are valuable, if risky, assets on the corporate balance sheet. To protect against risk of loss, companies can purchase trade credit insurance from an insurer covering any credit defaults by the trade debtors. Some insurers then securitize their pools of trade credit insurance through the ILS mechanism, issuing structured notes linked to the performance of a pool of insured trade credits, thus creating capacity to underwrite more credit insurance, or otherwise diversify and rebalance their portfolios.¹³ Trade credit securitizations, like other noncat ILSs, still feature rather modest issuance activity.

Life insurance mortality risk and acquisition cost securitizations

The life insurance sector is emerging as a promising area of securitization, from both a mortality and policy acquisition cost perspective. Prospects are particularly strong given the absolute size of the global life insurance market, which is substantial and growing; indeed, it must be regarded as the most significant area of business opportunity within the non-P&C sector.

¹² The residual value bond debuted in 1998 through an issue by Toyota Motors. Others, including BAE, Saab, and Rolls Royce, followed with similar structures of their own. Each one of these firms faces the risk that the residual value of assets (cars, jet engines, aircraft) that they have leased to customers will fall below market value. Leases generally give lessees the right to purchase the underlying asset at the conclusion of the lease period, which they will do if the market price is below the residual value. Accordingly, the companies are exposed as lessors to losses equal to the difference between the residual and resale values, and transfer the risk through the ILS market. For instance, Toyota, in a note issue arranged through the Grammercy Place Insurance SPV, launched a \$566 m three-tranche securitization covering 260 000 vehicle leases serviced by Toyota's financing arm, Toyota Motors Credit Corporation (TMCC). Under the terms of this residual value ILS, Grammercy provided three years of annual protection against residual value losses (with TMCC bearing a 10 % copayment and approximately 9 % deductible). Every year TMCC submitted its residual value claims (e.g. losses above the deductible, less the coinsurance) to Grammercy. Initial investor proceeds held by the trustee in a collateral account were used to repay TMCC's claims, and any remaining balance was then used to repay investors. If TMCC had no residual value claims in a given year, investors received full principal and enhanced coupons; if, however, there was a shortfall, they absorbed a fractional portion of the loss (not unlike a CDO investor; in fact, the only difference between a CDO and a residual value ILS is that the trigger event shifts from a pure credit claim to an asset value claim). Saab's transaction, conceptually similar to Toyota's, provided for \$1.17 b of 15-year lease risk protection.

¹³ For instance, in 1999, German insurer Gerling issued several tranches of SECTRS (Synthetic European Credit Tracking Securities) to transfer risks in pools of European corporate trade credits that it had insured; the portfolio included 92 000 randomly selected businesses. Under the terms of the transaction, Gerling launched €450 m of ILS in three tranches through the SECTRS 1999 SPR (with Goldman Sachs as placement agent). Namur Re, interposed between Gerling and SECTRS, provided Gerling with reinsurance cover on an excess of loss basis (in three different portfolios, i.e. one for each tranche). The trigger under the issue was based on annual and cumulative default counts in each of the three portfolios. Retrocession was activated when the annual count exceeded an annual attachment point, or the cumulative count (starting in year 2) exceeded the cumulative attachment point. When triggered, the payment due from SECTRS to Namur Re was obtained by multiplying the excess over the attachment by a recovery rate. Principal repayments to investors were reduced according to claims from Namur Re (which were based on claims from Gerling). Through this ILS structure, Gerling obtained extra capacity to write more credit risk business.

Under a life insurance mortality securitization, an insurer or reinsurer transfers to the capital markets a portion of the exposures it has written under life insurance or reinsurance policies. Though such transactions are relatively new,¹⁴ and there is no single established structure or framework, the essence of such a bond is to transfer excess mortality risk through the use of a trigger that references mortality indexes in a variety of countries where the insurer writes policies. In order to be effective, the structure cannot exclude very low probability events (e.g. pandemics, terrorist strikes with devastating – chemical/biological weapons). Though the securitization is in the earliest product evolution stage, it is likely to become an interesting product for insurers (i.e. another mechanism to shift risk) and investors (i.e. another vehicle that can provide returns that generally are uncorrelated with other asset classes).

Life acquisition cost securitizations, which are transactions that permit insurers to transfer the costs associated with writing life insurance policies, were introduced in the mid-1990s. The intent behind these unique issues is to transfer a portion of the costs associated with originating and servicing life insurance business, which often includes front-loaded expenses for the insurer (e.g. broker fees, sales distribution fees, and so on). In some instances, securitization is driven by specific regulatory requirements; for example, German regulations do not permit acquisition costs to appear as assets on the balance sheet, creating pressures on the financial position of any large underwriter of life policies. Through the life acquisition cost bond, the insurer grants another insurer (or parent company, joint venture partner, or financial institution) the right to receive future profits from a particular pool of life policies in exchange for the present value of future cash inflows, which can be used to cover upfront costs (and reduce the impact on the income statement and balance sheet). In fact, these securitizations can be viewed as versions of risk financing, rather than pure risk transfer (though much ultimately depends on the specific structure).¹⁵

Other ILS variations have been developed, including those related to mortgage insurance securitization and event cancellation risk securitizations (e.g. a terrorism-related risk¹⁶). The ILS market is still very much a niche within the structured asset sector, certainly when compared to MBS, ABS, CDOs, and structured notes on “mainstream” asset classes. Issuance in some sectors of the ILS market has been steady, a dedicated investor base has developed and spreads have tightened (making them an even more compelling alternative to traditional reinsurance mechanisms during some market cycles). Activity in other segments is still relatively quiet and will take time to develop. Not surprisingly, secondary liquidity in all ILS sectors is still extremely thin; the securities must be considered “buy and hold” investments.

7.3.2 Contingent capital structures

Contingent capital represents the second major class of structured insurance products. These instruments are contractually agreed financing facilities that are made available to a company

¹⁴ Swiss Re launched the pioneering VitaRe bond, \$250 m, 3½ years, referencing mortality indexes in the US, UK, France, Switzerland, and Italy. Analytic estimates suggested that in order for the bond to attach (e.g. for investors to start bearing losses), two major devastations would have to occur. The bond did not exclude pandemics or terrorist attacks.

¹⁵ For instance, in 1996 and 1997, American Skandia Life sold its parent company the rights to future mortality and expense charges for a present value payment of expected future claims. The parent company securitized future fees via an SPV collateralized by the receivables. National Provident, a UK insurer, securitized future profits on its life policies through the Mutual Securitization SPV, which issued two tranches of bonds (featuring final maturities in 2012 and 2022), with principal and interest tied to the surplus on the insurer's life policies. This permitted the insurer to crystallize, on a present value basis, the surplus embedded in life policies that would otherwise only be realizable over a long period of time. Others, including Hanover Re, have structured similar deals.

¹⁶ For instance, FIFA's Golden Goal securitization, provided coverage for any natural or manmade (e.g. terrorism-related) event causing a cancellation of the World Cup football tournament.

in the aftermath of a loss event. As with other capital market products, the contingent capital structure helps link the insurance and financial markets by raising funds from capital market providers/investors upon the trigger of an insurance-related event. Unlike ILSs, which contain aspects of insurance/reinsurance and securities financing, contingent capital facilities are arranged purely as funding facilities or securities transactions, with no element of insurance contracting.

Though the contingent capital facility is not yet as prevalent as the ILS, companies developing broad risk management programs must consider its use as an element of post-loss funding. In this section we consider the most common structures, including contingent debt and contingent equity. Within these broad classes, we subdivide contingent debt into committed capital facilities, contingent surplus notes, and contingency loans, and contingent equity into loss equity puts and put-protected equity.

Legal and structural issues

Contingent capital allows a firm to raise capital during a defined commitment period if a specific loss-making event occurs. Since such facilities are arranged in advance of any loss, their cost does not reflect the risk premium that may become apparent in the aftermath of distress (i.e. lower creditworthiness and less access to liquidity, leading to a higher cost of capital). This makes the facilities cost-efficient across a range of financial scenarios. A firm that attempts to arrange funding after a disaster has weakened its financial condition will pay a higher cost of funds; this is especially true if its credit condition has been lowered to sub-investment-grade levels. A firm that has been impacted by the same disaster, but arranged its capital access ex-ante, will be indemnified and recapitalized at the cost of capital agreed pre-loss.

Through a generic contingent capital structure (illustrated in Figure 7.4 as a securities issue, though it can be structured easily as a banking facility), a company identifies an amount of capital that it wishes to raise in the event it suffers a loss, determines the events that can trigger

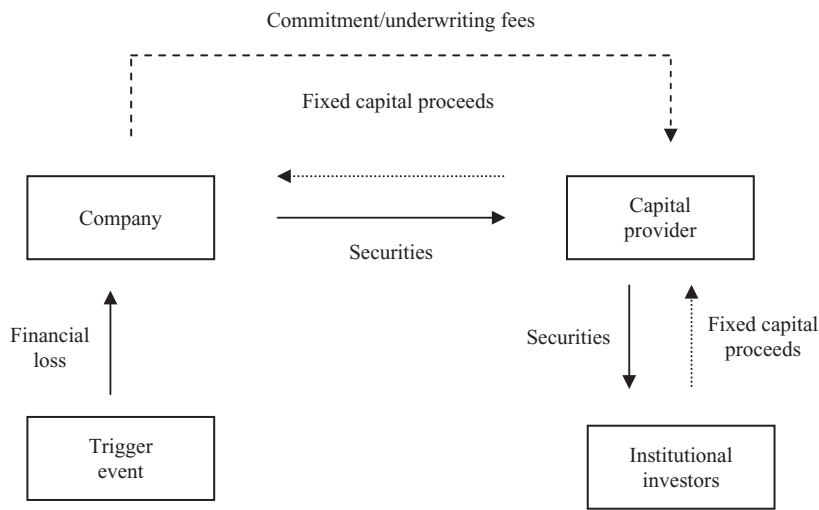


Figure 7.4 Generic contingent capital structure

the loss, and the specific form of securities it will issue in order to raise capital. If the event occurs, the capital provider supplies funds by taking up securities issued by the company at the ex-ante price. In return, the company pays the capital provider a periodic (or upfront), nonrefundable commitment fee (payable whether or not securities are ever issued), as well as an underwriting fee (payable only if securities are floated). Though the legal commitment to provide funds rests with the capital provider, in practice it will almost certainly distribute the securities to a base of institutional investors. However, if the provider cannot place the securities with investors, it must still supply the company with funds. Therefore, the capital raising effort, in underwriting parlance, is considered a “firm commitment” or “bought deal” (contingent on the triggering event), rather than a “best efforts” or “agented” transaction. Reverting to our building blocks from Chapter 2, we can view the generic structure as an option on a future financing held by the company; the strike price and notional size equate to the issue price and proceeds that will be raised in the event of exercise. If exercise occurs, the company invokes its right to sell the capital provider securities in exchange for capital proceeds. The exercise of the option is, of course, dependent on the occurrence of the trigger event; it cannot be exercised at will or at maturity, as would be common under an American option or a European option. The commitment fee payable by the company can be viewed as the premium any option buyer would pay a seller. The structure can also be viewed as a package of a knock-in option and a bond, where the breach of the knock-in barrier condition leads to the creation of a new bond.

Contingent capital products incorporate triggers (barriers) that are activated by a stated loss level. Like ILSs, the triggers can be created on a customized basis in order to match a company's exposure to a specific loss-making event, or they can be based on market indexes. The terms of the resulting securities, negotiated in advance between the company and the capital provider, can vary widely. Securities can be issued as common equity, debt, or preferreds. If issued as equity, dilution issues must be considered, and if issued as debt or preferreds, specific details related to leverage, subordination, maturity, coupon (or dividend), callability, and dividend treatment must be addressed. Similar issues must be addressed if debt funding occurs via bank lines rather than securities, along with specific details on drawdown features, material adverse change clauses and covenants, and so forth.

Post-loss financing products, such as contingent capital, can be used in conjunction with traditional insurance or financial hedges. Since contingent capital is focused primarily on low-frequency disaster events, rather than high-frequency/low-severity insurance events, they are meant to supplement, rather than replace, other forms of risk transfer and financing (e.g. a firm would use an insurance policy to cover more probable risks, and a contingent capital facility to cover less probable risks). Contingent capital products also have the advantage of giving a company the ability to manage risks that might not be possible through other traded instruments (e.g. losses arising from a particular catastrophe that does not lend itself to reference monitoring via a standard contract, losses emanating from certain forms of credit risk, and so forth). For instance, a bank might arrange a facility that is triggered by unexpectedly large credit losses; the infusion of capital arising from the breach of the trigger can be used to replenish capital and reserves. An insurer or reinsurer might use contingent capital to provide additional funding in the event of a large catastrophic loss; this can serve as a complement to any other ILS or excess of loss reinsurance coverage the insurer/reinsurer might have. The structure can also be applied to broader events. For instance, if a company is highly sensitive to economic growth rates – perhaps it fears it will be downgraded to sub-investment grade if a recession strikes and causes production and sales to decline sharply – it might arrange for a contingent capital facility that will allow it to borrow at rates determined today, before any recession hits. If economic growth

slows and weakens revenues, causing the downgrade, the firm will not be exposed to higher borrowing costs associated with its weaker credit status; the trigger in this example is based on a macroeconomic indicator, such as gross domestic product.

It is important to remember that contingent capital is not insurance, but a balance sheet and cash flow arrangement (that actually shares structural similarities with various finite risk programs), and does not therefore provide earnings protection or feature the same tax deductibility characteristics of insurance policies. Furthermore, a company arranging an issue of contingent financing relies on the provider of capital to supply funds when called on to do so. The company thus assumes the capital provider's credit risk on a contingent basis. We can imagine an extreme scenario where a company and a bank (as capital provider) agree to a \$500 m capital infusion in the event the company suffers a severe loss. If the trigger is breached and the company loses more than \$500 m, it expects the compensatory equity infusion from the bank. However, if the bank fails to perform (e.g. perhaps it has encountered financial distress of its own, or has actually defaulted on its obligations), the company is left without the vital capital injection it expects, which may create financial distress or insolvency. Credit risk issues are thus a central issue in any contingent capital structure. Though most transactions are arranged and funded by top-rated counterparties, some may be arranged by medium-rated entities (or higher-rated entities that deteriorate over time).

Contingent debt

We consider several structures within the general category of contingent debt, including committed capital facilities, contingent surplus notes, and contingency loans.¹⁷ Though each features slightly unique characteristics, all have the same end goal of providing the company with prenegotiated post-loss debt financing.

Committed capital facilities

The committed capital facility (CCF) – funded capital arranged prior to a loss and accessed when two trigger events are breached – is one of the most common forms of contingent capital. Under a typical CCF, a company creates a financing program, defining the specific bond it intends to issue upon triggering, e.g. seniority/subordination, maturity, repayment schedule, and coupon. The insurer/reinsurer arranging the facility acts as the capital supplier, providing funds in the event of exercise, which only occurs if the triggers are breached. The first “option”

¹⁷ An associated structure, the financial guarantee, serves a similar purpose but does not fit our classification of a structured product, so we choose not to include it as part of our primary discussion. In fact, financial guarantees, which have existed for many decades, are most commonly used to transfer risk; however, by virtue of their construction, they also represent a form of contingent financing. In its most basic form, a company and a financial guarantor (e.g. an insurer or reinsurer) agree to a loss trigger that, if breached, allows the company to access funds from the guarantor. Guarantees of this type are used commonly to protect companies and SPEs against credit losses or residual value claims. They are also used by exchanges and clearinghouses, who want to ensure sufficient capital in the event of a low-probability, loss-making event (e.g. a massive counterparty default). Bond insurance issued by monoline insurers for the benefit of SPEs that issue CDOs or ABS, essentially is a financial guarantee. In exchange for a fee, the insurer “wraps” the SPE's CDO with a guarantee, so that it can achieve a higher credit rating (and thus be salable to a broader range of investors). If credit loss experience is greater than expected – thus breaching the trigger embedded in the guarantee – the SPE receives a capital infusion from the insurer, which it passes on to investors holding the guaranteed tranches (e.g. the original super AAA or AAA tranches; the holders of subordinated tranches or residual equity will not receive the benefit of the capital infusion). Residual value guarantees perform a similar function by giving a company minimum protection against the residual value inherent in leased assets (e.g. airplanes or aircraft engines). If a company finds that the residual value at the end of a lease is lower than originally anticipated, it might suffer a capital shortfall and become vulnerable to financial distress. The residual value guarantee ensures some minimum value for the assets, meaning the company receives a capital inflow if a shortfall occurs. This again represents a contingency, as if the residual value remains at, or above, the scheduled value, no infusion is required.

trigger is often implicit – that is, the option will not be exercised unless it has value, and it will only have value if a loss occurs and the company cannot obtain cheaper funding from other sources. The second trigger generally is related to the exposure that the company is seeking to fund (in order to minimize basis risk), but the specific trigger event is unlikely to be under the company's control (in order to eliminate moral hazard).

As with other contingent capital structures, the CCF generally has a fixed maturity date, and is intended as a form of contingent financing rather than pure risk transfer. The CCF may contain covenants that can be used to protect one or both parties. These may include material adverse change clauses, change of control, financial strength/ratios, and so forth. The intent is to ensure that if the facility is triggered, the insurer/reinsurer providing funding does not become subordinate to other lenders. In more complex structures, the insurer/reinsurer writing the contingent option might join with a bank (or syndicate of banks or securities firms) to provide funding; this eases the financing burden on the insurer/reinsurer and places the funding in the financial sector. The bank syndicate might then choose to hold the funding instruments or sell them to institutional investors.

Consider the following example: a bank seeks to protect its reserve levels from low-probability/high-severity loan losses and arranges a CCF with an insurer. Under the facility, the bank can trigger additional funding through the issuance of up to \$750 m of fixed-rate notes if its loan portfolio suffers exceptional credit losses. This mechanism allows the bank to replenish its reserves when needed, but also to manage its balance sheet more efficiently by not having to hold greater reserves than necessary. Since the bank's credit portfolio is sufficiently well diversified, an external loan index comprised of cross-industry and country credits is selected as the reference trigger; this helps eliminate the possibility of moral hazard.

Contingent surplus notes

Insurance and reinsurance companies seeking protection against exceptional losses in their portfolios are issuers of contingent surplus notes (CSNs).¹⁸ Under a typical CSN structure, an insurer contracts with a financial intermediary to establish a trust, which is capitalized by outside investors through trust-issued notes paying an enhanced coupon. The trust invests proceeds in high-grade investments (e.g. AAA-rated bonds) until (or if) the contingent capital is required. If the insurer breaches a predefined loss trigger (e.g. the option moves in-the-money), it issues CSNs to the trust. The trust liquidates the AAA-rated bonds and delivers cash to the insurer. In exchange for providing the initial commitment, and contingent capital, the investor achieves an all-in yield that is greater than that found on similarly rated corporate securities. The insurer obtains a post-loss funding commitment in advance at a price that might prove compelling in an expensive reinsurance market. The commitment fee that the insurer pays to the trust can be viewed in the same light as the option premium in the CCF. Figure 7.5 summarizes a generic CSN structure.

Consider the following example: an insurance company arranges a \$500 m, five-year CSN issue that will be triggered in the event losses in its P&C portfolio exceed \$500 m over the next two years. The arranging bank identifies several institutional investors that “prefund” a trust

¹⁸ The first CSN, a \$400 m ten-year note, was issued by National Mutual in 1995. Under the terms of the transaction, National Mutual was able to access up to \$400 m in fresh capital through a trust, which issued surplus notes to investors as needed. In a variation on the standard structure, the insurer was able to raise funds whenever it wanted, and not just in the event of a loss. The risk to investors (for which they received an above-market coupon) was the possibility that the state insurance commissioner would instruct National Mutual to cease paying principal and/or interest if policyholders were deemed to be disadvantaged or prejudiced by such payments.

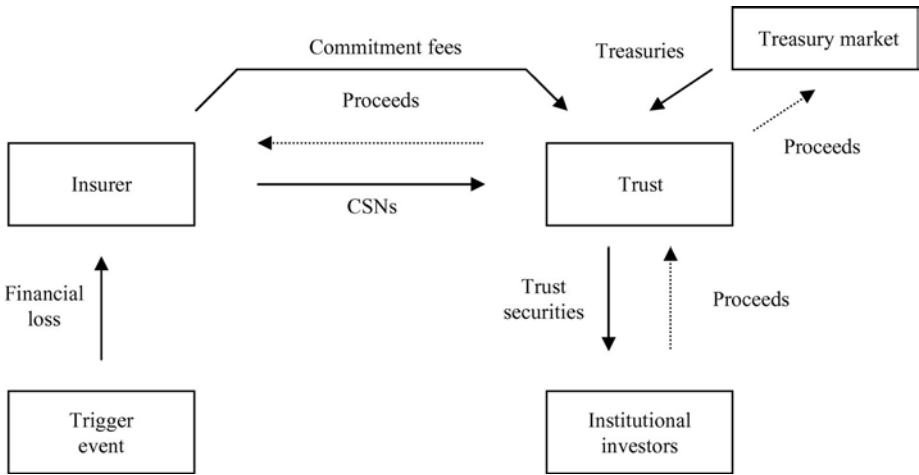


Figure 7.5 Contingent surplus note structure

for the full \$500 m in exchange for an all-in yield equal to the commitment fee plus the return on five-year US Treasury notes. The \$500 m in prefunding proceeds are used to purchase US Treasuries; the investors receive notes in the trust, reflecting an enhanced coupon. Assume that one year from now, loss experience is greater than expected, causing the trigger to be breached. The insurance company issues \$500 m of five-year notes to the trust; the trust liquidates its Treasury position and uses the proceeds to acquire the CSNs. The trust thus holds the insurance company’s CSNs, the insurance company receives \$500 m of cash to help manage its financial position, and end-investors continue to receive the enhanced yield on their trust-issued notes.¹⁹

Contingency loans

The contingency loan, a variation of the CCF, is a bank line of credit that is arranged in advance of a loss and invoked when a trigger event occurs. From a structural perspective, this can be viewed as an option on a bank line. Unlike a traditional line of credit, which can be used for any purpose and accessed at will, the contingency loan is only available to cover losses arising from a defined event, and can only be drawn when that event occurs. Since the company has far less flexibility in drawdown, it pays less than it would for a standard bank line. In addition, because the probability of drawdown is lower, a company may be able to negotiate a larger borrowing amount, giving it greater comfort that financial flexibility will not be impaired in the event a loss is triggered. As with the CCF above, key terms of the contingency loan are defined in advance, including maximum amount, fixed or floating rate, maturity, repayment schedule, and option trigger(s). The bank called on to fund the facility upon exercise can retain the loan commitment or syndicate it to other institutions.

¹⁹ Insurers may also issue surplus notes, which are subordinated securities that function much like the CSN, except that they are issued directly by the company rather than through a trust. Surplus notes typically have maturities of 10 to 30 years, and must be approved by insurance regulators. Importantly, these notes increase statutory, though not accounting, capital.

Consider the following example: an auto manufacturer that produces and sells primarily in the Americas wants to arrange contingent financing that it can access in the event of a slowdown in economic growth. Through intensive factor analysis, the company computes its revenue sensitivity to each basis point decline in economic growth, and is aware that the resulting loss of revenues is likely to lead to financial pressures, a credit rating downgrade, and a rise in its cost of funding. The company then develops (with its bank) a contingency loan that allows it to draw-down up to \$750 m of a multi-year loan if growth rates in North America and Latin America fall below certain average levels. Since the growth rate measure serving as the trigger references external indexes, there is no possibility of moral hazard. The contingency loan is, of course, structured with an appropriate rate, maturity, and repayment schedule.

Contingent equity

In some instances a company prefers, or requires, contingent funding in the form of either common or preferred equity. This helps ensure that the post-loss recapitalization effort does not increase the debt burden and negatively impact leverage ratios; since the infusion comes in the form of equity, leverage is preserved or lowered. However, any contingent equity structure that involves the issuance of new common shares results in earnings dilution; in addition, since equity capital generally is more expensive than debt capital, the economics of the post-loss capital-raising exercise might not be as compelling. We consider two different forms of contingent equity – the loss equity put and put protected equity.

Loss equity put

The loss equity put (LEP) (sometimes known as the catastrophe equity put²⁰ when related specifically to a natural disaster trigger), is a contingent capital structure that results in the issuance of new shares in the event a predefined trigger is breached. This, again, can be viewed as an equity issue with an embedded option. In fact, the structure and mechanics of the LEP are similar to the committed facilities noted above, except that equity, rather than debt, is issued if a trigger is breached. In a typical structure, a company purchases a put option from an intermediary that gives it the right to sell a fixed amount of shares (often on a private placement basis) if a particular loss trigger is hit during the life of the contract. In exchange, the company pays the intermediary an option premium. Since the terms of the put option are fixed (e.g. number of shares to be issued and strike price), the post-loss financing is arranged and committed in advance of any loss. If the option becomes exercisable, the company issues new shares to the intermediary, pays any additional underwriting fees, and receives agreed proceeds. In order to avoid dilution issues that arise from the issuance of new common shares, LEPs may result in the issuance of preferred, rather than common, equity. They may also be issued as convertible preferred shares (such as we discuss in Chapter 9), with an implicit understanding between the two parties that the preferreds will be repurchased by the company at a future time, prior to any conversion (thus avoiding dilution).

Each LEP transaction is characterized by standard terms and conditions, including exercise event, form of securities, minimum amount of securities to be issued on exercise, the time period

²⁰ Insurance broker AON, which developed the original structure in 1996, has coined (and registered) the product name CatEPut®, which seems to have entered the financial vernacular.

of coverage, the maximum time allotted for issuance of securities, strike price, and specific warranties (e.g. minimum net worth (or statutory capital) on exercise, change in control, minimum financial ratios, and so forth). To reduce the possibility of moral hazard arising from an indemnity structure, LEPs often have two triggers. The first trigger is the company's stock price, which must fall below the strike price in order to become effective; this is consistent with the option framework we described in Chapter 2. The second trigger relates to specific loss levels that must be breached in order for exercise to occur. Thus, it is not sufficient for a company's stock price to decline – it must be accompanied by a loss event. In fact, the two triggers are likely to be related if losses are large enough; that is, the company's stock price is more likely to fall through the strike price if the marketplace becomes aware that the firm has sustained large losses. LEPs can be structured with an index or parametric trigger instead of an indemnity trigger.

In addition to the obvious benefit of “locking in” post-loss funding at predetermined levels, LEPs feature at least two other advantages. First, unlike debt facilities that often contain material adverse change clauses that limit or prohibit funding in the event of disruption (either in the market or with the ceding company), LEPs have no such limitations (except for maintenance of minimum net worth), meaning they are certain to be available when needed. Second, the cost of an LEP can compare quite favorably to a standard reinsurance contract, because the option purchaser must remain financially viable in order to claim access to funding (the same is not true in a standard reinsurance contract, where the cedant can become insolvent and the reinsurer must still perform on its reinsurance obligation). Specifically, to ensure that the option writer is not forced to invest in a financially distressed company, transactions generally include minimum net worth covenants; if the option buyer's net worth falls below a predetermined threshold, it cannot exercise the option and raise new proceeds. For instance, if a cedant has net equity of \$100 m and a loss equity put that will give it access to a further \$50 m upon exercise, a catastrophic loss leading to net equity of only \$10 m would place the put seller in a dire situation, forcing a \$50 m investment in a company that only has a current net worth of \$10 m. The mechanism avoids this situation.

Figure 7.6 illustrates the flows of an LEP pre- and post-trigger. Note that, as in the contingent debt structures above, the intermediary ultimately is responsible for taking up new shares and delivering proceeds if the LEP is exercised. In practice, it would turn to its base of institutional investors to distribute the shares.

Consider the following example: insurer ABC, whose stock is currently trading at \$32/share, is concerned about risk concentrations in its catastrophe portfolio, and wants to be protected in the event total losses over the next underwriting season exceed \$500 m. If this occurs, ABC will need to issue new common stock in order to replenish its capital base and remain comfortably within regulatory capital requirements (its current equity level is \$1.5 b). Accordingly, it purchases a \$500 m LEP from Reinsurer XYZ, struck at \$30/share (e.g. 16.6 m shares), with an indemnity trigger of \$500 m and a maturity of 12 months. In order to exercise the option, ABC must preserve a statutory equity base of at least \$850 m. Assume the following two scenarios: Scenario 1: within the next 12 months, ABC's catastrophic risk portfolio continues to perform reasonably well; though the firm sustains losses of approximately \$75 m, these can be managed within its loss reserves, and the insurer has no need for new equity; in fact, ABC's stock price trades at \$36/share. The LEP expires worthless. Scenario 2: a bad hurricane season leaves ABC with \$600 m of losses in its portfolio; the market, concerned with the news, pushes the stock price down to \$20/share. ABC exercises its LEP against XYZ, delivering 16.6 m new

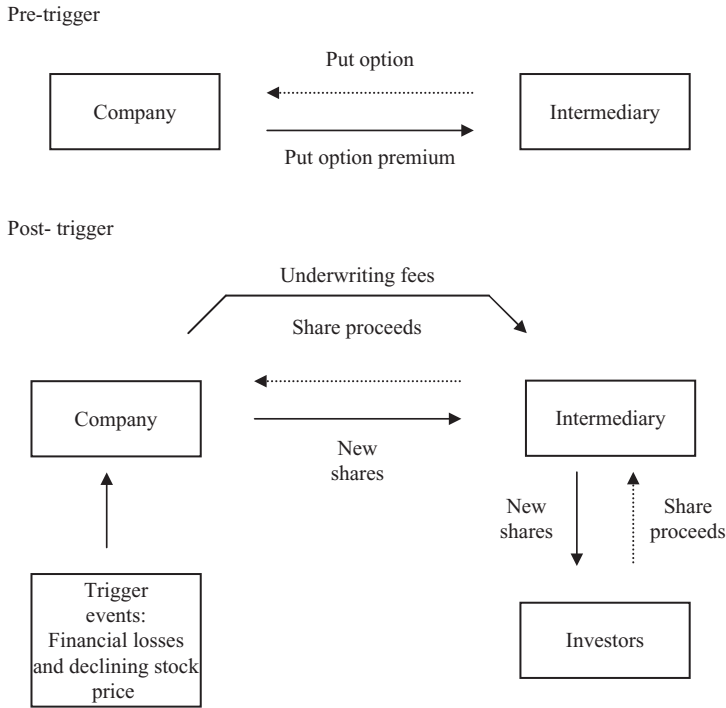


Figure 7.6 Loss equity put structure

shares at the strike price of \$30, for gross proceeds of \$500 m. The increase in capital helps stabilize ABC's financial condition, and the stock price eventually rebounds.

Put-protected equity

A second form of contingent equity is put-protected equity (PPE), where a company buys a put on its own stock in order to generate an economic gain as its stock price declines in the aftermath of a loss. Under a typical PPE, a company purchases a put from an intermediary, defining the number of shares, strike price, and maturity (as in any option). If the company suffers a major loss from the event it is seeking to protect itself against, there is a significant chance that its stock price will decline. The company then exercises against the put seller for an economic gain. The economic gain can be used to increase retained earnings alone, or to provide a price hedge against the issuance of a new tranche of stock. If the PPE is used only to generate an economic gain, based on a decline in the company's stock price (e.g. the company is delivering to the intermediary shares that it purchases in the open market), the addition to equity capital is indirect, rather than direct. That is, it accrues to the retained earnings account rather than the paid-in capital account, meaning that the after-tax proceeds will be lower, and there will be no share dilution. If the PPE is used to protect the company's issuance of new shares, then the gain on the put serves to offset the increased number of shares that will occur when the stock is issued at the new lower price. As with the LEP, if new common shares are

issued, dilution is a factor; if preferreds are involved, no dilution concerns arise. Unlike the LEP, the PPE does not need to be governed by a specific loss trigger; that is, the company might simply purchase a put on its own stock under the assumption that any sizeable loss will be sufficient to cause downward pressure on the stock price. However, PPEs can be viewed negatively in the market. If investors become aware that a firm is buying puts on its own stock, they might be concerned that there is bad news ahead; the share price may thus be forced down as investors sell their shares – and not because of any specific loss.

Convertible Bonds and Equity Hybrids

8.1 INTRODUCTION

Convertible bonds are one of the oldest, and most popular, of the financial sector's structured products. The basic convertible bond, which combines a bond and an investor-owned equity call option in a single package, has been in use since the 19th century. The combination of fixed-income returns through coupons and potential credit spread tightening, coupled with the possibility of conversion into equity, has proven to be an appealing and enduring structure. In fact, the success and popularity of convertibles over a period of many decades has led to the creation of various structured equity hybrid securities that fuse elements of the fixed-income and equity markets.

In this chapter we examine a variety of structured equity products, including conventional convertible bonds and convertible variations (zero coupon convertible bonds, exchangeable bonds, reverse convertible bonds, and mandatory convertible bonds), bonds with equity warrants,¹ buy/write (covered call) securities, and other equity hybrids (adjustable-rate preferred stock, convertible preferred stock). These categories are summarized in Figure 8.1. These securities should be compared and contrasted with the general class of equity-linked structured notes discussed in Chapter 5. Though both sectors can provide investors with exposure to the equity markets through a fixed-income vehicle (and can simultaneously provide issuers with attractive funding opportunities), nearly all of the synthetic and structured products discussed in this chapter can result in the creation of new equity;² structured notes, in contrast, result only in the reallocation of equity exposure from one party to another.

8.2 DEVELOPMENT AND MARKET DRIVERS

The first US convertible bond, floated by the Erie Railroad in the 19th century, was a pioneering issue that paid investors a fixed coupon and granted them a call option to purchase Erie's common stock at a predetermined stock price (i.e. the conversion price). The instrument gave investors a minimum fixed-income return (along with potential capital gains from credit spread tightening) while the stock traded below the conversion price, along with the potential to convert the bond into new shares once the market price exceeded the conversion price. Once converted, investors relinquished the fixed-income investment, but had the opportunity to participate in the issuer's equity dividends and any further upside equity gains. The flexible features and multiple return possibilities proved popular, and paved the way for subsequent issues in the US, Europe, Japan, and the emerging markets. By the mid-20th century, standard convertible bonds had become a mainstay of corporate financing, primarily for large companies. During

¹ Bonds with other types of attached warrants, including those that are based on currency, and commodity references, are also issued in the marketplace; these, however, represent a very small share of the marketplace and tend to be floated on an opportunistic, private-placement basis as market conditions allow. Consistent with the equity focus of this chapter, we shall confine our discussion to bonds with equity warrants.

² The only exceptions occur with certain (though not all) synthetic PERCS, as well as the original PRIMES/SCORES structures.

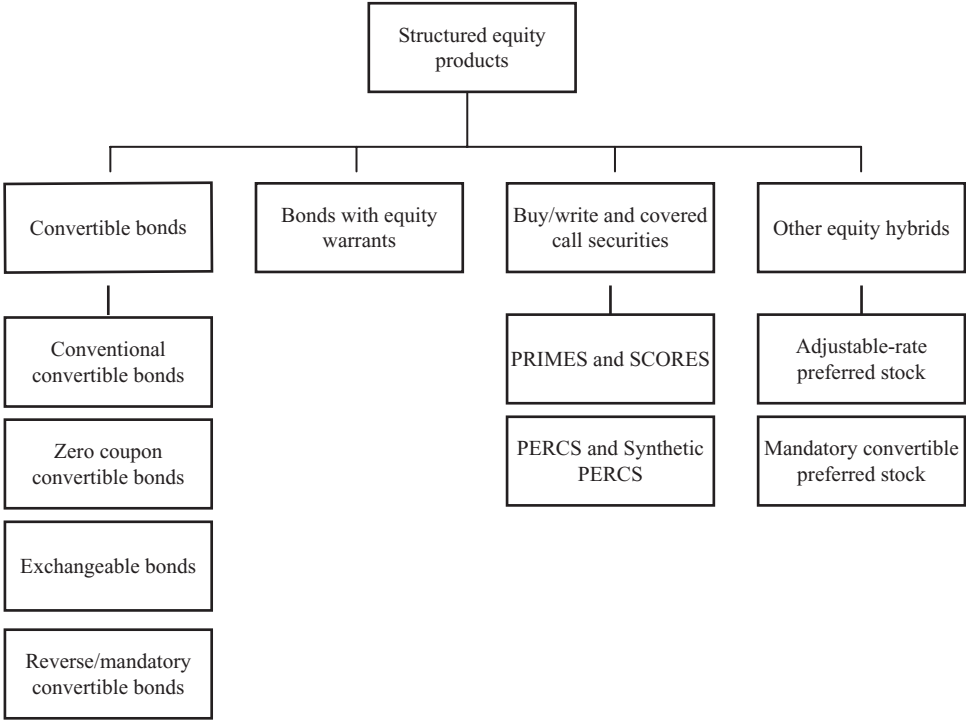


Figure 8.1 Structured equity products

the late 1980s and into the 1990s, financial intermediaries began experimenting with the basic convertible bond structure, developing a series of hybrid securities with fixed income and equity features that addressed specific funding and investment requirements.

Steady primary and secondary market expansion has occurred over the past few decades, and the convertible market has become global: issuers from many countries routinely float convertibles in a variety of currencies and place them on a public or private basis with institutional investors throughout the world. The secondary market for many issues is quite liquid, and the ability for intermediaries to short the stock of issuers allows them to operate significant market-making operations without absorbing an excessive amount of risk.

Bonds with attached equity warrants emerged in the early 20th century (e.g. American Light and Power’s inaugural issue), but became widely utilized during the 1980s and 1990s, as issuers sought to capitalize on cheap financing opportunities through the Eurobond market. Indeed, the Eurobond market has a long tradition of arranging and placing these types of issues, with issuance escalating during the strong bull markets of the 1980s. Japanese issuers were particularly significant participants, taking advantage of the rapidly rising stock market³ and certain accounting rules⁴ to obtain very favorable financing rates. Following this period of speculative excess, the market returned to more disciplined issuance, a process that continues

³ Unfortunately, the subsequent bursting of the economic bubble in the early 1990s caused stock prices to plummet and warrants to expire unexercised. This, in turn, forced Japanese issuers to seek alternate sources of refinancing for maturing host bonds, meaning the corporate sector retained a significant amount of leverage.

⁴ For instance, until 1994, Japanese accounting rules allowed bonds to be issued at a discount with no value ascribed to the equity warrant; since the annual coupon on the structure was smaller than that of a conventional par bond, financing appeared to be much

to the present time. Convertibles and warrant bonds led ultimately to the creation of additional structured products, including issuer-arranged buy/write (or covered call) securities and mandatory convertible stock. Though the earliest versions of the buy/write instrument appeared in 1983–1985, true activity did not take hold until the late 1980s. As equity markets regained strength following the 1987 crash, issuers discovered that buy/writes and mandatory structures provided cost-effective funding and the possibility of future equity proceeds, while investors realized the advantages of generating returns from multiple market variables. In fact, various market drivers have influenced activity in structured equity products; some of these drivers apply to all of the instruments, some more narrowly to one or two classes (e.g. convertibles, bonds with warrants). Specifically, the instruments considered in this chapter permit:

- lower issuer financing costs than standard fixed-income securities with equal maturity (e.g. convertibles);
- flotation of new equity capital upon conversion and subsequent balance sheet deleveraging⁵ – with less dilution than a transaction executed in the current market;
- return potential from both fixed-income and equity movements; the ability to claim a minimum fixed-income return while retaining potential equity market upside;
- access to multiple investor bases (e.g. equity investors, fixed-income investors, crossover investors), helping with primary and secondary placement, and increasing issuer name recognition.

Competitive cost of funding is a key motivating force for many issuers. Indeed, a convertible can be cheaper in the short to medium term than other debt alternatives, as the coupon on the bond is lower than the coupon on an otherwise equivalent straight bond. If conversion into equity occurs, dilution will follow,⁶ but so will a decrease in leverage; post-conversion deleveraging can lead to an improvement in the issuer's credit rating, which can serve as a catalyst for lower funding costs on subsequent debt-related issuance.

Convertibles, bonds with warrants, and other debt/equity hybrids generally are regarded as being marketable to a broad base of investors, as they can appeal to both equity and fixed-income buyers; this is particularly true in the new issuance phase, before a security becomes seasoned and assumes the characteristics of either of the two underlying assets (which is, of course, primarily a function of the direction of the issuer's stock price). In practice, equity investors tend to prefer instruments that are either at- or in-the-money; fixed-income investors, in turn, prefer out-of-the-money convertibles or warrant bonds, where the bond component accounts for 80–90 % of the value of the security and the option/warrant is seen only as an upside “kicker,” or where opportunities for credit spread tightening appear favorable.

8.3 PRODUCT MECHANICS AND APPLICATIONS

In this section we consider the mechanics and practical applications of the key structured equity instruments, beginning with the convertible bond, which arguably is the most popular instrument in the sector. We then analyze extensions of, and substitutes for, the convertible.

cheaper from an accounting perspective. The practice was disallowed in 1994, forcing companies to account for the potential value of the equity component.

⁵ Except for bonds with warrants, unless the equity proceeds are used specifically to retire the bonds.

⁶ Though dilution is an inevitable result of any issuance of common stock, there is some empirical evidence to suggest that, apart from the inherent conversion premium, equity raised through a convertible issuance may be priced at a higher level than equity floated via a rights issue or secondary offering; this means that, for a given dollar amount of equity raised, less shares are required – dilution is therefore lower.

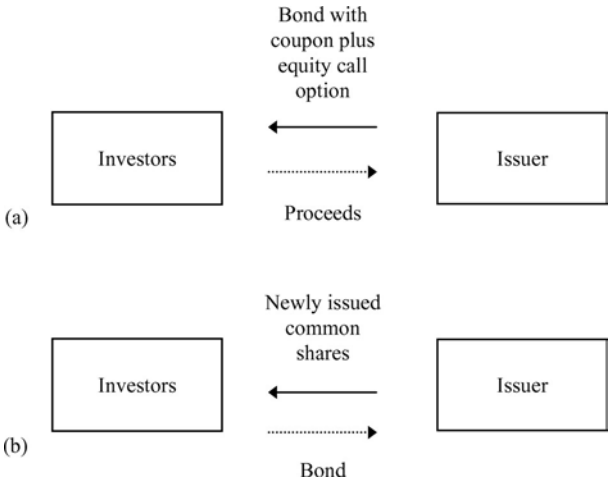


Figure 8.2 Convertible bond. (a) At issuance; and (b) at conversion

8.3.1 Convertible bonds and variations

Activity in the convertible bond sector centers primarily on the conventional convertible we have mentioned above. Most issuance, investment, and trading activity occur in standard domestic and international convertibles, and the market has become particularly deep and vibrant over the past few decades. The success of convertibles has also led to the development of certain convertible variations, including zero coupon convertible bonds, exchangeable bonds, reverse convertible bonds, and mandatory convertible bonds; these may be seen as substitutes and/or competitors to the standard instrument.

Convertible bonds

A conventional convertible bond is a debt/equity hybrid that allows the investor to convert a fixed-income instrument into a specified number of shares once a particular conversion price is reached; the issuance and conversion flows are illustrated in Figure 8.2. Assuming conversion occurs, the investor presents the bond for redemption (so forfeiting the right to future coupons, credit spread appreciation, and principal redemption) and receives shares that may pay a dividend and allow for further capital appreciation. From a building block perspective, the convertible can also be viewed as a package comprised of a share of common stock, a put to exchange the common stock for a bond, and a swap of the coupon stream for dividends until the conversion date. The issuer can generate an attractive cost of funds⁷ and minimize the impact of share dilution,⁸ while the investor may be able to earn a minimum

⁷ The actual cost of funds depends on the issuer's credit rating/credit spread (which dictates the fixed coupon level it must pay, generally as a spread to a risk-free benchmark) and the level of the stock price (which determines, ultimately, whether new equity will be floated). During the period in which the bond remains unconverted, the cost of funds is computed on a net basis, e.g. after deducting interest payments from taxable income; once converted, the payment of dividends is subject to double taxation in some jurisdictions (e.g. at the corporate and investor levels).

⁸ If conversion occurs, the issuer faces a lower level of dilution than if it were to issue common stock at the current market price; since most convertibles feature conversion premiums in excess of 20 %, an issuer raises more equity capital on conversion, and thus decreases the overall level of dilution.

interest rate return while preserving potential upside via capital appreciation. The investor, long a call option, must pay the issuer a premium; this can be done by establishing a below-market coupon, or setting a higher market premium over investment value (as discussed below).

While the standard convertible bond can provide an issuer with important equity funding and deleveraging opportunities, it also carries potential disadvantages. In particular, although new equity may ultimately result from the flotation of a convertible, there is no guarantee of conversion – unless the conversion price is set at artificially low levels, which increases the level of dilution for a given target level of equity financing. In addition, and unlike the flotation of common stock via a standard rights issue or public offering, the issuer has no control over the timing of the conversion/issuance. Furthermore, if conversion does not ultimately occur (i.e. the stock price never reaches the appropriate conversion level), the issuer must arrange for refinancing when the convertible is presented for redemption at maturity. It is also important to note that if conversion does not occur, the effective cost of funding may compare unfavorably with that of a standard straight bond; yield-to-maturity will increase as maturity draws closer, particularly for zero coupon and puttable structures. Accordingly, in order for the funding advantages to be realized fully, an issuer must be relatively confident that conversion will occur over the medium term. Note also that in some jurisdictions, issuers may benefit from favorable tax treatment that can lower all-in funding costs: coupon payments on convertibles that have not been converted into common equity are often tax deductible, while dividend payments on common equity may not be.

Since convertibles are true hybrid products, absorbing elements of the debt and equity asset classes, they expose investors to a range of risks including credit spread/default risk, interest rate risk, currency risk (for foreign currency issues), volatility risk, dividend risk, equity price risk, and call risk. Naturally, not all of these risk factors exhibit the same sensitivity at the same time; most depend on the state of the market and the value of the convertible in relation to the issuer's stock price. For instance, when an issuer's stock price is well below the conversion price, the convertible trades more like a bond than an equity, and is thus far more sensitive to interest rates and interest rate volatility than equity prices and equity volatility. The reverse is true as the stock price rises to the conversion price.

Convertible issuers must determine *ex ante* an appropriate conversion price and a strategy to manage balance sheet leverage should conversion not occur. Establishing a conversion price is a balancing process: the higher the conversion price (i.e. the farther out-of-the-money the equity option), the cheaper the funding and the lower the eventual dilution, but the lower the probability of conversion; the lower the conversion price (i.e. the closer-to-the-money the equity option), the more expensive the funding and the higher the eventual dilution, but the higher the probability of conversion. Naturally, most convertible issuers seek to have their bonds converted at some point, for at least two reasons: the act of converting lowers leverage through the retirement of the fixed-income portion of the bond and the issuance of new equity; and a convertible that remains unconverted for a long period of time (i.e. a so-called "busted" or "broken" convertible) can create uncertainty and stock price pressure for the issuer.⁹

A convertible bond is defined by a number of parameters that help crystallize issuer and investor value. Some of the most important include the conversion ratio, conversion price,

⁹ Although this can occur gradually as an issuer's stock price drifts lower over months or years, it can also happen at the launch of a new issue if the underwriter fails to price the deal properly. If this happens, the failed new issue can weigh on the issuer's stock price, and may require some form of refinancing.

conversion value, market conversion price, market conversion premium, and market premium over straight value.

- **Conversion ratio:** the number of shares that the investor will receive upon conversion, computed as bond par value divided by conversion price; the conversion ratio is adjusted for splits and dividends and may decline over time.
- **Conversion price:** the price at which an investor has the right to purchase shares of the issuer's common stock, computed as bond par value divided by the conversion ratio.
- **Conversion value:** the value of the bond if it is converted at the current stock price, computed as the conversion ratio times the current stock price. This value is sometimes also known as parity.

For example, if a bond is issued with a par value of \$1 000 and a conversion ratio of 50, its conversion price is \$20. Similarly, if the bond is worth \$1 000 and has a conversion price of \$20, its conversion ratio is 50. At a price of \$15, the conversion value is \$750. Both the conversion ratio and the conversion price remain unchanged once they are set, unless an issuer alters its capital structure during the life of the deal.

A typical convertible paying a regular coupon features the essential characteristics of a standard bond, meaning that certain minimum values can be ascribed to the security. For instance, the minimum price of a bond must be the conversion value or the straight bond value (i.e. a nonconvertible bond with the same coupon and maturity). Since the convertible bond contains the embedded equity option, it cannot sell for less than a straight bond; arbitrageurs will ensure this remains true by purchasing the undervalued convertible bond and converting it for a riskless profit. For instance, if the issuer's common stock is trading at \$17 and the conversion ratio is 50, then the conversion value of the convertible bond is \$850. If a comparable straight bond is trading at \$780, then the convertible cannot trade below this amount. If it does, the arbitrageur purchases the convertible (at, say, a price of \$780), and short sells 50 shares of stock; when the bond is converted and the short is covered with the new shares, an arbitrage profit of \$70 is generated.

- **Market conversion price:** an associated conversion indicator, computed as the market price of the convertible divided by the conversion ratio. This serves as a gauge of the security's breakeven point, because once the issuer's stock price exceeds the market conversion price, any further stock price increase raises the value of the convertible bond by the same percentage.
- **Market conversion premium:** a measure of the premium an investor pays over the current price of the stock as a result of the bond value (floor), computed as the market conversion premium per share divided by the market price of the common stock. A convertible investor generally receives a coupon on the convertible that is higher than the dividend yield on the common stock, because it has to absorb the market conversion premium. Note that the conversion premium can also be obtained by dividing the bond price by the conversion value.
- **Market premium over straight value:** a measure of the premium an investor pays for the embedded equity option, calculated by dividing the market price of the convertible by the value of the straight bond. The greater the premium, the less attractive the convertible investment (note that this measure can also be impacted by changes in interest rates). This amount is sometimes referred to as investment premium.

Since the convertible has an embedded call option, we can view the payoff profile of the security, from the investor's perspective, using a version of the option payoff diagram introduced




Figure 8.3 Convertible bond payoff profile

in Chapter 2. Knowing that the conversion value is simply the product of the conversion ratio and the price of the common stock, it follows that conversion value increases as the value of the stock increases; the actual angle of the axis emanating from the origin is a direct function of the conversion ratio. Note, further, that the bond value acts as a floor in establishing a minimum price. This relationship is illustrated in Figure 8.3.

The value of a convertible increases with rising stock prices, rising volatility, tightening credit spreads, falling rates, lengthening call protection, and declining dividend yields. By extension, it decreases with falling stock prices and volatility, rising rates, widening credit spreads, rising dividend yields, and shortening call protection. These relationships, summarized in Table 8.1, are intuitively and quantitatively supported.

There is no *ex ante* way of determining whether a convertible bond will offer a superior investment/funding return over common stock or straight bond alternatives: a convertible can increase investor returns if the common stock price rises, and provide downside yield protection if the common stock price falls. As noted, if a convertible is out-of-the-money, its value becomes less sensitive to the common stock price, but quite sensitive to rates and spreads; it will trade as a bond equivalent, making it more appealing to pure fixed-income investors. This characteristic changes once the stock price approaches, and ultimately exceeds, the conversion price: when the conversion value is well above straight value, the convertible will then trade as an equity equivalent (i.e. the market conversion premium per share becomes increasingly small.) However, it is important to note that if a convertible is well in-the-money, investors will not pay much of a premium to own the bond; as the common stock continues

Table 8.1 Convertible bond valuation variables

Convertible bond value increases with . . .	Convertible bond value decreases with . . .
Rising stock price	Falling stock price
Rising volatility	Falling volatility
Declining dividend yield	Rising dividend yield
Falling interest rates	Rising interest rates
Tightening credit spreads	Widening credit spreads
Lengthening call protection	Shortening call protection

to rise, the premium will continue to compress, and the bond will actually underperform the stock in value terms.¹⁰ The obvious disadvantage of a convertible bond is that the investor sacrifices some upside potential as a result of the market premium; the clear advantage is that the investor mitigates some downside equity risk through the bond floor.

Investors often examine the return tradeoffs that exist between an investment in a convertible and other alternatives. The process begins with the calculation of a flat yield, which is simply the coupon on the convertible divided by the market price of the bond. The yield advantage can then be determined by computing the difference between the flat yield and the gross dividend yield on the issuer's common stock. Next, a simple breakeven is determined by comparing the time it takes for the yield advantage to cover the conversion premium (in years); the shorter the breakeven period, the more attractive the investment, and vice versa. Since the simple breakeven measure does not account for the timing of cash flows or dividend growth rates, it is rather simplistic. More precise breakeven computations can be used, including those that focus on time value of money, dividend growth, credit spread movement over time, and the changing value of the embedded option over time.

From a basic and theoretical perspective, a noncallable/nonputtable convertible has a single investor equity call option, meaning that the value of the security can be decomposed into two components: straight bond value and call option value. Straight bond value can be derived from any reasonable fixed-income valuation model that accounts for both interest rates and credit spreads. The call option value can be modeled via a standard Black–Scholes or binomial process, such as we have summarized in Chapter 2. If the convertible is callable by the issuer (and in practice many of them are), or if it is puttable by the investor (which is rather less common), then additional options are inserted into the structure, and a binomial framework is required. A callable convertible may feature hard call protection, soft call protection, or both. Hard call protection is designed to protect the convertible investor by prohibiting issuer calls for a set period of time, generally three to five years. Soft call protection, in turn, protects the investor until the issuer's common stock price reaches a certain percentage of the conversion price (e.g. 130–150 %). When a convertible is called, investors can accept the call price or convert into shares. The optimal action is based on whether the call price plus accrued interest is above or below conversion value; if conversion value is above the call price, investors should convert into the underlying shares. Since issuers generally hope to have their bonds converted, they typically call when the conversion price is at least 5 % above the sum of the call price and accrued interest. There are, however, instances when issuers will call when conversion value is well below the call price; this occurs when interest rates have declined markedly, and it becomes cheaper to refinance with debt.

Some convertibles are issued on a puttable basis, obligating issuers to redeem the bond on one or more put dates. This feature gives investors the right to redeem the bond in cash terms (a so-called hard put) or shares (a soft put). Since the investor essentially purchases a put option from the issuer through this structure, the coupon received on the bond is lower than it would be on an otherwise comparable nonputtable convertible security. Puttable convertibles can be put back to the issuer at face value (par put), or at a premium to face value (premium put), on a specific date (or dates). The single premium put structure allows only a single exercise opportunity, while the rolling premium put structure permits multiple put opportunities, with the put premium typically rising on each subsequent put date.

¹⁰ Convertibles with call protection tend not to face the same premium compression, and can actually perform well versus the common stock.

A conventional convertible is defined by issuance market, issue price, conversion ratio, conversion price, conversion premium, coupon, currency, maturity, and issuer call/investor put features. Convertibles are issued generally as the equivalent of subordinated debentures, ranking below all other forms of debt, but ahead of any preferred issues (including convertible preferreds discussed below), in terms of default seniority. They may be floated in a national market or via the Euromarkets: the domestic US, UK, Swiss, and Japanese markets account for approximately 75 % of activity, while the Euromarkets account for the remaining 25 %. Regardless of market, bonds generally feature semi-annual or annual coupons; those floated in the Euromarkets are free from withholding taxes. Securities are often issued in bearer form, and trade on an OTC basis (but listed on a national exchange, e.g. New York, London, Luxembourg).

Though a convertible may be callable or noncallable, it must have a defined maturity date that forces repayment of investor proceeds (generally at par) if conversion has not occurred. While pricing of convertibles depends heavily on the creditworthiness of the issuer and the state of the overall market, standard structures typically feature conversion premiums of 15–25 %, and final maturities of 10–15 years; coupons generally are set well above the common stock's current dividend yield for the reasons noted above (but below those of otherwise comparable straight bonds). Low-premium convertibles, which are rather less common, are priced with premiums of approximately 5 % in order to promote rapid conversion; coupons on such structures are only slightly higher than current dividends. Secondary trading in many convertible bonds is very active, and liquidity can be good, though this is both market- and issuer-specific; financial intermediaries (including underwriters and risk-takers) act as market-makers, quoting two-way prices for many issues. The level of market interest rates and interest rate volatility, and the trend of equity markets and equity volatility, are important in determining the volume of secondary trading. Trading in foreign currency convertibles also depends on foreign exchange rates and interest rate differentials. From a micro perspective, liquidity is a function of an issuer's bond structure and the level and trend of its stock price and credit spreads; for large market-makers seeking to run hedged books, it is also dependent on the availability of issuer shares for shorting purposes.

Zero coupon convertible bonds

The zero coupon convertible is a deep-discount convertible security issued primarily in the US markets.¹¹ Like conventional convertibles, zero coupon versions of the security generally are priced with conversion premiums of 15–25 %, and final maturities ranging from 10 to 30 years. The typical structure is a package comprised of a zero coupon bond, an investor call option on the issuer's stock price, an investor option on the issuer's credit spread, and an issuer's call option on interest rates. The zero coupon convertible pays no periodic coupon (simply an accreted value payable at maturity or conversion), and redeems at par if not converted. Investors can benefit from both credit spread tightening and equity appreciation, and retain a rolling put feature; the put strikes generally rise steadily as maturity approaches, in order to preserve a constant yield-to-put. The issuer, through the callable feature, is protected from rising interest rates. All other benefits and risks found in conventional convertibles apply to zero coupon convertibles.

¹¹ Zero coupon convertible bonds debuted in 1985, when Merrill Lynch launched an issue known as a LYON (Liquid Yield Option Note) for Waste Management Inc. Since that time, many other intermediaries have launched similar types of structure.

Exchangeable bonds

The exchangeable bond is another variation on the convertible bond. This security functions just as any conventional convertible bond, except that exercise by the investor of the embedded equity option results in conversion into the shares of an issuer's affiliate, rather than the shares of the issuer itself. Thus, if Company ABC floats an exchangeable bond that is convertible into shares of affiliated Subsidiary XYZ, investors receive XYZ shares upon conversion (surrendering the ABC bond in the process). Such transactions are not particularly common, but they do appear in the marketplace on occasion. They are arranged typically for listed companies and their affiliates, but may also be floated on behalf of those in the pre-initial public offering stage.

Other convertible bonds

Certain other convertible bond variants have appeared over the years, including the reverse convertible bond and the mandatory convertible bond. These securities are quite similar to the conventional convertibles we have just described, except as related to the conversion option. In a reverse convertible, the issuer, rather than the investor, holds the conversion option. This means that the investor is long a fixed-income security and short a call option, and thus receives premium via an enhanced coupon. The reverse convertible allows the issuer to make the decision regarding conversion; this adds a degree of certainty to the timing of equity funding, though it still does not guarantee that an equity infusion will occur (e.g. the convertible may remain out-of-the-money). The mandatory convertible, in turn, is a security that requires the issuer to convert the fixed-income instrument into a preset amount of shares at some future time and price; neither the investor nor the issuer have any flexibility in the matter. We shall consider the mandatory convertible structure in greater detail below.

8.3.2 Bonds with warrants

The bond with equity warrant package, like the convertible, is a capital markets structure that allows investors to generate fixed-income earnings while retaining an equity upside. The issuer's sale of the warrant with the bond¹² generates a lower financing cost (e.g. an amortized receipt of premium that lowers the coupon on the bond) and allows balance sheet deleveraging if the issuer chooses to use new stock proceeds to redeem/retire the bond component; the package can thus be regarded as a form of synthetic convertible. Though the product first became popular with lower-rated companies seeking to raise financing by including an "equity kicker," it soon became widely accepted and utilized by a range of issuers, and is now part of the financial mainstream.

Most packages feature maturities of three to seven years, and contain warrants that are exercisable on either an American or European basis. In some instances, the warrants can be detached and traded separately. The issuer retains the initial decision on whether to make the warrants detachable from the host bond (in some instances, the detachability feature commands a small premium); the investor must then determine whether the structure should actually be broken apart, a decision that depends on specific goals and market outlook. Though bonds with warrants are similar to the convertibles discussed above, this is only true while the warrants

¹² In some instances, an issue of preferred stock, rather than a bond, serves as the host; this, however, is an unusual structure.

remain attached to the bond;¹³ once they are separated, the two component parts must be considered and valued separately (e.g. as a straight bond and a long-dated call option). In fact, much of the market is comprised of detachable structures, and is dominated by investors that specialize in holding either the credit or equity portions of the instrument.

The strike price of the equity warrant generally is set at a 25%+ premium above the market price at issuance, in order to increase future equity proceeds and decrease dilution. Most transactions feature antidilution provisions, so that investors are protected from any new equity issuance (e.g. new equity-linked issues, special dividends). In some instances, an issuer is permitted to decrease the strike level unilaterally if the market drifts away from the exercise level; such reset provisions are, however, relatively uncommon. If exercise occurs, the investor must deliver the warrants and a predetermined amount of cash in order to receive the new shares; the bond component remains outstanding until the issuer redeems it (which it may choose to do immediately with the new equity proceeds). Note that a relatively rare variation of the product, the usable bond with warrant, allows an investor to deliver the bond, instead of cash, upon exercise of the warrant (making it similar to a conventional convertible, and forcing the issuer to retire the bonds/lower leverage).

Most bond with warrant issues are floated in the Eurobond market; indeed, US companies seeking this type of financing issue securities offshore, where primary interest and secondary liquidity are far deeper. Secondary trading activity in bonds, ex-warrant bonds, and detached warrants is issuer-specific: the largest issuers tend to feature relatively active trading, particularly in the equity component; deals of smaller issuers may be quite illiquid. In certain jurisdictions, warrants are considered to be capital instruments, rather than derivative contracts, meaning that they must be registered with the relevant securities authority, and must adhere to trading rules related to margining and settlement.

8.3.3 Buy/write (covered call) securities

Buy/write, or covered call, securities have evolved into an important equity hybrid. These assets are built on the premise that equities can be decomposed into two separate components: one component that pays dividends, and a second one that pays capital gains. Since certain classes of investors may have a preference for one stream of income versus the other, an opportunity exists to sell the underlying equity into an SPE or trust, and allow the vehicle to issue structured securities that represent the individual components of the income stream: the dividend tranche pays all accumulated dividends and a redemption value (also known as a termination claim) that excludes any capital appreciation, while the capital gains tranche pays any accumulated value above the redemption value, but no dividends. These flows are highlighted in Figure 8.4.

Examining this structure in light of the building blocks discussed in Chapter 2, we note that the dividend tranche is precisely equal to a long position in the underlying stock, and a short call struck at the redemption value (as noted in Figure 8.5), while the capital gain tranche is equal to a long call struck at the redemption value. Though writing a covered call generally is considered to be a low-risk strategy (e.g. the investor writing the option has sufficient shares to cover the liability in the event of exercise), the strategy still has risks, including the risk that the security will underperform versus the naked equivalent, i.e. giving up some upside potential and

¹³ A conventional convertible can only be decomposed if it is sold into an SPE or trust, which then issues separate synthetic securities/receipts (or enters into a total return swap) that replicate the economics of each individual component. Such “convertible stripping” is, in fact, a common arbitrage technique, but requires the creation of the intermediation vehicle; decomposition of the bond with warrant requires no additional steps.

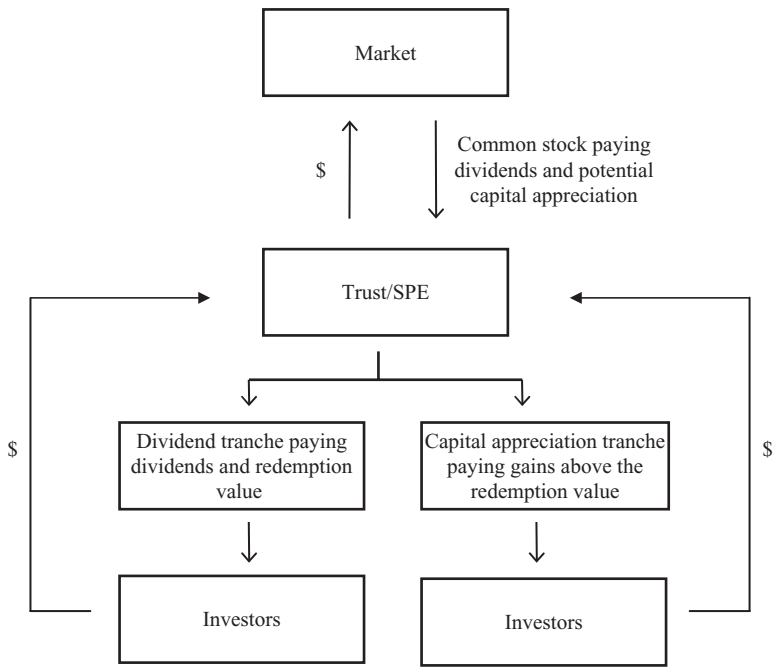


Figure 8.4 Two tranches of an equity security

assuming some downside risk that is only partly compensated by the “premium.” Nevertheless, various incarnations of buy/write securities have proven popular with sophisticated investors, as they permit the purchase of a precise risk/investment position that might otherwise be inefficient to create on a standalone basis.

Buy/write securities can be issued via standard underwriting methods or through a special distribution mechanism, and may be issued at a premium or discount to par. Dividends on the

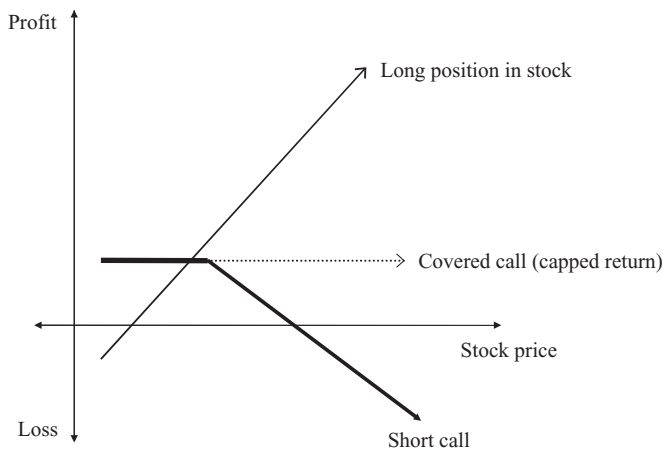


Figure 8.5 Covered call tranche

issue are generally payable quarterly, and some issues allow for the transfer of voting rights to the investor. Principal redemption is a function of the stock price: if the stock price is less than (or equal to) the cap price determined at launch, the investor receives one share per unit; if it is greater than the cap, the investor receives shares at a rate determined by the ratio of the cap price to the value of the common stock. Since the value is capped, the number of shares received declines as the stock price rises.

In the section that follows we consider several types of the buy/write security, including PRIMES and SCORES, PERCS, and synthetic PERCS. As noted earlier in the chapter, PERCS and certain synthetic PERCS result in the issuance of new equity upon exercise; PRIMES and SCORES, and various classes of synthetic PERCS result only in the reallocation of existing equity, and are thus nondilutive (indeed, such securities can be considered forms of the equity-linked structured note we describe in Chapter 5).

PRIMES and SCORES

The inaugural buy/write security appeared in 1983 via the Americus Trust, a special trust established to divide securities into separate cash streams.¹⁴ Though the PRIME/SCORE product is no longer issued, it is worth investigating its design, as it paved the way for a new generation of similar buy/write structures.

The Americus Trust was created as a redeemable unit investment trust, with units that could be split into a PRIME component (Prescribed Right to Income and Maximum Equity) and a SCORE component (Special Claim On Residual Equity); the two components could also be recombined at any time. The termination claim (strike) on the PRIME unit was equal to a dollar limit on the potential for capital appreciation: all value up to the termination claim, plus dividends, was distributed to investors in the PRIME unit, while all value in excess of the termination claim was directed to investors in the SCORE unit. The two units were thus equivalent to the dividend and capital appreciation securities noted in Figure 8.5 (i.e. PRIME as a covered call and SCORE as a long-dated call warrant, meaning the PRIME investor was long the stock and short the SCORE). Investors tendered shares into the trust and received a unit for each share tendered (with specific expiration and termination claim); the units could then be split into the two different components. Each unit had a five-year maturity and a strike set at approximately 25 % above issue price, and was collateralized by the underlying shares, which were held by a trustee. If the common stock price ended below the termination claim price, the PRIME investor received a share of stock, along with cash dividends, the dividend received deduction, and full voting rights; the SCORE investor's contract expired worthless. As the stock price traded further below the termination claim price, the PRIME traded more like equity, and as it rose above the termination claim price, it became more "debt like"; indeed, as the stock price moved well above the termination claim, the PRIME became a discount bond, paying a below-market coupon. (This behavior is, of course, precisely the reverse of a convertible's, which is logical, since the PRIME featured an embedded short, rather than long, call option.)¹⁵ If, however, the stock price ended above the termination claim price, the

¹⁴ The first deal, launched on AT&T stock, was arranged by Kidder Peabody; a second trust was floated by Furman Selz on Exxon stock in 1985. Alex Brown then decided to take a leading role by negotiating the distribution rights for the remaining trusts that had been filed with the SEC, and floated deals on a number of stocks in 1986 and 1987. The October 1987 stock market crash curtailed further activity, and new products, such as PERCS, eventually emerged to replace the role filled by the Americus Trust.

¹⁵ From a theoretical perspective, an out-of-the-money SCORE and a corporate bond floated by the same issuer could be repackaged to create a synthetic convertible bond.

PRIME holder received a fractional portion of shares, equal to the termination claim divided by the closing price. The SCORE investor, in turn, received value equal to the closing price minus the termination claim, divided by the closing price. The SCORE could be exercised at any time by combining it with a PRIME and submitting it to the trustee in exchange for a share of stock.¹⁶ SCORES were nondilutive (e.g. third party warrants rather than issuer warrants, as they did not result in the creation of new equity), but were considered by regulators to be securities rather than derivatives; the long-dated nature of the transactions gave investors the opportunity to purchase long-dated call options that were otherwise unavailable in the market.¹⁷ The exchange-listed PRIME and SCORE issues were of limited size, meaning they were prone to illiquidity.

PERCS and synthetic PERCS

The next generation of the buy/write structure appeared in 1988, when Morgan Stanley launched an inaugural issue of PERCS (Preference Equity Redemption Cumulative Stock) for Avon. Though PERCS evolved naturally from the PRIME/SCORE, the security is structurally different in several key respects: the issuer, rather than the investor, holds the equity option, and the exercise of the option results in the creation of new shares, and is thus dilutive.

PERCS are buy/write securities based on mandatory convertible preferred stock, and rank senior only to the issuer's common stock.¹⁸ As the name implies, the mandatory convertible requires the issuer to convert into common stock at a specific price and time, which resolves some of the uncertainty related to the funding exercise; indeed, the investor has no right or option over conversion. An investor purchasing a PERCS holds a synthetic instrument comprised of a long position in preferred stock, and a short call option that allows the issuer to convert into common stock once the conversion price is reached;¹⁹ as noted, the long call option is held by the issuer, rather than a different class of investors (as in the typical buy/write security). The investor receives an above-market dividend on the preferred in exchange for granting the right; this is equal to the premium payment amortized over many periods (rather than upfront, as is typical in a covered call transaction). The investor thus faces full capital depreciation and limited appreciation in exchange for a higher return. PERCS are often launched for issuers with common stock that appears to be undervalued. The conversion price generally is set at 30–45 % above the common equity price at the time of launch, and two conversion features are usually incorporated: a mandatory conversion clause that requires conversion on a 1:1 basis by some future date, and a fractional conversion clause that permits conversion on less than a 1:1 basis as soon as the conversion price is reached.²⁰ As the stock price of the issuer rises, fewer shares are needed to redeem the PERCS; this can be regarded as a form of reverse dilution. Note that early redemption does not eliminate the issuer's obligation to pay the investor the entire premium that is due.

¹⁶ Since a share of common stock held by the trustee backed the SCORE, the SCORE investor was not exposed to the counterparty risk of the PRIME investor (i.e. the option seller).

¹⁷ Though the Chicago Board Options Exchange listed LEAPS, or Long-term Equity Anticipation Securities, as long-dated options, the number of issues was limited and the exercise prices and styles were restricted. Long-dated OTC contracts were expensive, and were limited to the sophisticated institutional investor market.

¹⁸ Though PERCS are similar in design to the original Americus Trust securities, they form part of the issuing company's capital structure (just as a convertible or bond with warrant might).

¹⁹ In fact, this structure can be viewed as a form of a covered call or buy/write, with the investor delivering shares of preferred, rather than common, stock if the issuer exercises the call option.

²⁰ Since the offering price of the PERCS is set equal to the common stock at issue date, the issuer essentially sells equity at a price that is no lower than the price at issue date; any subsequent rise in the price from that point on simply decreases dilution.

The value of the PERCS is a direct function of the premium/discount of the security versus the common stock and the present value of the enhanced dividend stream. Stated differently, the value of the embedded call retained by the issuer is equal to the stock price, less the PERCS price, plus the present value of expected dividends. PERCS typically are callable; however, since investors rely on the enhanced dividends over time to generate their returns, the call price must be set at some appropriate premium to par value (i.e. it must be a function of the remaining unaccrued call premium due). In fact, the call price generally is equal to a cap price of 30–50 % above the common stock price at launch, plus expected incremental dividends as of the issue date (declining over time as the incremental amount accrues and converges to the cap price). Though PERCS are similar to common shares, ultimately they feature several key differences: PERCS rank senior to common shares in the event of bankruptcy, they are capped on the upside, and they pay an enhanced dividend in exchange for the cap; however, like common stock, PERCS can be an expensive form of funding, and dividend payments are not deductible.

Following the success of PERCS, various intermediaries began creating their own buy/write structures, many of them following the synthetic PRIME/SCORE/PERCS model. In fact, proprietary buy/write debt obligations, such as Debt Exchangable for Common Stock, Dividend-enhanced Convertible securities, Common-linked Higher Income Participation securities, Yield-enhanced Equity-linked securities, and so forth, emerged in the mid- to late 1990s and are still issued to the present time.²¹ Most carry three to five year maturities, and 30–50 % strike caps, but pay quarterly coupons rather than dividends, and are noncallable and nondilutive. Conventional and synthetic buy/writes are listed and traded on a major exchange; secondary liquidity is best described as moderate, with most activity occurring in the weeks and months following initial issuance.

8.3.4 Other equity hybrids

There are, of course, many other types of equity hybrids in the financial sector; like other financial instruments, some have proven to be enduring, others appear to fill a temporary role or take advantage of very specific market opportunities. Adjustable-rate preferred stock and convertible preferred stock are two of the most popular instruments of the sector.

Adjustable rate preferred stock

Adjustable-rate preferred stock (ARPS), sometimes also known as money market preferred stock or auction market preferred stock, is a synthetic instrument comprised of preferred stock that pays a standard fixed-rate dividend, and an interest swap that exchanges the fixed dividend for a floating interest rate, generally LIBOR. Though companies typically issue ARPS to investors directly, they can also be created by using an intermediate SPE that swaps fixed dividends on an existing preferred stock issue into floating dividends. Figure 8.6 illustrates the individual components that comprise the ARPS. Unlike fixed dividend preferred stock (whether cumulative or noncumulative), which pays a certain amount every quarter, the ARPS structure pays a dividend that floats with interest rates; it is thus appealing to investors that

²¹ Naturally, some products have been introduced and ultimately retired, while others have never actually been floated. For instance, Lehman Brothers proposed an “unbundled stock unit” structure in the 1990s that sought to split equity into various components, including a large debt element; after significant regulatory concerns, the bank decided to cancel the product before the premarketing stage even began.

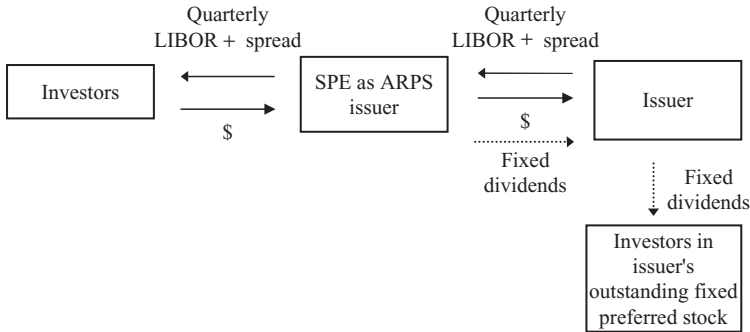


Figure 8.6 Synthetic adjustable-rate preferred stock dividend flows

believe interest rates are likely to rise over time. Like other preferred stock issues, ARPS are subordinated to all other forms of secured, senior, and subordinated debt, but rank senior to common stock. Liquidity in the ARPS issues of major companies can be quite good; since the variable dividend feature of the structure contains a direct link with floating rates, many money market funds are able to hold them in their portfolios.

Note that the reverse structure, known as an adjustable convertible note, also exists. Issues that provide investors with a fixed dividend in exchange for a floating dividend are launched from time to time, though these are far less common.

Convertible preferred stock

Convertible preferred stock is similar to the convertible bond, except that the host security is preferred stock rather than a debenture. Convertible preferreds can therefore be considered a package of preferred stock and an investor call option on the issuer's common stock. The preferred stock flotation vehicle generally pays a fixed quarterly dividend on a cumulative or noncumulative basis, and is subordinated to other forms of debt and convertible bonds; in fact, convertible preferreds rank senior only to the issuer's common stock. The securities provide investors with certain advantages over pure common stock, including a senior position in the issuer's dividend stream, and a ranking senior to common stockholders in the event of bankruptcy; there is thus some downside protection associated with the preferred position (though clearly not as much as with the convertible bond).

Unlike a standard convertible bond with a final maturity, the preferred version is issued as a perpetual security that remains outstanding until (if) conversion occurs. If conversion takes place, investors surrender the preferred stock and receive new common stock; effectively they exchange a subordinated equity-like instrument paying a fixed (and generally cumulative) dividend, for a standard equity instrument that provides potential capital appreciation and/or common dividends. Convertible preferreds are common primarily in the US and UK markets; they are not especially prevalent in other markets, and comprise a much smaller portion of the equity hybrid market than convertible bonds. Note that in some cases, convertible preferreds can be structured with mandatory conversion features, requiring the replacement of preferred with common stock at a future date.

It is important to note that the entire class of equity hybrids can raise legal, accounting, and regulatory challenges. Instruments that can vary between debt and equity depending on the price of an issuer's common stock can be complicated to regulate, value, and administer. Regulators rightly scrutinize any structured or synthetic instrument that attempts to convert equity into debt, and generally establish minimum standards related to classification of debt or equity on the corporate balance sheet.

Investment Funds

9.1 INTRODUCTION

Investment funds comprise the broad class of asset portfolios that are financed through the issuance of a fixed or variable amount of stock, or the distribution of partnership shares. We include funds in our discussion of structured assets, since they transform the equity capital of investors into pro rata interests in an entire portfolio of assets or financial contracts from the equity, fixed-income, currency, and/or commodity sectors. In fact, an interest in a fund referencing an equity market (or any other asset reference) shares economic similarities with a funded total return equity swap or a zero coupon bond with redemption linked to the same equity reference.

In this chapter we consider several of the most popular types of investment funds, including investment companies¹ (open-end funds, closed-end funds, and unit investment trusts), hedge funds, and exchange-traded funds. Open-end funds (known as mutual funds in some countries) comprise the largest segment of the investment fund market and are designed for the broadest group of retail and institutional investors. Hedge funds, designed primarily for sophisticated (or accredited) investors, can invest in a broad range of assets, and typically run aggressive risk profiles. Exchange-traded funds (ETFs) are registered hybrids that combine the structural features of open-end funds with the trading and liquidity features of actively traded corporate securities. Figure 9.1 summarizes the broad classes of investment funds.

9.2 DEVELOPMENT AND MARKET DRIVERS

Open- and closed-end funds, which have existed for many decades, have become an integral part of the investment management sector. The market for open-end funds dates back to the 19th century, when the earliest investment-based conduits were created. Though early efforts were sometimes managed unprofessionally and were often highly speculative, the market matured steadily throughout the 20th century; the sector has experienced significant, and accelerating, expansion over the past three decades, as global wealth has continued to build. Millions of retail and institutional investors around the world now use funds to help them achieve their investment goals. New types of funds have been introduced at regular intervals, including those specialized on individual stocks, bonds, indexes, and commodities. Funds that invest in individual stocks have been the perennial favorites over the long term; many investors are attracted to such investments as the possibility of earning returns that exceed those available in other market sectors is an appealing proposition. While stock funds have been in vogue for many years, index funds that track an entire market index are a more recent phenomenon, with significant growth commencing in the early 1980s. Since the introduction of the first index-based open-end fund by Vanguard in 1976, funds on narrow and broad market indexes

¹ In the US, the investment company classification also includes vehicles known as “face amount certificate companies” that issue fixed-income debentures; these represent a very small portion of the marketplace, and we shall not consider them in further detail.

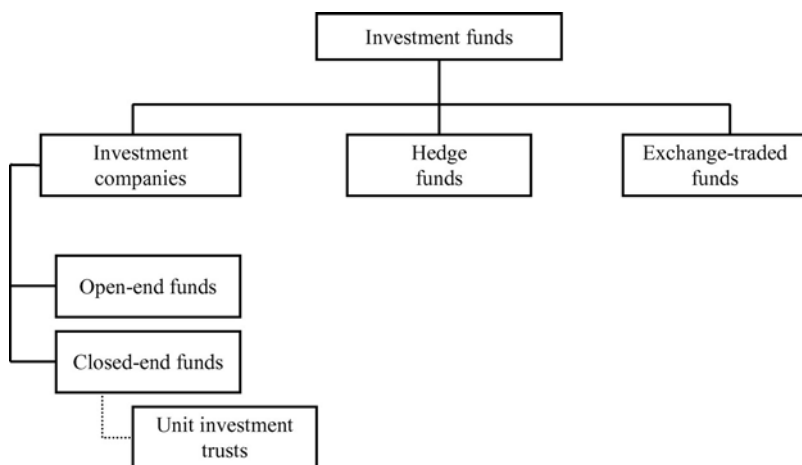


Figure 9.1 General investment fund classes

have become a significant component of the marketplace. Bond funds that invest in corporate and government securities have also emerged as an important part of the market. Money market funds, which debuted in the early 1970s, have become extremely popular with investors seeking relatively low risk/predictable return assets, a liquid conduit that generates income prior to reallocation to another sector, as well as those that expect short-term interest rates to rise.²

Certain closed-end funds date back more than a century. Despite such early beginnings, the sector went through a period of extreme losses during the stock market crash of 1929, which curtailed investor interest for many years. Indeed, the funds fell out of favor until the 1950s, when they began to re-emerge very gradually. Though closed-end funds benefit from some of the same drivers as open-end funds (e.g. diversification, efficiency), growth in the sector has been rather more measured – largely because they are structured more rigidly, as we shall discuss below.

Hedge funds are a more recent creation, having been developed in the 1940s and 1950s,³ but gaining true popularity in the 1970s (when more than 150 funds were formed). Early hedge funds were based on commodities and equities, and fund managers preserved a fairly strict single-asset focus. However, during the early 1980s, as inflation and interest rates rose sharply, hedge fund managers began investing across asset classes in order to offset the impact of poor equity returns. The multi-asset focus remains a key feature of many funds to the present time. Hedge funds have the ability to trade across a wide range of assets and strategies, and can freely employ derivatives, leverage, and short selling; this flexibility has attracted the attention of investors seeking larger returns than might be achieved through indexing or passive investing strategies.⁴ Naturally, the fact that hedge funds offer the possibility of generating above-market returns means that they often assume more risk than their more conservative open- and closed-end counterparts. Hedge funds can often charge significant fees: while all investment funds

² Indeed, money market funds investing in discount instruments maturing in less than 90 or 180 days (e.g. CP, CDs, BAs, government bills and similar short-dated instruments) gained such a rapid and widespread following in the decade following their introduction that many banking institutions were forced to create competing products (e.g. money market deposit accounts) in order not to lose an excessive amount of customer deposits.

³ For instance, in the US, A.W. Jones & Co. established a de facto hedge fund investment partnership in the 1940s that remained active until the 1970s.

⁴ By 2005, the number of outstanding hedge funds equaled the number of open- and closed-end funds for the first time; however, on an assets under management basis, the latter still outweighs the former by a significant margin.

apply front- or back-end fees to cover their expenses and management charges, hedge funds do the same and add on sometimes substantial performance-related “success” fees.

ETFs, like hedge funds, are a relatively recent addition to the financial markets. The earliest incarnations of ETFs, structured as index certificates/participation units, appeared on the American Stock Exchange (ASE) and Philadelphia Stock Exchange (PHLX) in the late 1980s, but drew only modest attention and activity. In fact, the original offerings by the ASE and PHLX were eventually abandoned after the Chicago futures exchanges and the Commodity Futures Trading Commission filed lawsuits claiming the instruments were equivalent to cash-settled futures (without any underlying securities). The breakthrough ETF, backed by securities and issued in the form of a trust in order to avoid any regulatory disputes, was the ASE’s 1993 S&P Depository Receipt (ticker SPDR); within five years, the size of the SPDR market had grown to more than \$12 b, proving investor appetite was real.⁵ The creation of other products soon followed, including ETFs on the NASDAQ 100 (e.g. QQQ/QQQQ).⁶ The Toronto Stock Exchange (TSE) launched its own physically settled TSE 35 index participation units in 1990, also gaining market share.⁷ As the advantages and mechanics of the structure became more widely understood, various other exchange and bank sponsors began creating their own ETFs, either on broad market indexes or on narrower country/industry sectors.⁸ The ETF market is not confined to the US. Apart from the pioneering work of the TSE, other global stock exchanges, including those in Hong Kong, Singapore, Tokyo, London, Frankfurt, and Sydney, have introduced their own ETFs.

Several factors have led to strong growth in the structured investment product sector over the past few decades. In particular, investment funds, hedge funds, and ETFs:

- provide investors with exposure to a professionally managed portfolio of securities, with risk and return characteristics synchronized to accommodate a broad range of investor styles and goals – those preferring more or less risk can identify a suitable investment vehicle easily;
- take advantage of portfolio management techniques to attempt to maximize returns for a given level of risk, in a manner consistent with financial theory;
- create transactional efficiencies by allowing execution of multiple transactions (e.g. the purchase of an index fund or basket) through a single trade – this generates portfolio management and tracking efficiencies, as well as cost savings;
- generate continuity in an investment portfolio, particularly for assets with final maturities or redemption dates – this allows investors to continue deploying their capital in a preferred strategy without being concerned about reinvestment;
- allow investors to self-direct their own asset allocations or trade their own fund shares – this has proven beneficial in an environment where investment/retirement assets are increasing and technology has become pervasive.

⁵ In fact, ETFs feature characteristics from both the index futures and the closed-end fund markets: like index futures, ETFs are quite standardized, liquid, and fungible, and trade throughout the day on an exchange; like closed-end funds, they feature a computable NAV that is determined through buying and selling forces.

⁶ For instance, the ASE’s NASDAQ 100 (ticker QQQ, or “cubes”) ETF, which has since been transferred to the NASDAQ (as QQQQ), is one of the most liquid securities in the entire financial marketplace, attracting hedgers, speculators, and investors from around the world to create an exceptionally deep market.

⁷ The TSE was a pioneer in ETF product development. Each share in the benchmark TIPS 35 ETF represents an interest in a trust that holds the basket comprising the TSE 35. The trust is permitted to lend its securities to generate fees that can be used to offset certain costs. The TIPS basket automatically adjusts if the index adds or deletes companies. The success of the TIPS 35 ETF led to the launch of the TIPS 100 in 1996.

⁸ For instance, in the mid-1990s, Deutsche Bank developed proprietary country index baskets and achieved a modicum of success for a period of time. However, the clear institutional leader has been Barclays, which in 1996 created its World Equity Basket Security (WEBS) product using Morgan Stanley’s well-established MSCI country indexes; WEBS, which were eventually rebranded “i-shares,” have become synonymous with active ETF product development and trading.

As we have noted, investment funds are suitable for a very broad range of investors. The lowest risk, index-tracking funds that seek to replicate, as closely as possible, market performance, are often considered to be desirable for those with a low risk-tolerance level. Highly leveraged and concentrated hedge funds exist at the other end of the spectrum, and are suitable for investors with a high level of risk tolerance and a desire to maximize potential returns. In between lies a range of open and closed-end funds and ETFs that cater to various other risk/return tradeoffs.

9.3 PRODUCT MECHANICS AND APPLICATIONS

Though the broad classes of investment fund outlined above transform investor capital into shares/partnership interests representing ownership in a portfolio of assets, product structure and mechanics can vary considerably. As we analyze the general structure of each fund class in the sections below, we highlight the specific product features that create these distinctions – and which dictate why certain fund structures are of greater or lesser appeal to particular groups of investors.

9.3.1 Structure, diversification, and management

The theoretical balance sheet of an investment fund vehicle, which may be structured as an investment company, SPE, trust, or limited partnership, can be viewed as a portfolio of securities on the asset side (along with a residual cash balance reflecting funds awaiting investment/liquidity to meet redemptions), and investors' paid-in capital in the equity account. Depending on the nature of the fund, the balance sheet may also contain a certain amount of leverage in the form of short-term or long-term debt, and may also feature off-balance-sheet assets or liabilities (primarily derivatives). The general structure of a fund's balance sheet is illustrated in Figure 9.2.

Diversification of risk across an entire portfolio of securities is a well-established mechanism for maximizing expected return for a given level of risk, or minimizing risk for a given target return. The Capital Asset Pricing Model (CAPM), developed by Markowitz, along with empirically tested versions of CAPM, such as Roll and Ross's Arbitrage Pricing Theory, demonstrate that it is possible to create a portfolio that maximizes return for a given level of risk (i.e. standard deviation of returns) by incorporating securities that are uncorrelated or negatively correlated. An efficient portfolio is one that generates high returns for a given level of risk, while an inefficient portfolio is one that generates lower returns for the same level of risk, as noted in Figure 9.3.

The inclusion of properly selected assets with appropriate correlation is thus central to creating an efficient, diversified portion. It is important to note that while diversification is widely recognized as a favorable trait, a fund may reach a point where excessive diversification is no longer beneficial; this tends to appear with managed funds that seek to outperform an index or benchmark, where an extreme amount of diversification may cause the portfolio manager(s) to be insufficiently "expert" truly to comprehend the individual components of the portfolio. In fact, it is hard for a portfolio manager to be expert in every country, sector, or security; some degree of specialization is ultimately necessary. Investors must be aware of the particular strengths of a portfolio manager and the underlying fund, to ensure that the philosophy and style are consistent with expertise and expectations.

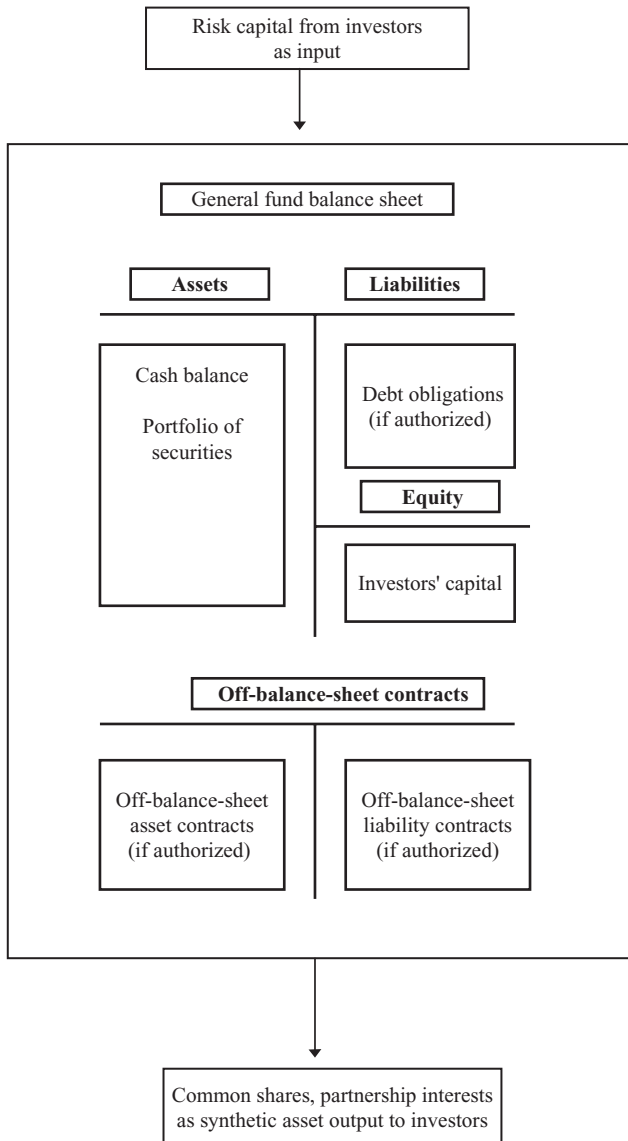


Figure 9.2 General fund balance sheet

Management of a portfolio of securities also gives rise to continuity, particularly with regard to assets that otherwise feature finite maturities; this means that fund managers can create investment horizons synthetically that are not achievable through the purchase of individual securities. For instance, an investor that is interested in continuous exposure to the bond market can purchase individual bonds, but must then reinvest principal continually as individual bonds mature. Continuous investment in a portfolio of bonds, in contrast, requires no such reinvestment by the investor; the portfolio manager simply acquires new bonds to replace




Figure 9.3 Risk, return, and portfolios

mature ones, all the while generating a steady stream of income for the investor.⁹ Similar techniques can be applied to other types of assets with specific final maturities (e.g. managed futures accounts/commodity pools, money market funds).

Professional management is significant for many retail investors that may lack the time, expertise, or resources to identify, track, and manage a group of diversified investments. Indeed, selecting and managing a portfolio of assets (including portfolios that are designed simply to track a broad index) can be a complicated task, and employing the services of a team of professionals can be of considerable benefit. However, the selection of a portfolio manager does not guarantee good investment performance, and comes at a price (annual fees and/or profit sharing fees are deducted from overall returns). Even seasoned professionals can perform poorly for one or more periods, leading to deterioration in investment returns. In addition, investors cannot control their investment capital directly once they purchase shares in a fund; all discretion lies with the fund manager. While investors obviously can target general markets or sectors of interest, and can indicate their desire for current income or capital appreciation, they cannot select specifically the securities that ultimately comprise the portfolio. The purchase of a position in a fund is an efficient transaction: a fund that tracks the S&P 500, for instance, requires the investor to purchase a single structured asset (e.g. a share in the mutual fund), rather than a weighted portion of 500 stocks. This efficiency entails a cost, of course: in addition to fund management fees, funds may also charge fees to cover various types of expenses, which again detract from overall returns.

9.3.2 Open-end funds

Open-end funds are well-established corporate vehicles that adhere to basic standards related to legal/regulatory structure, share creation, trading, redemption, operating costs, and tax treatment; the end goal is to use the structure to meet certain investment-related objectives. Before considering the specific structural characteristics of open-end funds, we examine the general types of funds that investors can buy/sell in order to meet their objectives (note that these same

⁹ Of course, certain bond funds are created with specific maturity dates. These “target maturity funds” make fixed principal payments to investors on particular dates through the use of zero coupon bonds, and thus synthetically replicate the payment profile of a portfolio of dated securities. Most bond funds, however, continue to reinvest principal in new securities, de facto converting dated assets into perpetual ones (until the investors liquidate).

general fund categories may be equally applicable to the closed-end funds and hedge funds discussed later in the chapter).

General fund categories

Open-end funds, which represent the single largest component of the global investment fund market, are designed to meet a broad range of investor goals related to earnings, geography, asset class, and subasset class.

- **Earnings.** The earnings component relates to the specific source of a fund's income, which may be derived from dividends and/or interest (so-called current income, which also has capital preservation features), capital gains, or both.
- **Geography.** The geographic component relates to a fund's specific market focus, and may be national (e.g. confined to a specific country), regional (e.g. based on a related group of national markets), or international. Those with a regional or international component may expose investors to currency risks.
- **Asset class.** The asset class parameter relates to the general markets in which a fund invests in order to meet its investment targets. These, broadly speaking, may include equity, fixed income, currency, commodity, and alternatives.¹⁰ Many funds limit their investments to a single asset class (though most funds also hold a certain amount of cash equivalents to meet redemptions). Equity funds represent the single largest component of the investment fund sector, with most major fund companies offering investors a range of investment opportunities across sectors, countries, and broad indexes; funds may be structured to provide investors with high current dividend yields and/or capital gains. Equity funds generally feature a considerable amount of diversification, with large funds holding positions in dozens or hundreds of individual issues in order to boost returns for a given level of risk. Bond funds focus on investment in domestic/foreign government, municipal, or corporate securities, or some mix of the three.¹¹ Unlike equity funds, bond funds generally require less diversification, as most fixed-income securities react to changes in interest rates in the same manner (the exception being those that are exposed to prepayments or amortization); however, those that focus strictly on credit spread products may require more diversification.¹² Balanced funds incorporate assets from different markets, including stocks, bonds, preferred stock, convertibles, and so forth (e.g. an asset allocation fund may feature a dynamic mix of investments and instruments, such as a 50 % allocation in equities, 20 % in bonds, 5 % in cash, 25 % in commodities).¹³

¹⁰ The alternative sector can be considered something of a "catch all," and may include any nontraditional investments, such as catastrophe bonds, inflation- or real-estate-related structures. It also includes funds that invest in other funds, i.e. the growing "funds of funds" market, where fund companies allocate their capital to the purchase of affiliated or nonaffiliated funds; these funds can be considered the equivalent of asset allocation funds that alter their composition through the purchase of other funds, rather than underlying equity or fixed-income securities.

¹¹ Bond funds may specialize in taxable securities (e.g. corporate bonds, high-yield bonds, MBS) or nontaxable securities (e.g. municipal bonds), and they may be constructed to provide specific exposure to credit quality or duration. Each specialty bond fund tends to feature distinguishing characteristics and risks: corporate bond funds, for instance, may focus on high-grade bonds, high-yield bonds, or emerging market bonds, each providing a different level of exposure to credit spread/volatility; MBS funds, constructed with a range of MBS (such as those we discussed in Chapter 4), contain many of the same prepayment risks that characterize the underlying securities; high-quality government bond funds may be required to invest a minimum amount of their assets in government securities, in order not to expose investors to credit default risk; and so forth.

¹² Bond funds generally are more transparent than other classes of fund, and tend to feature more predictable monthly inflows of interest-based dividends. They also facilitate retail participation in the bond markets by allowing small denomination purchases (versus institutional size orders of physical bonds).

¹³ Though certain balanced funds may appear attractive, investors must examine cost structure to determine whether better results can be achieved by assembling the individual components directly through other specialized funds.

- Subasset class. For funds that focus on a specific asset class, a subasset class distinction may also be relevant. Thus, a stock fund may invest in large capitalization stocks (e.g. \$5 b+ in market capitalization), mid-cap stocks (up to \$5 b), small-cap stocks (up to \$1 b), or micro-cap stocks (up to \$500 m), or some mix of the four. Similarly, an index fund may replicate a specific market/sector, such as the S&P 500, Russell 2000, Nikkei 225, or FTSE 100.¹⁴ Sector funds attempt to capitalize on potential gains from a single industrial or market sector (e.g. real estate investment trusts, pharmaceutical stocks, technology stocks, emerging market stocks); since their focus is much narrower, security concentrations can be much more significant (e.g. some funds are permitted to invest up to 25 % of their assets in a single stock or bond). Similar subasset class sector distinctions can be created with fixed-income funds (e.g. a focus on municipal securities, government securities, high-grade corporate bonds, money market securities, and so forth), currency funds (e.g. G10 currencies, emerging market currencies, and so on), and commodity/alternative funds (e.g. precious metals – physical or via the stocks of gold/silver producers, energy – physical or via the stocks of producers/refiners/marketers, catastrophe/weather, and so on).

The largest global fund companies offer investors choices in a wide variety of funds from different asset classes and subsectors. This provides investors with flexibility as well as lower fees/costs (i.e. fund companies will often provide a preferential fee structure for investors allocating assets within the same family of funds).¹⁵ The broad investment parameters are summarized in Figure 9.4.

All funds, regardless of defining characteristics, contain one or more risks that can impact overall investment performance/returns. Common risks that investors must consider include default risk, sovereign risk, credit spread risk, currency risk, interest rate risk, equity price risk, dividend risk, volatility risk, liquidity risk, reinvestment risk, basis risk, and prepayment risk. For instance, a fund that invests in a portfolio of large-cap equities exposes investors to equity price, volatility, and dividend risk; one that invests in high-yield bonds exposes investors to default, credit spread, interest rate, volatility, and liquidity risk; one that invests in emerging market bonds creates default, credit spread, interest rate, volatility, liquidity, sovereign, and currency risks; and so on. Awareness of such potential risks is essential prior to committing capital.

Legal/regulatory structure

Open-end funds that are intended for public distribution must adhere generally to specific legal and regulatory parameters. In many countries, this means that fund companies must register with the relevant national securities regulator and comply with minimum standards

¹⁴ Index funds attempt to mimic, but not replicate precisely, a particular benchmark, by optimizing their portfolios with constituent securities. Performance of an index versus the benchmark is thus impacted by differences in portfolio holdings, as well as the level of cash held to meet redemptions and overall expenses. The first US index funds debuted in 1976, and overall growth in the sector has been dramatic as a result of Markowitz's modern portfolio theory contributions and growing belief in market efficiency (i.e. market efficiency indicates that investors balance rationally between risk/return, there are no freely available risk arbitrage opportunities, and the past is not a good predictor of the future). Indeed, empirical research has provided greater support for the medium-strong form of the efficient market hypothesis (suggesting that the market incorporates all publicly available information, making it difficult for portfolio managers to outperform the market). Well-constructed index funds also feature lower expense ratios, and remove the specter of portfolio manager underperformance. Investors seeking to allocate capital to an index fund must be aware of the fund's replication strategy and tracking error to understand potential deviation from the benchmark. Some studies suggest that approximately 20 % of global equity funds utilize some form of indexing strategy.

¹⁵ Though funds are purchased generally in unleveraged form, some accounts allow investors to purchase shares on margin. In addition, special accounts can be established to short sell a fund position; this, however, can be a risky transaction, since funds only trade once per day at the closing NAV (rather than instantaneously, as in the case of ETFs).

Earnings source Current income/capital preservation Capital gains Both	Geography National/single country Regional International
Asset class Equity Fixed-income Currency Commodity/alternative Balanced	Subasset class Market type Index Sector

Figure 9.4 Broad investment parameters

of disclosure, diversification, reporting, dividend policy, and income distribution. Depending on the specific jurisdiction, regulatory authorities may monitor a fund and its activities in a formal or informal manner to ensure compliance; further oversight may be provided by the listing exchange. For instance, in the US, investment companies operating as open-end funds are regulated via the Investment Company Act of 1940 and the Investment Advisors Act of 1940.¹⁶ Any registered investment company is required to file annual reports with the SEC and maintain specified accounts and records; in addition, all customer assets must be held in custody. However, as a result of the additional oversight provided via registration, the funds are permitted to make public offerings of shares to retail and institutional investors, and are allowed to market their services and products to the public at large.

Though an open-end fund can be organized as a corporation, trust, or partnership, in practice most select the corporate structure. A fund is managed typically by a sponsoring investment management company, which is responsible for organizing the fund and making administrative/operating and investment decisions. Each fund is governed by a specific investment mandate, which delineates the types of investments that may be undertaken, the specific classes of instruments/strategies that can be employed to fulfill investment goals, and the kinds of risk that can be taken. The portfolio manager responsible for operating a fund on behalf of the investment management company must adhere strictly to the mandate. For instance, some funds are only permitted to invest in cash-based securities, while others can use derivatives either to hedge or speculate; some are permitted to use a small degree of leverage (i.e. through general bank borrowing¹⁷), while others cannot;¹⁸ some must hold a minimum amount of cash for redemptions, while others are granted greater discretion; and so forth. Details on a fund's investment mandate and other items of interest are conveyed typically through a prospectus that must be distributed by the underwriter (or fund sales distributor) to each investor. Though the form of a prospectus varies across national systems, coverage often includes a fund's goals,

¹⁶ The original Investment Company Act of 1940 was updated significantly in 1970, and has been refined on several occasions since that time. The Investment Company Act of 1940 covers all companies that invest, reinvest, or trade as a primary form of business, as well as those that are engaged in those businesses and have 40%+ of their assets in investable securities.

¹⁷ In general, an open-end fund cannot issue any senior security, apart from notes evidencing bank borrowing.

¹⁸ The use of leverage and short selling in a fund may be restricted by formal asset/liability tests, e.g. liabilities must be less than one-third of assets (where liabilities are defined to include borrowed money and securities), or gross long and gross short positions must be less than 150% of net assets.

strategy, risk factors, sales charges, operating expenses, financial history, annual return history, purchase/redemption mechanisms, portfolio manager (fund advisor) experience and fees, and shareholder distribution mechanisms.

An individual fund forming part of a broader family group of funds may “outsource” certain control and infrastructure operations to the investment management company in order to avoid duplication and reduce costs; such outsourcing arrangements generally must adhere to strict rules in order to avoid any conflict of interest. However, each individual fund generally must have its own independent board of directors to consider and approve all fund-specific matters – just as any other publicly listed company might. Investors in an open-end fund are allowed to vote on matters related to the fund, but not on those related to any underlying equity securities the fund may hold in its portfolio; the fund retains the proxies and can thus vote in support of, or in opposition to, particular issues.

An independent custodian, such as a bank or trust company, must be nominated by the investment management company to hold a fund’s assets and monitor its cash inflows/outflows. A transfer agent, in turn, is appointed to track share purchases and sales, maintain shareholder records, compute the daily NAV, and arrange for dividend/interest payments and capital gains disbursements. Many of the largest fund companies also feature fund distributors (underwriters), which are affiliates of the investment management company responsible for marketing fund shares to investors (directly or through brokerage firms). The legal vehicles and participants involved in open-end fund creation and management are summarized in Figure 9.5.

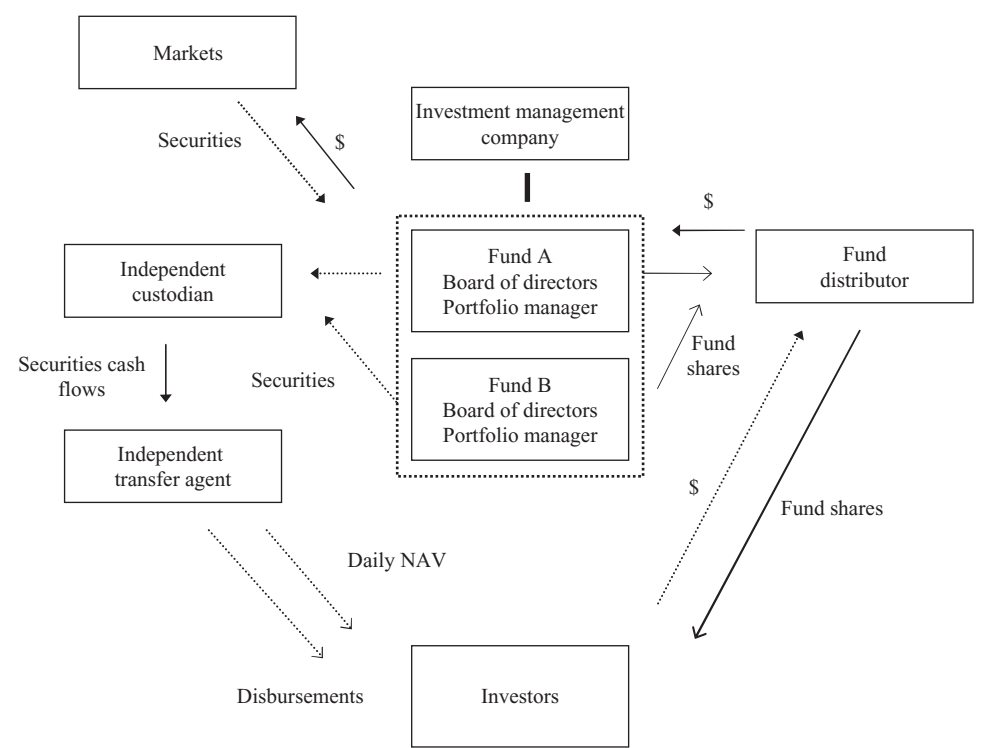


Figure 9.5 Open-end fund legal vehicles/participants

Share creation, trading, and redemption

An investor wishing to buy shares in an open-end fund transfers a specific minimum amount of cash to the fund distributor, which forwards the cash to the fund; the fund then buys the underlying assets that form part of the investment strategy. Since the fund is structured as an open-end vehicle, it can create new shares at the NAV without limitation (subject to the exceptional circumstance noted below). An investor selling shares submits appropriate redemption instructions to the fund, which uses its cash reserve or liquidates a pro rata amount of securities to meet the redemption.

Open-end funds react to order flow, and order flow generates cash flow. A fund's cash flow is defined as the net amount of new money invested in a fund, or the excess of purchases over redemptions. Funds generally preserve a cash buffer of 5–10 % of assets in order to meet redemption orders; in fact, fund managers attempt to calibrate the size of the cash buffer to meet expected redemptions, as an excess amount of cash reduces fund earnings. While most successful funds face a greater amount of purchases than redemptions, there are times when net redemptions exceed the cash buffer, meaning portfolio assets have to be liquidated. In the most extreme situations, continuous net redemptions can become problematic, causing a liquidity squeeze and forced/distressed sales of assets.

Every publicly traded fund is valued at the end-of-day NAV, which is computed as:

$$\begin{aligned} &\text{Cash and equivalents} \\ &\quad + \text{market value of securities held in the portfolio} \\ &\quad - \text{current liabilities (including accruals)} \\ &= \text{total net assets} \\ \\ &\text{Total net assets/shares outstanding} = \text{NAV per share} \end{aligned}$$

Given these equations, it is clear that NAV must fluctuate on a continuous basis: the composition of the market portfolio and current liabilities can change, and the market value of the underlying portfolio will change. Open-end funds may trade at a discount, and occasionally a premium, to NAV; unlike the ETF structure described below, there is no true arbitrage mechanism that can efficiently take advantage of potential discounts/premiums to force the funds to trade at, or close to, NAV (i.e. there is no facility for exchanging shares of the underlying portfolio for fund trust shares, and vice versa). Although open-end funds generally accept new funds from investors, some eventually accumulate an excess of assets and may close to further subscriptions, either temporarily or permanently. This is a defensive mechanism that helps ensure the portfolio managers do not suffer from “style drift” – investment outside of a given area of expertise that can result in a misbalancing of risk/return.

Operating costs and fees

Open-end funds generally charge investors one or more participation fees; though these vary by market, the general approach tends to be consistent across markets. Using the US market as a guide, a given fund may feature one or more of the following investor fees:

- A front-end load fee, payable by investors when they initially purchase shares in a fund. These may range from 3 % to 6 %, though certain discounts are often applied once certain volume thresholds (known as breakpoints) are reached.

- A 12(b)-1 fee, payable by investors on an annual basis to cover a fund's distribution, advertising, and marketing costs. These fees can range from approximately 0.25 % to 1 %.
- A back-end load fee, or a contingent deferral sales charge, payable by investors when they exit the fund. In some cases these fees decline with the holding period, e.g. 5 % exit fee if shares are sold during the first year, declining by 1 % per year thereafter. The intent is to urge investors to hold their fund shares for the long term.

Certain funds feature one or two of these charges, while others feature all of them, meaning investors contemplating a purchase must investigate thoroughly the fee structure disclosed in the fund prospectus. In the US market, funds tend to be categorized into A, B, and C classes, where A shares have front-end load fees and small 12(b)-1 fees but no back-end load fees, B shares have no front-end load fees but feature higher 12(b)-1 and back-end load fees, and C shares have no front- or back-end load fees, but high 12(b)-1 fees. Investors buying funds through a broker may also have to pay commissions; those buying directly from the investment management company's fund distributor can often avoid such commissions.

The fees summarized above are used to cover a fund's expenses, which include administration, marketing, transaction execution, reporting, and management. Investors evaluating a fund's expenses can examine historical performance as a benchmark.¹⁹ This can be done by computing a fund's total return over time, which is simply the sum of capital gains distributions, dividend/interest income distributions, and changes in the NAV, net of expenses. It can also be done by comparing total expenses to net assets under management, which is commonly referred to as the expense ratio. A review of outstanding funds suggests that expense ratios range between 0.5 % and 3 % per annum; those achieving ratios below 1 % are considered very efficient, and may represent the best value for investors attempting to maximize returns. Certain observations can be made regarding fund expenses: large funds generally have lower expense ratios than their smaller counterparts, as they can take advantage of economies of scale and obtain tighter bid-offer spreads on portfolio transaction orders; larger account sizes enjoy a smaller expense ratio for the same reasons; index funds have lower expense ratios, as they seek simply to replicate a benchmark; international funds have higher expense ratios as a result of global research and due diligence costs; and, bond funds have lower expense ratios than equity funds, as their need for detailed research and investment in many different securities typically is more limited, and turnover is lower.

Portfolio turnover can have a significant impact on expense levels: the higher the turnover, the greater the transaction costs, and the higher the overall expense ratio. Accordingly, an examination of fund turnover is useful. Portfolio turnover rate is the most common measure; for instance, 100 % turnover indicates that a security is held, on average, for one year, 50 % turnover reflects a two-year holding period, 200 % turnover indicates a six-month holding period, and so forth. Funds tend to average turnover rates of 40 % to 100 %; those with rates of 200 %+ are likely to generate greater expenses, and may detract unnecessarily from overall performance.

Tax features

Open-end funds can distribute income/earnings to investors as ordinary income (e.g. dividends/interest), capital gains, and nontaxable distributions (e.g. dividends paid out of principal);

¹⁹ In some jurisdictions, such as the US and UK, regulators require fund companies to disclose their fee structures in the prospectus/marketing circular. This may include details on shareholder transaction expenses, operating expenses, and so forth. The requirement is not, however, uniform across countries, meaning some investors are forced to rely on incomplete information when making their investment selections.

depending on national tax laws, each of these distributions can attract different tax treatment. For example, in the US, long-term capital gains generated by a fund (generally those covering a period of at least 18 months) may be taxable at the preferential capital gains rate rather than the higher ordinary income rate applied to short-term capital gains and other current cash inflows. In certain markets, such as the US and UK, a minimum amount of a fund's income must be paid to investors in order to ensure proper tax treatment. For instance, in the US, 98 % of net investment income must be paid to investors (where net investment income is defined as dividends plus interest, less expenses); in addition, 98 % of net realized capital gains must be distributed (when capital gains are distributed, a fund's NAV falls by some per share amount, reflecting the distribution). Equity-based funds generally distribute earnings to investors on a quarterly or annual basis, while fixed-income funds do so on a monthly basis. International funds that generate income for nonresident investors may be subject to foreign tax withholding that reduces dividends/returns, meaning investors may then have to use a foreign tax credit to avoid double taxation.

Fund investors are liable for taxes on dividends/interest generated by the fund, and for any capital gains that are triggered when share redemption by any fund investor leads to net asset liquidation (or when the portfolio manager voluntarily liquidates a position that generates a gain). Investors have little or no control over the timing of any gains. Capital gains distributions generally are declared prior to the end of the fiscal year, and are distributed in the early part of the following year, but are taxable to investors in the year they are declared. Investors that retain their shares when assets are liquidated may be impacted by an unexpected tax liability. Conventional open-end funds may choose to redeem the highest cost basis assets in the fund first, in order to generate the smallest amount of taxable gain; since the lowest cost basis securities remain in the portfolio, the unrealized tax liability grows larger. Active funds (e.g. portfolio turnover of 100 %+) may generate significant tax liabilities versus index funds that only turn over through rebalancing activities; similarly, funds that pay out smaller dividends and/or have a steady inflow of cash may be more tax efficient than those paying out large dividends or those that suffer from unstable cash flows. Investors choosing between competing fund alternatives must consider these characteristics properly.

9.3.3 Closed-end funds

Closed-end funds represent a much smaller segment of the investment fund sector. Like their open-end counterparts, closed-end funds are professionally managed portfolios of securities that are listed and traded on exchanges, and tend to feature only a limited amount of leverage.²⁰ Unlike open-end funds, however, closed-end funds have only a limited number of outstanding shares, and cannot create new ones at will. In addition, they feature two separate prices: the quoted NAV, and a distinct market price. The closed-end NAV is the same as the NAV of an open-end fund, and is computed in the same manner. The share price, which is quoted separately, can often diverge from the NAV, leading to the premium and discount scenarios described below. As a result of their unique structural and trading characteristics, closed-end funds generally are considered to be suitable for investors with a medium- to long-term investment horizon. In the US, the most active segment of the closed-end market centers on municipal bonds; offshore, many closed-end funds provide investment opportunities in the equities of companies located in separate countries (so-called "country funds").

²⁰ A closed-end fund generally cannot issue more than one class of debt securities and not more than one class of preferred stock; before any debt or preferred stock issuance, the fund must adhere to minimum asset coverage levels.

The general class of closed-end funds also includes the unit investment trust (UIT), which can be considered a passive fund that accepts an initial (and limited) amount of investor capital contributions, purchases a target portfolio of securities, and holds the securities until the maturity of the fund. The portfolio of a UIT is not rebalanced dynamically; in fact, the trustee simply holds the assets that have been purchased, and collects and distributes cash flows. In practice, most UITs are invested in fixed-income, rather than equity, securities, with a finite life (i.e. a defined maturity/redemption date). Once the final securities mature, the UIT is liquidated and investors receive their principal. Since UITs are passive portfolios, transaction costs are negligible once the initial portfolio is set, and ongoing fees generally are lower than on conventional closed- or open-end funds.

Share creation, trading, and redemption

Closed-end funds do not issue new shares on a regular basis (hence the name “closed-end”). Indeed, new share issues of closed-end funds are extremely rare, primarily as a result of significant fees that investors must bear. As a result, funds nearly always arrange for a single capital issuance and then preserve a relatively constant level of capital. There are two exceptions to the constant capitalization of closed-end funds: rights issues and leverage. In certain cases, closed-end funds may opt to increase capital through a rights issue placed directly with existing investors,²¹ and/or they may borrow against fund assets to acquire additional assets (leveraging of this type tends to occur primarily in bond, rather than stock, funds). Closed-end funds do not redeem outstanding shares; all shares issued remain outstanding until (or if) a fund is liquidated or restructured.²² Investors seeking to exit the fund must therefore sell their shares to other investors seeking to buy them, and do so through a standard brokered transaction; this typically means secondary liquidity is much lower, and is a key reason why closed-end funds often trade at a discount to NAV. While the constant level of capital may, in some respects, be seen as a limiting factor, it can also be regarded as an advantage: fund managers need not be concerned about continuous cash inflows and reinvestment, and are unlikely to suffer from style drift. In addition, they do not need to hold a significant cash buffer (e.g. 5–10 %) of assets, since they need not worry about redemption; this places less pressure on potential returns.

As noted, all activity in closed-end funds is conducted through brokers/dealers, rather than investment management companies and fund distributors, meaning investors buy and sell at a fund’s bid-offer share price, rather than at the close of business NAV. In fact, there is no guarantee an investor will be able to buy or sell shares in a fund at any particular time without crossing a potentially wide bid-offer spread. The share price of a closed-end fund fluctuates throughout the trading day, while its NAV is established only at the close of business: if the share price is lower than the NAV, the fund is trading at a discount, and if it is higher than the NAV, it is trading at a premium (the actual premium or discount can be computed via $[\text{share price} - \text{NAV}]/\text{NAV}$). Closed-end funds are particularly prone to discount valuations for reasons that may include the illiquidity just mentioned, poor management performance

²¹ Rights issues typically are distributed at the rate of one new share for each existing share, with existing investors able to purchase new shares at a 5–10 % discount to the quoted share price. Rights may be transferable or nontransferable; investors not wanting to take up transferable rights may sell them to third parties, while those not wanting to take up nontransferable rights effectively limit the ability for the fund to raise additional capital.

²² Closed-end funds ultimately may continue operating in closed-end form, or they may merge with other funds, sell off all of their assets and wind down, or convert into open-end form.

against a benchmark, wide bid-offer spreads associated with purchasing risky assets, high expense levels, and/or unrealized appreciation (leading to tax liabilities). Certain investors employ strategies that are designed to take advantage of these discounts, buying when share prices fall well below theoretical NAVs. Trading at a premium is an extremely rare event, and may only appear in special circumstances (e.g. when a closed-end fund focused on an “in-favor” single country draws the attention of many investors eager to allocate capital at any price).

Other aspects of closed-end funds, including legal/regulatory structure and tax issues, follow from the discussion on open-end funds above.

9.3.4 Hedge funds

The hedge fund sector has developed very rapidly over the past two decades, and has become the center of investor and regulatory interest since the turn of the millennium. Though US management and capital originally drove the market, the sector has become quite global, with hedge fund groups now established in countries such as the UK, Germany, France, and Spain, among others, and with investment capital coming from investors located around the world.

Hedge funds are flexible investment conduits, and are intended primarily for wealthy and sophisticated investors that are able to absorb greater risk than the average retail client. Unlike other classes of investment funds considered in this chapter, hedge funds attempt to generate absolute, rather than relative, returns. The primary goal of hedge fund managers is thus to create pure returns through active management, regardless of what is occurring in a specific market.²³ Hedge funds do not replicate indexes or engage in passive investment strategies, and may compound, rather than diversify, risk, by taking positions in positively correlated assets. They may also add to their risk (and potential returns) through leverage, concentrations, and short selling,²⁴ and by buying and selling illiquid assets and derivatives. Some portfolio managers employ systematic techniques to make investment decisions; this suggests a strong focus on technical indicators and systematic rules for entering or exiting a market. Other portfolio managers use discretionary techniques, which involve a combination of fundamental²⁵ and technical factors.

Contrary to popular belief, not all hedge funds are active traders, buying/selling positions repeatedly throughout the trading day; some feature relatively low portfolio turnover, as they seek to crystallize value over the medium term.

We have noted that hedge funds enjoy considerable flexibility in investing capital across markets and strategies. Though some funds are created specifically with a relatively narrow focus, most have the option to invest in a particular range of markets through general or proprietary strategies, on a leveraged or unleveraged basis. It is important to remember that the use of leverage increases the risk of a fund, but may/may not increase returns (depending on how the fund managers select and manage their portfolios). Hedge funds generally attempt to

²³ We may note a spectrum of management styles across investment funds: pure index tracking (target returns equal the index), enhanced index tracking (where target returns may be up to 1% away from a benchmark index), constrained active management (2–4%), unconstrained active management (4%+), and hedge funds (no index relationship).

²⁴ Short selling emerges as a significant advantage for the hedge fund sector. When traditional fund managers operating an open- or closed-end fund detect a poorly performing asset or sector, they can only avoid or underweight the contract; hedge fund managers can take the extra step by actually capitalizing on a negative move by shorting.

²⁵ Key fundamental variables may include economic/sector growth, inflation, interest rates, trade/capital flows, asset valuations, and so forth.

“lock in” the terms of their leverage to the greatest extent possible, so that they are not prone to margin calls and asset liquidation. Popular hedge fund strategies include:

- **Directional asset allocation:** directional positioning in one or more asset classes²⁶ (i.e. buying²⁷/selling/shorting²⁸ equities,²⁹ commodities, debt, and/or other assets, in order to take advantage of specific changes in market direction or volatility).
- **Relative value:** market neutral pairs trading,³⁰ bond arbitrage, or convertible arbitrage (i.e. buying/selling stocks, bonds, or convertibles to take advantage of discrepancies between a spread or basis, generally with little or no exposure to outright market direction³¹). Relative value fund managers try to remain hedged/balanced at most times, and tend to operate in only one or two markets at a time; the overall strategy, in simple terms, is to sell a rich/overvalued asset while buying a cheap/undervalued asset in the same class.
- **Event driven:** including risk arbitrage (i.e. buying/selling/shorting equities that are takeover targets or are “in play”),³² and distressed debt positioning (i.e. buying/selling debt that is deeply distressed or already in default, in order to capitalize on potential recapitalization or restructuring opportunities).³³

Some fund managers combine all of their strategies into a single partnership, while others prefer to suballocate them across individual partnerships so that investors can select from the available choices. Historical analysis suggests that relative value and event-driven strategies

²⁶ Global asset allocators that invest in multiple sectors often do so on the basis of macro themes; this means that positions may be medium to long term in nature.

²⁷ Some funds are pure momentum buyers (purchasing assets on strength), others are pure contrarian buyers (purchasing assets on weakness), and still others are opportunistic.

²⁸ Note that while virtually all hedge funds take short positions in particular assets from time to time, their net portfolios are still likely to be market neutral or net long. Some funds, however, only take short positions and are thus directionally bearish at all times.

²⁹ As noted in the section on open-end funds, some equity-based hedge funds focus on growth companies, others on value companies, and so forth. Unlike other funds, however, hedge funds can, and do, actively short stocks, which provides considerable flexibility.

³⁰ Though pairs trading is a popular relative value strategy, it can suffer from capacity constraints when many funds are focused on a single sector. Accordingly, some fund managers choose to apply their pairs trading algorithms across industry sectors. Those using fundamental, rather than technical, tools, may develop a style bias; in the equity markets this often means purchasing value stocks (undervalued) and selling growth stocks (overvalued). Pairs trading is also applied in the fixed-income markets, where fund managers attempt to capture a narrowing of the yield differential between two bonds. Bond spreads are, of course, driven by several risk factors, including inflation, credit defaults, prepayments, and liquidity, so the relative value algorithms are applied to each risk factor.

³¹ Fund managers attempt to purchase convertibles that are undervalued relative to the common stock, and then sell the underlying stock; this provides both a relative value and a volatility position. The resulting common stock hedge is managed dynamically on a delta-weighted basis (e.g. an at-the-money convertible that can be converted into 100 shares has a delta of approximately 0.5, and is thus hedged by shorting 50 shares of stock). In order for this strategy to be successful, the market must be volatile (i.e. if volatile on the upside, the fund gains on the convertible and loses on the short stock, and if volatile on the downside, it gains on the short stock and loses on the convertible); a quiet market limits trading opportunities on the short side, and the embedded call options in the convertibles erode in value. Many hedge funds employing the convertible arbitrage strategy choose to eliminate interest rate risk and credit risk associated with the convertible, in order to concentrate on equity volatility. Hedge funds have become dominant players in convertible arbitrage.

³² Risk arbitrage strategies generally are relatively high frequency/small-profit propositions that require specialized legal knowledge related to mergers and acquisitions and takeover rules. The risk arbitrageur seeks to take advantage of positive news on a takeover that has been announced by taking a long position in the takeover stock, or by going long or short the acquirer or target as an announced deal appears likely to run into difficulties. Unlike other classic arbitrage strategies, the risk arbitrage process can be viewed as a time-driven arbitrage: the stock of the target company typically trades towards the takeover price (premium) over time, but will not converge on it precisely, as the deal can collapse – even in its final days/hours. The intent is thus to capture as much of the premium over time as possible. The gains in risk arbitrage are relatively limited (i.e. the spread between the purchase price and the intended takeover price), while the losses are potentially large (i.e. a collapse in the target's share price if the deal fails to materialize); this risk profile is akin to that found in selling naked options. Takeovers that are structured as stock, rather than cash, deals feature an additional element of risk, as the acquirer's stock price can rise or fall over time, changing the value of the proposed deal.

³³ In distressed debt strategies, hedge funds attempt to invest in good companies that are plagued by bad capital structure. Funds may target high-yield bonds with very high debt service requirements that trade at significant discounts, hoping to identify those obligors that can continue to remain current over time so that discounts narrow and begin to move back towards par value. They may also trade in bonds/bank loans that have gone into default, attempting to profit from value generated via the reorganization process; this requires skills in identifying which class of debt will fare best under various reorganization and post-reorganization scenarios. Some funds investing in defaulted bonds/loans take an “activist” role on creditor committees in order to attempt to influence the eventual outcome; this, however, prevents them from trading in the defaulted obligations while negotiations are underway.

feature less risk than directional asset allocation strategies. It is worth noting that managers operating conventional open- and closed-end funds also use some of these strategies; however, any strategy that seeks to profit from small market discrepancies generally is unsuitable for funds that cannot use a significant amount of leverage.

Another important segment of the hedge fund market centers on funds of funds and funds of funds of funds (so-called F3s), which account for approximately one-third of all hedge fund asset investments. Following concepts of portfolio theory, these entities attempt to manage and diversify hedge fund performance risk by using investors' capital to invest in 5–20+ different hedge funds (or funds of funds). While the concept is appealing and has proven to be a sound strategy in some cases, funds of funds and F3s add an extra layer of costs that detract from investor returns (i.e. 1–3 % p.a. for each additional “management” layer). Thus, to be truly appealing to investors, fund of fund and F3 managers must have proven skills in selecting above average hedge fund performers.

Legal/regulatory structure

Though variations can appear within and across borders, many hedge funds are organized as private limited partnerships that are largely unregulated (though in some countries, certain minimum regulatory requirements must now be met, such as registering with the national securities regulator). Hedge funds historically have used regulatory exclusions and “safe harbors” to conduct their operations. This means that they are exempt from the same regulations applied to open- and closed-end funds; in fact, hedge fund managers are considered to be unregistered investment advisors, and cannot therefore advertise or market their products or services directly to the public. That said, hedge funds must still make available to prospective institutional investors an offering memorandum (rather than the more extensive prospectus that open- and closed-end funds must prepare) that contains essential details on the fund, its activities, and its risk factors. In recent years, some hedge funds have opted to register with their local regulators (e.g. the SEC in the US). Though this requires them to increase their level of disclosure, it gives them the ability to accept more investors and lower minimum investment amounts (though the “accredited investor” requirements remain the same).

Hedge funds are often structured as private partnerships rather than investment companies, using the standard general partner (GP)/limited partner (LP) framework. The fund management team acts as the GP, making all investment management decisions and retaining unlimited liability in the business (in practice, many GP interests are reincorporated in the form of limited liability companies (LLCs) in order to limit the liability of the GPs to capital invested). Investors in the fund serve as the LPs, holding passive interests in the fund and retaining only limited liability. Since hedge funds typically are formed as LPs, they can only have a small number of large investors (e.g. less than 100 limited partners)³⁴ and the placement of partnership interests cannot be done on a public basis. All investors in the LP must be accredited (which is taken to mean a firm or individual of substantial means, i.e. \$1 m in liquid net worth or \$5 m in assets).

In order to cater to both domestic and offshore investors, hedge funds often set up multiple vehicles that permit capital to be directed to the appropriate locale. For instance, a US-based hedge fund may establish an offshore partnership (e.g. in the Caymans, British Virgin

³⁴ In the US, offerings under Section 3 (c)(7) of the Securities Act can be directed at up to 500 investors, including individual investors with \$5 m in investable assets, and institutions with \$25 m in investable assets. The more common Section 3 (c)(1) offerings limit the number of investors to 100.

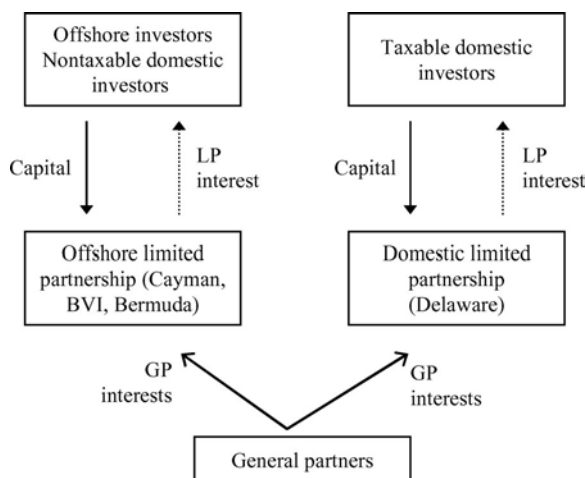


Figure 9.6 Multiple LP structure

Islands, or Bermuda) to accommodate the requirements of offshore investors and nontaxable US investors (e.g. pension funds, endowment funds, foundations³⁵), and an identical US partnership (e.g. Delaware) to service taxable US investors. Though multiple partnership vehicles may be employed, the capital is invested in the same strategies. A UK or German hedge fund may opt to follow a similar structure, selecting a tax-friendly locale for the establishment of an offshore partnership, and a domestic partnership for local investors. Figure 9.6 summarizes this process.

Share creation, trading, and redemption

The supply of shares in a given hedge fund is strictly limited. Partnership shares generally are arranged as private placements that are sold directly to accredited investors (in the US, this occurs under Regulation D as a “Regulation D offering”).

Since most hedge funds are structured as private partnerships, the level of liquidity accorded to investors in the fund (or prospective investors seeking to enter the fund) is minimal. In fact, partnerships may only allow for quarterly, semi-annual, or annual liquidity, rather than the daily liquidity characteristic of open- and closed-end funds, or the instantaneous liquidity of ETFs. Partnership share trading and redemption must therefore be viewed as virtually nonexistent, suggesting that direct investment in a hedge fund is a medium- to long-term investment decision. The fact that hedge funds need not worry about daily liquidity affords them greater flexibility in managing cash balances; rather than retaining a certain percentage of funds to cover potential redemptions (e.g. 5 % of assets), hedge funds can optimize their portfolios by holding very little cash.

Operating costs/fees

The income statement of a typical hedge fund may reflect interest income (from fixed-income instruments), interest expense (from borrowings and other forms of leverage), realized gains

³⁵ Indeed, nontaxable institutions have become the single most important clients within the hedge fund sector. Government pension funds, university endowments, and other nonprofit entities invest directly in hedge funds, or do so through fiduciaries.

and losses (from positions closed out during the fiscal period), and unrealized gains and losses (from positions still held on the books that are marked to the current market); these are supplemented by general operating costs related to the specific cost of doing business, including transaction execution, clearing/settlement, administration, and reporting. Since hedge funds tend to operate as very small partnerships, with a minimum of staff, many of these functions are outsourced to specialist intermediaries and prime brokers that are capable of handling the operational and infrastructural aspects of the business. Indeed, prime brokerage services offered by banks and investment banks to the hedge fund community, which include custody, settlement, reporting, and financing, have become an important revenue stream.³⁶

As we have noted, one of the distinguishing features of hedge funds is found in fee structure. Hedge funds command the highest fees within the general investment sector, meaning that fund performance generally must be much stronger than in other classes in order for investors to achieve the same after-tax returns. Hedge funds generally charge asset-based fees (just as any other fund might) and performance fees. Though the magnitude of fees varies across countries, fund types, and market/competitive conditions, a typical hedge fund may charge between 1–2 % p.a. on assets under management, and a 20 % performance fee on any gross earnings generated. Thus, the general partners of a hedge fund with \$100 m under management, that posts \$20 m in returns, will receive \$5 m (i.e. 2 % of \$100 m and 20 % of \$20 m). Importantly, the returns paid to limited partners come from the earnings that remain after all asset and performance fees have been paid (i.e. in this example, distributable earnings amount to \$15 m). Some hedge funds protect investors by including a minimum hurdle rate that must be exceeded before managers receive any performance bonuses. They may also include a “high water mark” test that must be passed before determining performance fees; the high water mark means that in order for fund managers to be compensated, they must continue to create value year over year. For instance, a fund with \$100 m under management and a 20 % payout that generates \$20 m of earnings in year 1, pays out \$4 m in performance fees to the managers. If markets are difficult in year 2 and the managers lose \$10 m, they receive no payout. If performance improves in year 3, and the funds generate \$10 m of earnings, the managers still receive no payout as a result of the high water mark: the fund has not yet exceeded the performance position held in year 1 (note that without the high water mark, the fund managers would be entitled to performance fees in year 3).

Tax features

Hedge funds, which are pass-through vehicles, may or may not be tax efficient; much ultimately depends on the strategies employed to generate earnings. In fact, most funds are not created to generate tax efficiencies, but to maximize gross returns; this is often accomplished through active turnover, short-term holding periods and/or long-term retention periods (i.e. many hedge funds choose to retain earnings over multiple years, rather than making distributions to limited partners³⁷). If a fund performs in such a tax-inefficient manner, then the optimal investors are those that are tax-exempt or nontaxable. In order to be considered tax efficient, a hedge fund

³⁶ Open- and closed-end funds tend not to use prime brokerage services as much, as they do not require financing/leverage, which is the single most profitable element of the process for banks/investment banks.

³⁷ Limited partners in a fund each receive a profit allocation when earnings are generated. When a limited partner departs, the profit allocation typically is drawn from realized gains, leaving the remaining limited partners with a disproportionate amount of unrealized gains.

generally must feature an asset portfolio with a long-term horizon and a minimal amount of annual turnover; if these characteristics exist, then investors may face tax consequences similar to those of other classes of investment funds.

9.3.5 Exchange-traded funds

ETFs, which have become extremely popular with investors since the late 1990s, have several unique structural features that distinguish them from open- and closed-end funds; in fact, some of these features have been instrumental in fueling growth in the sector. Key features include: the ability for an authorized participant and the ETF trust bank (sometimes known as a custody bank) to develop and liquidate interests in the fund quickly, allowing new shares to be created and existing ones to be redeemed; continuous trading throughout the trading day, rather than at the end-of-day closing NAV; and lower cost structure, allowing investors to create a more economically attractive investment. Efficient allocation and deferral of taxable capital gains to investors in the fund again generates economic benefits.

Legal/regulatory structure

An ETF can be created as an investment management company (e.g. Barclays' i-shares), a UIT (e.g. SPDR, QQQQ), or a grantor trust (e.g. ML HOLDERS). Each structural form has slightly different characteristics.

ETFs created using the investment management company structure are similar to open-end funds, with a fund manager responsible for all coordination and investment allocation decisions. Under this scheme, the investment company tracks, rather than replicates, a reference index, and can use derivatives instead of, or in addition to, physical securities, in order to achieve desired investment results. Indeed, there are times when it is difficult or impossible to replicate the index, meaning the fund manager must attempt to optimize the portfolio within permissible operating parameters.³⁸ UIT-based ETFs are similar to investment companies, but are less flexible – primarily because they must replicate a target index very closely. Despite this fact, some of the largest and most liquid ETFs in the market are structured as UITs (e.g. SPDR S&P 400 Midcap, DIAMONDS, QQQQ). Grantor trust ETFs are used primarily for smaller baskets that track specific sectors.

Regardless of the specific legal vehicle used, however, all are backed or sponsored by major financial institutions or exchanges, and must meet certain basic exchange/regulatory requirements. These include disclosure by the trust bank of the next day's creation/redemption baskets, purchase/redemption of ETF receipts with the underlying securities, and use of a central clearing/depository.

Share creation, trading, and redemption

ETF shares can be created without limitation. In a typical ETF, a portfolio of securities representing the underlying reference index is purchased and deposited by an authorized participant

³⁸ Investment company ETFs can also lend securities held in custody through the stock lending markets in order to generate additional income. This lending, which occurs on a short-term basis, allows fee income to be used to defray some of the fixed operating costs. However, not all ETFs permit such activity; for instance, UIT and grantor trust ETFs typically are barred from lending securities.

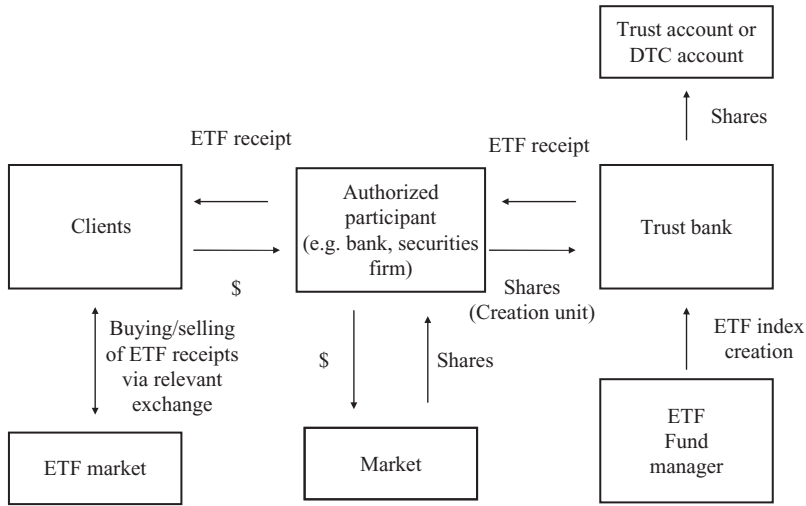


Figure 9.7 The process of creating/trading an ETF

(generally a financial institution that acts as an intermediary on behalf of many smaller investors) with a trust bank,³⁹ which places the securities in trust⁴⁰ and issues divisible depository receipts at an established ratio contained within the ETF's disclosure.⁴¹ The authorized participant may submit an order (generally in 50 000 share blocks, with each block known as a "creation unit") as a result of either client demand or perceived arbitrage opportunities. Once the receipts are issued, they can be subdivided and purchased and sold freely via an exchange, just as any other listed corporate security. If an investor chooses to sell ETF shares it has purchased, the shares are not redeemed by the fund as in a standard open-end fund transaction, but are simply sold to another investor. If overall demand for the ETF declines – that is, no further buying interest exists – the authorized participant can accumulate the divisible receipts, reassemble them, submit the single receipt to the trust bank, and receive the underlying portfolio of securities in return. The authorized participant can then liquidate the securities, thus returning to its original cash position.

Figure 9.7 summarizes the process of creating and trading an ETF via an investment company.

As a result of the continuous creation, trading, and redemption of ETF shares, funds do not typically trade at a discount or premium to theoretical NAV. This, of course, is in contrast to open- and closed-end funds, which frequently trade at substantial discounts (and occasionally premiums) to NAV. Arbitrage activity exists to force the value of the share receipt to track the value of the underlying securities throughout the trading day; the presence of large institutional investors and financial intermediaries, which monitor price discrepancies, helps ensure that prices remain relatively matched. Let us examine how this occurs: if the price of the ETF shares rises above the quoted value of the securities portfolio (i.e. the ETF is trading at a premium),

³⁹ The delivery of shares may be accompanied by a cash payment that serves as a fee to the trust bank for creating the new receipts, and to cover any accrued dividends/interest associated with the underlying portfolio securities.

⁴⁰ For instance, in the US, the portfolio is lodged directly with DTC.

⁴¹ It is worth noting that open-end funds technically can create new shares in the same fashion, but rarely do. In fact, new shares are almost always created through the deposit of additional cash funds by investors.

arbitrageurs purchase the underlying shares,⁴² deliver them to the trust bank in exchange for new receipts, and sell the receipts in the marketplace at the higher price – until the value is driven down to an equilibrium level.⁴³ If the price of the ETF shares falls below the value of the securities portfolio, the reverse occurs: the arbitrageur purchases the ETF shares in the open market, delivers them to the trust bank, receives the underlying securities in exchange, and sells them in the open market at the higher price – again, until the equilibrium point is reached. Since trading, creation, and redemption occur almost instantaneously, the ETF price generally is the arbitrage-free price; in fact, certain exchanges that list and trade ETFs publish separate real-time tickers reflecting the prices of the underlying ETF securities (plus any estimated cash value) on a per share ETF basis, in order to highlight potential arbitrage opportunities. That said, there are instances when ETFs trade at some discount to NAV; this tends to occur when the reference index includes securities that are difficult to value or are extremely illiquid (e.g. those that might be encountered in emerging markets).

ETFs can be traded without limit on a daily basis. Investors can purchase and sell shares in the fund multiple times during the trading day, or they can hold them for the intermediate or long term (note, however, that each purchase and sale attracts a bid-offer spread and a brokerage commission charge).⁴⁴ ETF trading also accommodates a broader range of buy/sell orders than open-end funds. For instance, shares in an ETF can be bought/sold with market or limit prices, at open, at close, on a good-till-canceled basis, and so forth (open-end funds, in contrast, generally can only be bought/sold at the market close). Since ETFs are tradable at current market prices, they are eligible for short selling; this, again, stands in contrast to mutual fund shares, which generally cannot be sold short (unless special arrangements are made). Furthermore, since an ETF share represents an ownership interest in a portfolio of securities (rather than an actual direct ownership of a single security), exchange uptick rules are often suspended;⁴⁵ that is, a short sale can be initiated on a downtick, rather than the uptick required for other common stock transactions. ETFs can thus be used to express outright bearish views, or they can be used as a hedging tool against long positions in securities, index futures, or call options, or against short positions in put options. The existence of continuous pricing and the depth of liquidity in the most popular indexes has led to the creation of derivative contracts on select ETFs (e.g. put and call options are available on QQQQ, DIAMONDS, SPDR). This, in turn, helps promote further activity and liquidity in the ETF, in a self-fulfilling cycle. It is important to note that the voting rights attached to the underlying common shares are not conveyed to investors through the structure; the trust bank or fund manager responsible for assembling the tracking portfolio retains the proxy. Any dividends (or interest) generated by the underlying assets are payable to investors in the ETF (i.e. there is no dividend reinvestment).

It is worth noting that greater customization of reference indexes and baskets is now available, and the turn towards actively managed ETFs has commenced; such “customized” structures have much less liquidity than major index ETFs (suggesting greater bid-offer spreads), and the

⁴² If the ETF's reference index features a futures contract, such as the S&P 500, NASDAQ 100, FTSE 100, Nikkei 225, TOPIX, and so forth, the same arbitrage can be executed: the arbitrageur buys the relevant futures contract, rather than the underlying index shares, and takes delivery in an “exchange for physical” transaction.

⁴³ In practice, the arbitrageur (generally a large, sophisticated institutional investor or intermediary) advises the trust bank of its intent to create and sell a specific number of ETF shares. The trust bank issues the shares, which the arbitrageur delivers against the short position.

⁴⁴ Certain open-end funds, in contrast, limit the number of times an investor can buy and sell fund shares during a given quarterly or yearly timeframe.

⁴⁵ Some exchanges continue to enforce uptick rules on ETFs; the TSE, for instance, only permits short selling of its ETFs on an uptick or a zero plus tick.

inclusion of an actively managed component also means greater fees payable to the asset manager. Bond ETFs are also emerging. The “end game” in the ETF sector ultimately may be a gradual migration from passive index tracking to a more dynamic, short-term positioning strategy.

Operating costs/fees

ETFs generally feature lower costs than equivalent open-end funds; management fees and annual expenses generally are very modest, and processing fees associated with the creation/redemption of ETF shares are small (e.g. $\frac{1}{2}$ basis point). In addition, ETFs are traded on a book-entry basis (e.g. in the US, this occurs through DTC), which helps lower operating costs.⁴⁶ In fact, certain studies suggest that the costs associated with ETFs are approximately half of those attributable to equivalent index-tracking mutual funds, and up to 80 % cheaper than managed equity funds. Naturally, investors buying or selling ETFs through their brokers must pay relevant commissions. They must also cross the bid-offer spread before they can begin generating returns. However, ETFs on global, large-cap indexes are extremely liquid, and thus feature narrow bid-offer spreads. Spreads widen for less liquid baskets or emerging indexes, or when friction costs rise (e.g. costs associated with settlements, research, and so forth). In addition, customized, actively managed ETFs, which seek to outperform, rather than replicate, a particular benchmark, carry investment management fees that are comparable to actively managed open-end funds.

Tax features

ETF investors enjoy certain tax benefits that do not accrue to open- and closed-end fund investors. Specifically, ETF investors do not bear capital gains taxes when other investors in the fund sell their shares, because the underlying securities in the fund are traded, and not sold. The ETF structure is an in-kind exchange of securities for shares, meaning that redemption is not a taxable event within the fund, as it is in a conventional fund. This means that ETF investors that continue holding their shares face no tax consequences until they choose to exit.⁴⁷ This, as we have noted, stands in contrast to the tax treatment of open-end funds, where share redemption by any investor that leads to net asset liquidation triggers an allocation of capital gains to the remaining investors in the fund; investors still holding their shares may thus be impacted by an unexpected tax liability.

⁴⁶ However, settlement periods are often longer than they are for open-end funds, e.g. in the US, ETF settlement is $T + 3$, versus only $T + 1$ for open-end funds.

⁴⁷ Naturally, in certain jurisdictions, the holding period of the investor has an impact on overall tax liability (e.g. ETFs held on a short-term basis – less than one year – are taxed at ordinary income rates, while those held over the long term – more than one year – receive the lower capital gains rate).

Derivative Replication, Repackaging, and Structuring

10.1 INTRODUCTION

In Chapter 2 we introduced derivative contracts as one of the fundamental building blocks of the synthetic and structured asset marketplace. We have built on that foundation in subsequent chapters, demonstrating how derivatives can be paired with capital market securities/loans to create a broad range of structured debt and equity products, such as callable and puttable bonds, synthetic CDOs, CMOs, convertible bonds, commodity- and insurance-linked securities, and so forth. In this chapter we expand on the topic by considering how the contracts can be used to produce additional synthetic and structured contracts. This process follows one of two approaches: using multiple derivatives to create entirely new synthetic assets, or using derivatives to “repackage” or “restructure” assets already in existence. We shall not, of course, discuss the practice of embedding derivatives in host bonds, loans, or securities, as this has already been discussed in previous chapters.

Given the creativity of financial engineers, it is not surprising that the number of combinations in this category is particularly large; even slight variations in product structure can lead to the creation of instruments that produce markedly different investment and risk management results. In the interest of space, we must therefore limit our discussion to several of the most common replication, repackaging, and structuring techniques (providing additional guidance on extensions and variations where relevant). Our focus in this chapter is on synthetic long and short option and swap positions, multiple option/swap positions, callable, puttable, and extendible swaps, and credit derivatives/synthetic credit positions, along with various subcategories within each sector. These derivative structures are summarized in Figure 10.1.

10.2 DEVELOPMENT AND MARKET DRIVERS

Soon after listed option contracts were introduced via the US exchanges, intermediaries and investors began combining short and long put/call positions to create synthetic long and short positions to take advantage of yield enhancement and arbitrage opportunities. Similarly, they began packaging options in particular combinations to express views on market direction and volatility. This process deepened throughout the 1980s, and has since become an essential element of synthetic investment and risk management.

Interest rate and currency-based asset swap packages appeared in the mid-1980s, as institutional investors sought specific asset-related flows that were not directly available in the capital markets, or which were not accessible as a result of regulatory restrictions. Financial creativity and investor demand coincided with attractive credit arbitrage conditions that provided economic benefits for participants, and a very active market in synthetic \$ LIBOR¹ and

¹ Synthetic \$ LIBOR FRNs remain the single largest component of the asset swap market, though considerable growth has occurred over the years in other sectors.

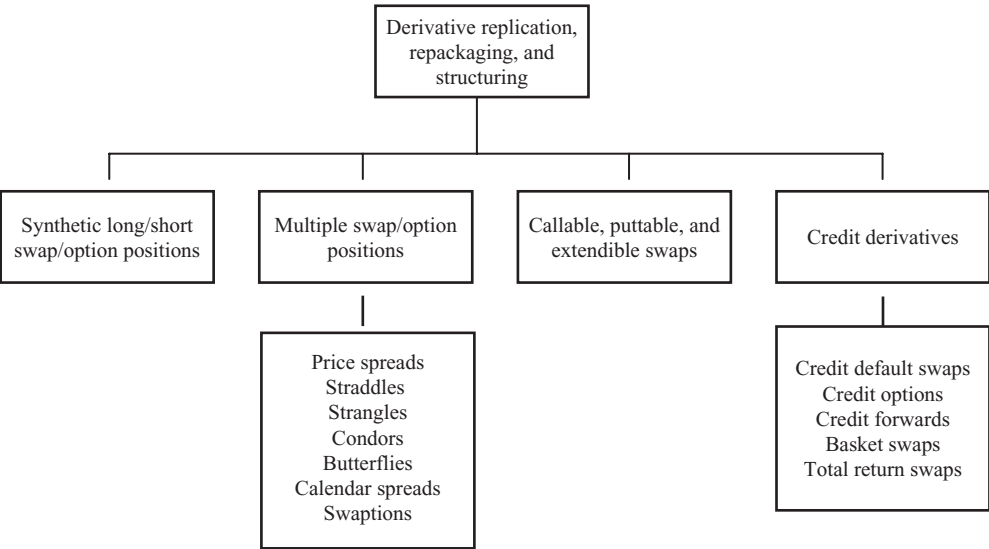


Figure 10.1 Derivative replication, repackaging, and structuring instruments

non- $\$$ LIBOR structures soon appeared. Swapping FRNs into synthetic fixed rate securities, and $\$$ securities into synthetic non- $\$$ securities, also became popular, if rather more opportunistic. During the late 1980s, Japanese speculative financial engineering (i.e. the “zaitech” movement) fuelled a great deal of market flows. Local companies regularly issued bonds with equity warrants – using the equity option to lower funding costs – and reinvested the proceeds in yen assets or synthetic yen assets created through the asset swap market. The process continued until the country’s speculative asset bubble burst in the early 1990s. During the early to mid-1990s, as asset swaps became firmly entrenched in the marketplace, and new variations on the theme emerged, including repackaged/structured notes with embedded derivatives (as discussed in Chapter 5), and callable and puttable asset swaps.

Equity-related derivatives have followed a similar trajectory, starting with the first listed stock options introduced on the CBOE in 1973. The publication of the seminal Black–Scholes model at about the same time, along with the introduction of index futures/options several years later, created a critical mass of interest in the equity derivative sector. Equity warrants and convertibles, including those we discussed in Chapter 8, made use of the same derivative technologies and helped expand the issuer and investor market. OTC structures built on these early successes; OTC equity, basket, and index options and swaps gained prominence during the bull market runs of the 1980s, paving the way for increasingly sophisticated and customized structures by the 1990s.

In some instances, synthetic derivative product development has only been possible once a critical mass of liquidity has been attained in the underlying derivative contract. Indeed, there is little point in creating a synthetic asset from one or more derivatives if the underlying derivatives are too illiquid; any such constraint is likely to reduce or eliminate potential economic benefits. Default swaps and total return swaps are excellent examples of this process. Though the credit derivative market began to form in the early to mid-1990s, default and total return swaps were not widely quoted or traded for several years. In fact, investors

or speculators seeking exposure to a specific corporate bond or bond index were more likely to purchase or short the reference bond or index directly. As market makers began managing their credit portfolios more actively, and quoting two-way prices on a range of credit derivatives, activity began to build, and opportunities for investors to participate in synthetic (leveraged) credit positions via default swaps and total return swaps improved. Indeed, as a robust two-way market began to form, bid-offer spreads on these synthetic credit instruments continued to compress, attracting more end-users eager to assume or transfer credits synthetically. The market is now able to support a broad range of credit references, because the underlying credit derivative market is liquid, active, and well supported. The same is true of other types of synthetic structures true activity builds once the component derivatives are well understood.

Synthetic derivative structures have emerged in a gradual and methodical fashion since the 1970s, as a result of various key market drivers. In particular, the derivative strategies discussed below allow:

- development of very specific risk management and/or investment management goals – the bespoke nature of the instruments and strategies is a direct function of the needs of end-users;
- creation of synthetic assets or liabilities to provide market access, lower funding costs, and/or enhanced returns;
- participation in markets that might otherwise be restricted by regulatory barriers, allowing end-users to express a desired view efficiently and securely;
- monetization of funding and asset strategies resulting from capital mobility and arbitrage conditions;
- implementation of multi-asset hedging strategies in the most efficient way possible.

The OTC derivative market is dominated by institutional dealing, with professional intermediaries and end-users accounting for the majority of all activity. Corporate end-users may be active from either a hedging or investment perspective. Aggressive institutional investors, including hedge funds and pension funds, tend to use the markets to create relative value, yield enhancement, and speculative strategies. Intermediaries that are responsible for developing and executing many of the end-user requirements utilize many of the same instruments and strategies to risk manage their own operations.

10.3 PRODUCT MECHANICS AND APPLICATIONS

In order to consider the structure and applications of various derivative instruments and strategies, we begin with an analysis of the functional aspects of synthetic long and short option and swap positions. This allows us to gain an appreciation of the fundamental components used to create more advanced synthetic packages, including multiple-option strategies, callable and puttable swaps and swaptions, and credit derivatives.

10.3.1 Synthetic long and short option and swap positions

In Chapter 2 we discussed forwards, options, and swaps, and indicated that these often form part of structured/synthetic products. The same contracts can also be combined with one another to produce synthetic long and short positions, synthetic put or call option positions, or synthetic swap positions. This is a very powerful characteristic that ultimately serves to expand investment and risk management opportunities.




Figure 10.2 Payoff profile of long asset position

To demonstrate the process, we return to the payoff profile of the long asset position introduced in Chapter 2, which we replicate as Figure 10.2. We know from this diagram that, as the price of the underlying asset increases, the profit attributable to the position rises, and as the price falls, the loss increases.

We also recall from Chapter 2 that a long call position generates a profit as the price of the underlying asset moves above the strike price. If we overlay the call option payoff profile from Chapter 2 on Figure 10.3 (with the strike set at the origin and premium paid ignored for simplicity), we note that the profit components of the two positions are identical.² This leads us to conclude that half of the long asset position is comprised of a long call option position.

We can extend the discussion by considering an option position that generates a loss as the price of the underlying asset declines. A short put option, which creates a growing liability as the price of the asset falls, features just such a characteristic. If we overlay the short put option payoff profile on the long asset position, as in Figure 10.4 (with the strike again set at the origin and premium received ignored), we obtain an identical payoff profile – again leading us to believe that the second half of the long asset position is comprised of a short put option. Indeed, combining the two options together in a single diagram, as in Figure 10.5, illustrates this clearly.

This simple exercise is meant to demonstrate how a long call and a short put can be paired to create a synthetic long asset position. Naturally, in order for the correct position to be obtained, the strike prices and the maturities of the two options must be identical; if they are not, a mismatch will result.

We carry this process further by considering the opposite scenario: the short asset position. The payoff profile of the short asset position, depicted in Figure 10.6, reflects an increasing profit as the price of the underlying asset falls, and a loss as it rises.

² For simplicity we also ignore the fact that options feature a certain degree of gamma, or convexity, and are thus not perfectly “linear.”




Figure 10.3 Payoff profile of long asset position and long call position

Using the same logic, we can select option positions that generate profits as the market declines, and losses as the market rises. In fact, this can be accomplished with a long put option and a short call option (struck at the same price and expiring at the same time). The combined payoff profiles of the short asset, long put, and short call are depicted in Figure 10.7.




Figure 10.4 Payoff profile of long asset position and short put position




Figure 10.5 Payoff profile of long asset position and long call/short put positions

We recall from Chapter 2 that through put–call parity, the value of a call option less the value of a put option is equal to the present value of the asset price less the strike price. By rearranging these terms, we can develop the same synthetic asset positions reflected in the diagrams above. A call can thus be viewed as a leveraged position in an asset, protected by a put, while a put can be regarded as a short position in an asset, protected by a long call.




Figure 10.6 Payoff profile of short asset position




Figure 10.7 Payoff profile of short asset position and long put/short call positions

It is worth noting that long and short option positions can be used to create synthetic forward contracts, since the long forward is equivalent to a long asset position, while the short forward is equal to a short asset position. In a synthetic forward the strike price, rather than the forward price, becomes the relevant indicator of value; depending on the level of the market in relation to the option strike, there may be a net premium payment (this would suggest the creation of a synthetic off-market forward).³ These relationships can be extended to other investment/arbitrage positions. For instance, buying an asset, selling a call, and buying a put is equivalent to buying an asset and short selling a synthetic asset; there is no market risk in the transaction, indicating that this is equal to the creation of a synthetic Treasury bill position (or some other risk-free asset). Table 10.1 summarizes these findings.

Given these relationships, it comes as no surprise that synthetic option positions can be constructed in the same way. For instance, a long call option can be synthesized from a long asset and a long put, while a long put option can be created from a long call and a short asset.

³ Since shorting the asset is equivalent to borrowing, we can also convey the same concept using the terminology introduced in Chapter 2. For instance, to create a long synthetic forward in stock S , we note that the payoff at maturity T is:

$$S_T - F_{0,T}$$

This position can be purchased on a dividend discounted basis for

$$S_0 e^{-\delta T}$$

This will provide a share of stock at time T . However, in order not to generate an initial investment in the position (e.g. to create a synthetic prepaid, rather than conventional, forward), we assume that the entire amount can be borrowed at a rate of r . At maturity, we take delivery of the stock and sell it for S_T , repaying:

$$S_0 e^{(r-\delta)T}$$

This is equivalent to borrowing to buy a stock position equal to the expiration payoff of the forward. In simplified terms, this means that the forward is simply a package of long stock and a short zero coupon bond (representing the borrowing). We can extend this by rearranging the terms, e.g. a long position in the stock is simply a long forward and a long position in the zero coupon bond, and so forth.

Table 10.1 Synthetic asset positions

Option position	+ Option position	= Synthetic asset position
Long call	Short put	Long asset
Long put	Short call	Short asset

These make intuitive, as well as theoretical, sense: if a company owns an asset and the price rises, it will post a gain; if it also owns a put, it will lose on the asset if the price falls, but will gain on the put – this, of course, is precisely equal to the payoff profile of a long call. Short call and put options can also be created from combinations of options and assets. For instance, a short call can be created from a short asset and a short put, while a short put can be created from a long asset and a short call. These synthetic option positions are summarized in Table 10.2. Note that the discussion is not limited to vanilla options and assets; certain other exotic options, such as those described in Chapter 2, can also be created by combining one or more contracts. For instance, a knock-in barrier option and a knock-out barrier option can be combined to form a standard option.

Other derivative contracts, such as interest rate swaps and cross-currency swaps, can be created synthetically by assembling the components that create particular cash flow streams. Let us consider the case of a vanilla interest rate swap, which requires the fixed-rate payer to pay a fixed rate and receive a floating rate (e.g. LIBOR) on a periodic basis until maturity, and vice versa for the fixed-rate receiver. Decomposing the structure, we note that issuing a fixed-rate bond and investing in an FRN creates a fixed-payer swap. In other words, the fixed-payer swap is a package of a short position in a fixed-rate bond (e.g. the fixed-rate payer must pay the fixed coupon from the bond), and a long position in an FRN (e.g. the fixed-rate payer receives a LIBOR stream from the FRN). A fixed-receiver swap, by extension, is simply a long position in a fixed-rate bond and a short position in an FRN. As indicated in Chapter 2, we can also view a swap as a bundle of interest rate forwards, with each discrete contract calibrated to a maturity that corresponds to the settlement date of a swap, as illustrated in Figure 10.8.

We can also consider a swap’s floating-rate LIBOR coupons in terms of a strip of Eurodollar futures (or Euroyen, or EURIBOR, or some other relevant deposit futures contract). A strip is defined as sequential quarterly contracts out to a maturity of several years (depending on the market), which allows each quarterly rate to be locked in. Since the price of the Eurodollar contract is determined by subtracting the current dollar LIBOR rate from 1, there is a direct link between the floating LIBOR coupons on a swap leg and the Eurodollar futures strip. In fact, the Eurodollar strip rate and the quarterly LIBOR payable/receivable in a swap are the same, apart from slight differences related to quoting conventions/day counts, and more obvious differences related to reset dates. Though convergence between the two rates is not always perfect, the relationship is close enough that the rates are quite fungible.

Table 10.2 Synthetic option positions

Asset position	+ Option position	= Synthetic option position
Long asset	Long put	Long call
Short asset	Long call	Long put
Short asset	Short put	Short call
Long asset	Short call	Short put

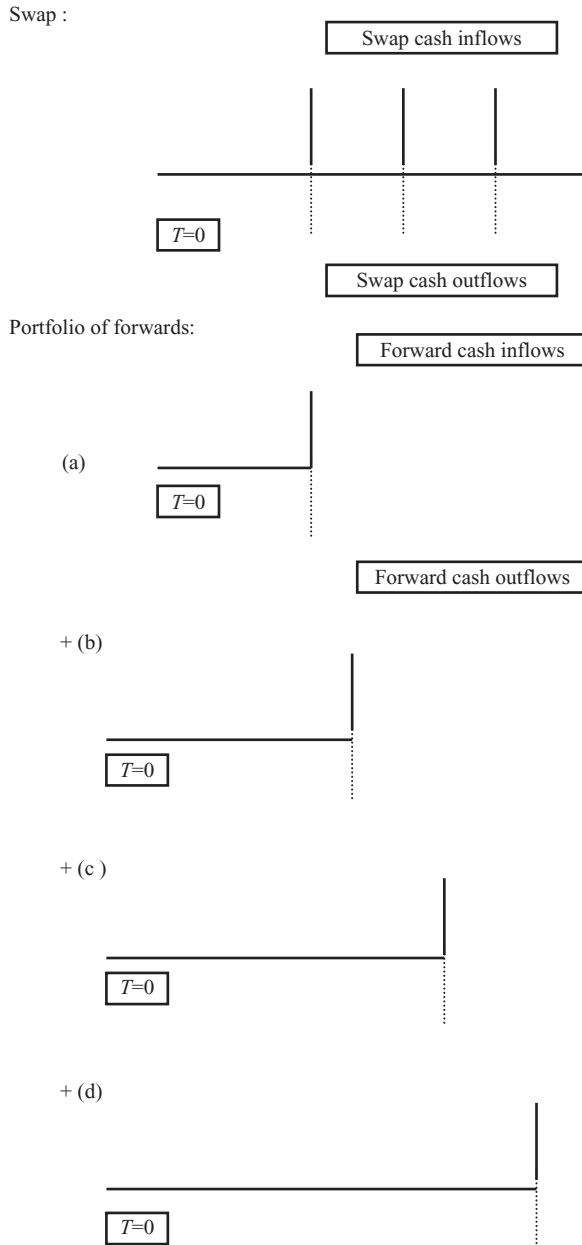


Figure 10.8 Interest rate swaps and a portfolio of forwards

Consider a firm that wants to enter into a \$10 m three-year swap with a bank, where it pays fixed and receives LIBOR flat. This is equivalent to the company selling a series of Eurodollar futures contracts with settlement coinciding with each swap settlement date. Each contract requires the company to deliver \$10 m of a Eurodollar deposit on each settlement

date, receiving \$10 m in exchange. We assume the bank issues a Eurodollar deposit at LIBOR flat at each settlement period in order to raise the necessary funds. This means that if LIBOR is greater than the fixed rate, the company has a net cash inflow for that settlement period, while if LIBOR is less than the fixed rate, the bank has a net cash inflow – precisely what we would expect on a standard swap. In fact, given the relationship between short to intermediate swaps and deposit futures, arbitrageurs ensure that swap levels trade at the same levels implied by the futures strip. If any discrepancies arise, arbitrage forces will soon equalize them. Thus, if the Eurodollar strip is trading rich, it is economically favorable to sell the strip rather than buy the swap; similarly, if the strip is cheap, it is better to buy the strip than sell the swap. The richness and cheapness will be exploited, bringing the relationships back in line.

The actual computation of the Eurodollar strip rate, which can be used to derive the implied swap rate, can be determined via:

$$\begin{aligned} & [1 + R_0(D_0/360)]^* [1 + F_1(D_1/360)]^* \dots [1 + F_n(D_n/360)] \\ & = [1 + R(365/360)]^{N*} [1 + R(D_r)/360] \end{aligned}$$

where

R is the Eurodollar strip rate expressed as a money market yield

R_0 is the spot LIBOR rate

F_1 is the lead futures rate (100 – price)

F_n is the last futures rate in the strip

D_i is the actual number of days in each period

D_r is the total number of days in the strip

N is the number of whole years in the strip.

A swap can also be viewed as a collar, which we describe in greater detail below, where a firm buys a strip of Eurodollar put options (e.g. a cap), and sells a strip of Eurodollar calls (e.g. a floor) at three-month intervals for the entire maturity of the swap transaction (or vice versa). Though this approach yields the desired swap results, it involves extra expense in the form of additional transaction costs. Not surprisingly, intermediaries that manage swap books often use deposit futures and options, collars, and OTC FRAs⁴ to price and manage their positions. The synthetic interest rate swaps are summarized in Figure 10.9.

We have indicated previously that currency swaps can be created synthetically by assembling a series of spot and forward foreign exchange contracts. Thus, a standard currency swap can be created by combining a US\$ interest rate swap and a strip of long-dated forwards, or a foreign currency interest rate swap and a series of rolling short-dated forwards. It can also be assembled by combining foreign currency interest rate swaps and foreign currency basis swaps. For instance, a pay fixed yen/receive \$ LIBOR currency swap is a package of a standard pay yen fixed/receive yen LIBOR interest rate swap and a pay yen LIBOR/receive \$ LIBOR currency basis swap.

10.3.2 Multiple swap/option positions

There are instances when multiple derivative contracts can be combined to create synthetic structures with unique payoff profiles. These multiple swap/option positions can be

⁴ It is important to note that OTC FRAs contain elements of counterparty credit risk, and are thus not perfectly fungible.

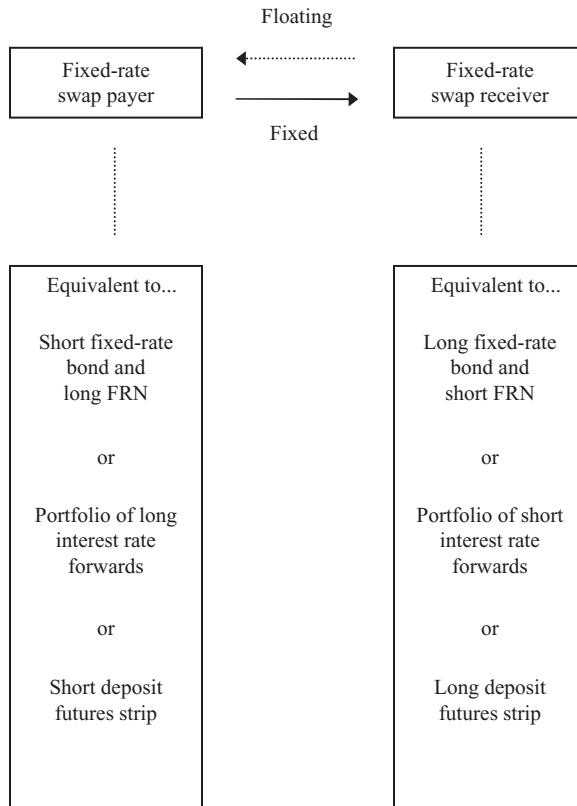


Figure 10.9 Synthetic interest rate swaps

used to create synthetic funding, achieve investment/portfolio goals, speculate on the direction/volatility of rates, facilitate capital markets access, monetize embedded contract value, and aid in hedging and gap management practices. The use of options in multiple-derivative packages provides buyers with structural flexibility, and sellers with additional income. As we know, selling options generates premium income that can be booked into current income or used to lower all-in funding costs. As in any option sale, the transaction may be covered or naked; if covered, then the sale serves as a yield enhancement or funding management transaction, and if naked, then it is most likely to be a speculative transaction.

Though the marketplace features many different multiple-derivative combinations, we consider several of the most common in this section, including multi-option strategies, swaptions, callable, puttable, and extendible swaps, and other structured derivatives.

Multi-option strategies

We have noted immediately above that multiple options can be used to create synthetic long and short positions. They can also be combined to create very specific synthetic exposures to the volatility and/or direction of a particular asset or market. Several of the most common strategies include bull spreads, bear spreads, straddles, strangles, butterflies, condors, and

calendar spreads. Other variations on the theme (e.g. strips, straps, backspreads, and so on) can be developed along similar lines.

Price spread

The price spread (sometimes known as a vertical spread or money spread), the single most actively used multi-option strategy, is used to express bullish or bearish views on market direction. The strategies can be created in four different forms: bullish call spread, bearish call spread, bullish put spread, and bearish put spread. Through the bullish call spread, an investor buys a call option and sells a higher-strike call option; the purchase of the closer-to-the-money option results in a net outflow of premium, but locks in a profit equal to the difference between the two strikes. The bullish put spread is created when an investor sells the high-strike put and buys the low-strike put; in this case, the sale of the closer-to-the-money put creates a premium inflow, and locks in a liability equal to the difference between the strikes as the market sells off. Bearish spreads are simply the opposite: in the bearish call spread, the investor sells the low-strike call and buys the higher-strike call, generating a premium inflow and locking in a known liability as the market rallies; in a bearish put spread, the investor purchases the high-strike put and sells the low-strike put, paying a net premium but locking in a gain equal to the strike differential as the market trades down. Price spreads are designed primarily to express a view on market direction; this stands in contrast to the volatility-driven strategies discussed immediately following (which are volatility-sensitive, but market neutral).

Straddle

The straddle is a common multi-option position, taken in order to profit from, or protect against, volatility. The basic long straddle consists of the purchase of a call and a put on an underlying reference at the same strike price and with the same maturity. In exchange for this delta-neutral, long-volatility position, the buyer pays the seller a premium. By creating a straddle, the buyer seeks to take advantage of market volatility and is relatively indifferent (or uncertain) as to the direction of the market; the intent is to capitalize on a large market movement expressed through rising volatility. Thus, if the market moves up sharply (i.e. volatile on the upside), the call leg of the straddle will increase in value; the expectation is that the upward movement in the call will more than offset the downward movement in the put, resulting in a net profit. The reverse is true in the case of a sharp downward movement in the market (i.e. volatile on the downside). The seller, in contrast, sells a straddle in expectation of a relatively tranquil market. If the reference trades in a narrow range, neither leg of the straddle will appreciate in value, meaning that premium received from selling the package of options will be preserved. In fact, premium decay in favor of the seller will accelerate as time to maturity draws closer, particularly if the market remains range bound. The straddle can thus be seen as a simple synthetic strategy that is used to capitalize on volatility, rather than market direction. Figures 10.10 and 10.11 highlight the payoff profiles of long and short straddles.

Strangle

The strangle is a variant of the straddle. The strangle, like the straddle, seeks to take advantage of movements in volatility, rather than market direction, by using a put and a call with the same expiry date. Unlike the straddle, however, the strike prices on the put and the call are different, suggesting that a greater amount of volatility is required before the position moves




Figure 10.10 Payoff profile of long straddle position

in-the-money. This means that the buyer pays less premium for the position than it would for a straddle, and the seller can withstand a greater amount of market movement before a potential liability begins to accrue. The basic long strangle thus consists of the purchase of a call and a put on an underlying reference at different strike prices, but with the same maturity, as illustrated in Figure 10.12; the seller's position is reflected in Figure 10.13.

Butterfly

The butterfly is another volatility-based multi-option strategy comprised of long and short call/put options. A long butterfly is a package consisting of long low- and high-strike options and two short mid-strike options; this can also be viewed as a short straddle and a long strangle (e.g. a short straddle without the extreme downside risk). The short butterfly is a package

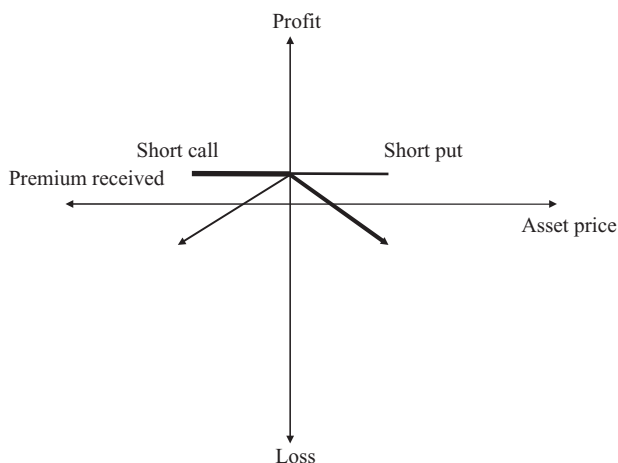


Figure 10.11 Payoff profile of short straddle position




Figure 10.12 Payoff profile of long strangle position

comprised of short low- and high-strike options and two long mid-strike options; this again is simply a long straddle and a short strangle. Figure 10.14 illustrates the payoff profile of the long butterfly; the payoff profile of a short butterfly is simply the “mirror image.”

Condor

The condor is an extension of the butterfly, with the middle strikes set further apart – thus requiring greater volatility to move the position in-the-money (or further in-the-money). The long condor is thus a package consisting of long low- and high-strike options, a short mid-strike low option, and a short mid-strike high option; the short condor, by extension, is a package of short low- and high-strike options, a long mid-strike low option and a long mid-strike high




Figure 10.13 Payoff profile of short strangle position




Figure 10.14 Payoff profile of long butterfly position

option. The payoff profile of a short condor is illustrated in Figure 10.15; the payoff profile of a long condor is simply the reverse.

Calendar spread

The calendar (or time) spread is a multi-option strategy that seeks to take advantage of market volatility as related to time. Since options on the same underlying reference with identical strikes but different maturities trade at different levels, the use of long/short calendar spreads allows an institution to capitalize on the differences. A short calendar spread, for instance, is created by buying a short maturity option and selling a longer maturity option, both with the same strike. The time value of the shorter maturity option decays more quickly than that of the longer maturity option; the seller thus hopes for a calm market in order to capture and preserve




Figure 10.15 Payoff profile of short condor position

Table 10.3 Sample swaption quotes

Swaption structure	Fixed-rate strike (%)	Receiver spread (bps)	Payer spread (bps)
1 × 3	5.00	5–10	70–75

Premium dealer pays for buying the fixed-payer swaption

Premium dealer demands for selling the fixed-payer swaption

the net premium inflow. A long calendar spread is created in the opposite fashion, and requires a volatile market in order to generate value for the buyer.

Swaptions

A swaption, as the name indicates, is a combination of a swap and an option. We have already discussed the two derivative contracts separately: a swap involves the periodic exchange of interest/currency coupons on a fixed/floating-rate basis, while the option grants the purchaser the right to buy or sell a particular asset at a set strike price. The swaption, then, is a synthetic derivative contract that grants the buyer the right to enter into a swap with a fixed rate defined by the strike. It can be described further as a payer (or put) swaption, and a receiver (or call) swaption:⁵ the payer swaption grants the buyer the right to enter into a swap where it pays a fixed rate (i.e. the strike level) and receives floating (e.g. LIBOR); the receiver swaption conveys the right to receive the fixed rate and pay the floating rate. If the buyer chooses to exercise the option, all terms related to the swap come into force, including rates, payment frequency, notional amount, and final maturity.

The market quotes swaptions in terms of a strike rate, option maturity, and swap maturity. For instance, a 1 × 3 5 % US\$ payer swaption grants the buyer a one-year option to enter into a three-year swap, where it pays US\$ fixed at 5 % and receives US\$ LIBOR flat. If rates rise above 5 % during the 12-month period, it becomes economically rational for the buyer to exercise the swaption to pay 5 % and receive LIBOR for a period lasting three years; conversely, if rates remain below 5 %, the contract remains unexercised. Exercise may be defined in American or European terms, and settlement may be arranged in cash or swap terms. Swaption quotes are available on a range of global rates and a series of common structures related to the tenor of the option, the tenor of the underlying swap, and the fixed-rate strike (note that the fixed-rate strike is quoted as an absolute rate rather than as a spread to a benchmark). Dealers provide bid-offer premium spreads for both payer and receiver flows for each of these structures, as noted in Table 10.3; these premiums are applied to the notional value of the transaction to generate a dollar value premium.

The value of a swaption, like any other option, is influenced by the strike level, time to maturity, volatility, and risk-free rate. For instance, a receiver swaption with a high strike is more valuable than one with a low strike, as it provides the buyer with a larger cash flow upon

⁵ The put and call swaption terminology can be compared with that of a standard bond option: a call option on a bond gives the purchaser the right to receive the bond's cash flows once a certain strike level has been reached; similarly, a call swaption grants the right to receive fixed cash flows under a swap. The reverse is true for bond put options and put swaptions.

Table 10.4 Swaption exercise states

	Rates rise through strike	Rates fall through strike
Payer swaption	Exercise	No exercise
Receiver swaption	No exercise	Exercise

exercise; conversely, a payer swaption with a higher strike is less valuable than one with a lower strike, as it requires the buyer to pay a larger cash flow on exercise. Swaptions with a longer time to maturity are more valuable than those with shorter maturity, as they provide more time for the contract to move in-the-money. Similarly, a more volatile market increases swaption value, as greater volatility can force a contract in-the-money (or further in-the-money).

Actual exercise of the swaption depends on the market level of interest rates. If rates fall through the strike of the swaption, the buyer of a payer swaption will not exercise, as it is economically favorable to pay the lower market rate, while the buyer of a receiver swaption will exercise, as it will receive an above-market rate. If rates rise above the swaption strike, the reverse occurs: the buyer of the payer swaption will exercise in order to lock in a maximum payment level, while the buyer of the receiver swaption will choose not to exercise. These four states are summarized in Table 10.4.

Swaptions are highly flexible and can be used to meet a range of goals. An institution may purchase a swaption in order to lock in a known or anticipated funding, hedge, or investment level for a future date, and it might sell a swaption to take advantage of a view on the movement and/or volatility of interest rates, or to monetize value in a callable or puttable bond structure. We will consider several of these applications in more detail below.

An institution believing rates will rise can purchase a payer swaption or sell a receiver swaption; one believing that rates will fall can sell a payer swaption or buy a receiver swaption. For instance, an institution believing rates will rise can buy a payer swaption; if rates rise, it can exercise the right to pay fixed at below-market levels and receive higher-market floating rates. Alternatively, it can exercise and then sell the swap in the market for its intrinsic value, or enter into a reverse market transaction, where it receives the higher-market fixed rate against floating (thus locking in the spread on the fixed legs). Or, it may simply sell the in-the-money swaption in the marketplace to another party. If rates decline, the institution will let the swaption expire unexercised and will simply lose the premium it paid to acquire the contract. It may also sell the receiver swaption to capitalize on its expectation of rising rates. The institution believing that rates will decline can sell a payer swaption and collect premium; if rates remain below the strike, the contract will expire unexercised and it will have earned premium income (naturally, if rates rise above the strike, the institution's liability increases). Alternatively, it can buy a receiver swaption, exercising to lock in a higher fixed rate if market rates fall below the strike.

Swaption-driven liability and asset strategies are popular. For instance, a floating-rate borrower that does not believe that interest rates will decline can sell a receiver swaption to generate premium that reduces funding costs; a borrower that believes rates will fall can sell a payer swaption to achieve the same result. Or, an institution that owns an interest-rate-sensitive asset can buy a payer swaption as protection, while one that is managing an interest gap created by short-duration assets and long-duration liabilities can buy a receiver swaption. Swaptions can also be used to protect or leverage a volatility position: the issuer of a callable bond can sell a receiver swaption to monetize value, while the issuer of a puttable bond can buy a payer swaption to protect against rising rates. Regardless of the specific strategy, the premium inflow or outflow must be considered as part of the all-in funding cost or investment return computation.

Let us consider the example of an FRN issuer that sells an out-of-the-money payer swaption (where fixed rates are below the forward rates associated with the option's final maturity): if rates fall and the buyer of the payer swaption exercises, the issuer/seller will have converted a portion of its FRN funding into fixed rates, which is a perfectly acceptable goal; if rates do not fall, the debt remains floating, but the issuer will have earned premium from the swaption sale that it can use to offset its all-in funding costs. We can imagine the opposite scenario: an investor holding an FRN can sell a receiver swaption, taking in premium on the sale. If the swaption remains unexercised, then the investor enjoys the benefit of enhanced portfolio yield via the earned premium; if the swaption is exercised, the investor locks in a fixed rate on a portion of its FRN position. In a further example, the issuer of a fixed-rate puttable bond receives a lower funding cost from the initial sale of the put option, and can supplement this by entering into a fixed/floating interest rate swap to exchange its funding into LIBOR, and buying a receiver swaption from the put date to the final bond maturity. If rates fall and investors do not exercise their puts, the issuer can exercise the receiver swaption to keep its borrowing at lower floating rates (which offsets the sum paid for the swaption).

Swaptions can be used to monetize value in the callable and puttable bonds described in Chapter 3. Monetization generally arises from an arbitrage opportunity that makes it possible to buy "cheap" options and sell "rich" ones. The arbitrage works because of pricing discrepancies that periodically appear between issuers, investors, and intermediaries. Empirical evidence suggests that the implied volatility of options in the capital markets is often lower than in the derivatives market at large, as investors tend to focus on yield, rather than total return, when determining value. This means that investment decisions are made on a yield-to-worst basis, ascribing zero value to the optionality contained in puttable or callable bonds. Furthermore, investors tend to lack information on corporate posture regarding callability and the probability of a call action, generally believing that no such calls will occur. Further pricing discrepancies can arise from relative value differentials among synthetic or structured securities, and between the swap curve and the corporate credit curve. For example, if the swap curve is flat to inverted versus the corporate curve, implied forward rates are lower in the swap market – this means that receiver swaptions in the swap market are worth more than call options in the corporate market, leading to potential arbitrage opportunities. An issuer of a callable bond *de facto* purchases a cheap call option from investors via the bond, and can then sell a receiver swaption to an intermediary, which gives the intermediary the right to exercise into a swap where it will receive fixed and pay the issuer floating. If rates fall, the intermediary will exercise the swaption, paying LIBOR and receiving fixed; the issuer, in turn, will call the bond, and refinance via a floating-rate note at a new lower rate. Rather than paying fixed rates to the original bond investors, it is now paying fixed rates to the swaption intermediary; the floating cash flows it receives from the intermediary are redirected to the new FRN investors. Naturally, if rates rise, the entire bond/swaption package remains outstanding until maturity. Figure 10.16 summarizes this example. Note that this structure works if the swaption can be sold for more on a present value basis than the market yield on straight debt (e.g. ex-call option); if it cannot, then the arbitrage does not work and the issuer is better off issuing noncall debt directly.

When forward swap rates are high, payer swaptions are worth more than put options in the corporate market, again leading to profit opportunities with puttable bonds. For instance, an issuer might float a 15-year fixed-rate bond that is puttable at par after five years. The issuer, which is short the put to investors, can buy a five-year payer swaption on a ten-year fixed rate – and achieve an arbitrage if the swaption is cheaper than the option sold via the puttable bond. If this occurs, then its all-in cost will be lower than on a straight bond, and value will have been

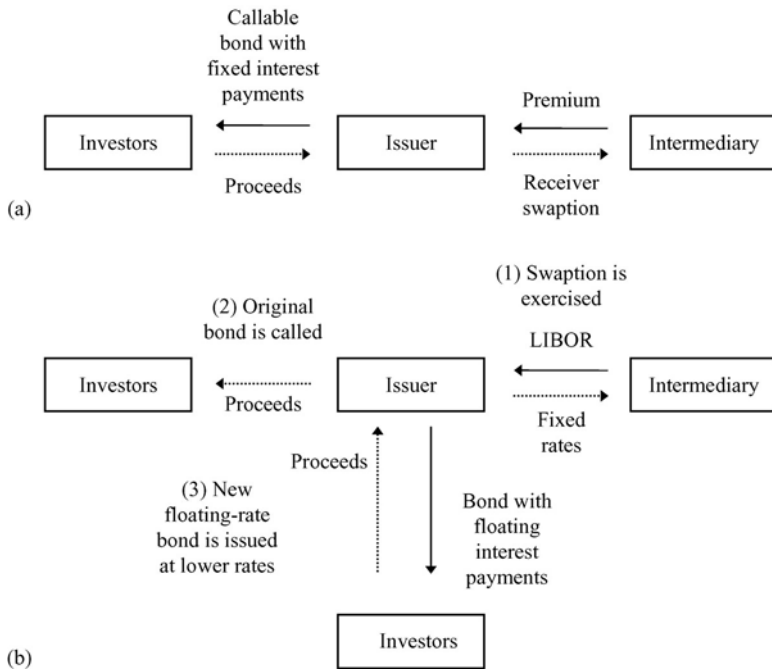


Figure 10.16 Callable bond/receiver swaption arbitrage. (a) Initial position and rising-rate scenario; (b) Declining-rate scenario

monetized. If rates fall below the issuance yield in year 5, the bond will remain outstanding and the swaption will remain unexercised. If rates rise, investors will put the bond back to the issuer, and the issuer will exercise the payer swaption to continue paying a known fixed rate versus LIBOR.

In general, swaption arbitrage tends to appear when interest rates are volatile (which increases the value of the options), the yield curve is relatively flat, and credit spreads are either tight (callable bonds) or wide (puttable bonds). The swaption arbitrage does not always exist; in fact, it appears to be a cyclical phenomenon driven by issuance activity in specific credit sectors, the needs of investors, and the creation of new structured liability opportunities.⁶ It is also important to stress that the actual net swaption arbitrage that is achieved is a function of an issuer's future creditworthiness; if credit deterioration occurs, so that the spread over LIBOR, or a government benchmark, widens, the actual funding arbitrage will compress.

Put-call parity can be applied to the swaption market to create a series of synthetic positions. For instance, an investor can purchase a payer swaption or purchase a receiver swaption and a forward starting swap; the two results are identical. Synthetic construction can also be used to create different forms of funding. Noncallable/nonputtable bonds, for example, can be converted into callable/puttable equivalents, and callable/puttable bonds can be converted into fixed or floating noncallable and nonputtable equivalents – all through the use of the appropriate

⁶ While investor undervaluation of embedded options historically has been a market driver, greater sophistication within the investment community, as well as access to better modeling tools and corporate information, means that the arbitrage may begin to compress or appear less frequently.

underlying asset and a swaption. This is again a powerful application, particularly when issuers or investors are seeking to obtain funding or investment results that are not directly available in the marketplace.

Consider that an issuer can create synthetic noncall debt by issuing a callable bond and selling a receiver swaption, which creates fixed-rate funding to the call date and synthetic fixed-rate funding from the call date to the maturity date. Thus, if a company issues a fixed-rate callable bond with a five-year final maturity, and sells a payer swaption, it will have a fixed funding profile regardless of what happens to rates: if rates rise, the bond will remain outstanding and the swaption will go unexercised; if rates decline, the bond will be called and the swaption buyer will exercise the swaption to receive fixed in a swap with the issuer. These scenarios make it appear as though the callable bond is, in fact, noncallable for its entire life. Synthetic callable bonds can also be created using noncallable debt and swaptions.

An issuer can also synthesize a puttable bond from a package that includes a callable bond and the sale of back-end receiver and payer swaptions. The buyer of the swaptions thus has the option to exercise either swaption and force the issuer to pay fixed or receive fixed. One of the two swaptions purchased is used to offset the option embedded in the callable bond, while the other serves to create the puttable structure. Thus, if the issuer floats a seven-year bond that is callable after year 5 and rates rise, the buyer of the swaptions will exercise the contract to receive floating for the final two years of the issue. The issuer thus has fixed funding for the first five years and floating-rate funding for the last two. If rates fall, the buyer will exercise the contract to receive fixed for the last two years; the issuer then calls the outstanding bond and funds via a new FRN in the lower-priced market. Puttable bonds can also be converted into nonputtable securities using similar techniques. Note also that callable/puttable and noncallable/nonputtable bonds can be created synthetically using callable and puttable swaps, as discussed immediately following.

Callable, puttable, and extendible swaps

Callable and puttable swaps, like swaptions, are structured derivatives created by combining swaps and options. A callable swap is a swap with an embedded option that allows the fixed-rate payer the right to terminate the transaction at some future point. The puttable swap, in contrast, allows the fixed-rate receiver to terminate a swap. We can view these synthetic structures in terms of their constituent components. For instance, a five-year swap that is callable in three years is equal to a standard five-year swap and a three-year receiver swaption on a two-year swap. Similarly, a five-year swap that is puttable in three years is equal to a five-year swap and a three-year payer swaption on a two-year swap. Since callable and puttable swaps are constructed from swaps and options, put-call parity relationships hold, primarily for European structures. For example, knowing that a callable swap is simply a package of a swap and a receiver swaption, we note that a callable swap and a short receiver swaption are precisely equal to a swap, as are a puttable swap and a short payer swaption.

Let us consider the puttable swap in more detail. A puttable swap gives the fixed-rate receiver the right to terminate the underlying swap transaction at some future date, meaning that the fixed-rate payer is selling a receiver swaption on an underlying swap that starts during the life of the swap and terminates at maturity. In exchange for selling the right to cancel, the fixed-rate payer receives a premium in the form of an upfront cash flow or a higher fixed rate on the underlying swap. The purchaser of the puttable swap may choose to enter into this type of transaction if it is unsure of its need to preserve the swap at some point in the future, or

if it expects that interest rates will decline. One common puttable swap structure is the index amortizing rate (IAR) swap, where the notional principal of the swap amortizes as interest rates decline through particular thresholds (i.e. strike levels). This can be viewed as a strip of partial cancellations of the swap tied to the downward movement of rates. In this instance, the fixed-rate payer pays the fixed-rate receiver the higher rate (premium) for the right to cancel portions of the swap as rates fall. IARs have proven to be quite popular as hedges for securities/transactions with uncertain cash flows that depend on the movement of rates (e.g. MBS, ABS, leases).⁷

Callable and puttable swaps can also be used to synthesize callable/puttable bonds. For instance, an issuer may find it economic to convert a callable bond into a synthetic FRN using a callable swap. Indeed, the all-in funding cost can be lower on the synthetic FRN, because the issuer can purchase the call embedded in the callable bond at cheap levels and sell it in the swap market at higher levels. We can consider a scenario where a firm issues a fixed-rate callable bond and enters into a callable swap where it receives fixed and pays LIBOR. If rates rise, no change in the structure occurs, meaning the issuer continues to fund on a floating-rate basis. If rates fall, however, the callable swap counterparty will terminate the swap transaction, the issuer will call the fixed-rate bond, and will then reissue in the FRN market directly, again achieving its desired floating-rate funding. Puttable bonds can also be converted into synthetic FRNs using the puttable swap. Consider, for example, a company that issues a ten-year fixed-rate bond that is puttable after five years, and enters into an associated puttable swap where it pays LIBOR against fixed rates. If rates rise after year 5, investors in the puttable bond will put the security back to the issuer, which will terminate the swap. It will then issue new floating-rate funding in the market directly, maintaining the same LIBOR-based funding profile. If rates fall, the bond and the swap remain outstanding for the full ten years, meaning that the issuer achieves its goal of maintaining floating-rate financing.


The reverse scenario is also possible: noncallable debt, for instance, can be created synthetically by pairing a callable bond with a callable swap. This opportunity again arises when investors demand less for the call embedded in the callable bond than its true worth; the issuer, long the call, can issue the bond, and sell a callable swap granting the counterparty (which pays fixed) the right to terminate the underlying interest rate swap associated with the funding. If rates decline, the issuer calls the bond and funds at a new lower floating rate; the callable swap counterparty, in turn, calls and terminates the swap. Throughout this transaction, the issuer continues to pay floating rate (e.g. via the callable swap until call date, and via the FRN market when refunding occurs), as if the underlying callable bond was actually noncallable.

The extendible swap, like the callable or puttable swap, is a structured derivative package comprised of an interest rate swap and an option that grants one party the right to require its counterparty to continue a previously contracted swap under existing terms for an additional period of time.⁸ The payer extendible swap, for instance, is a combination of a fixed-payer interest rate swap and a payer swaption. Assume that the purchaser enters into an underlying swap, where it pays 5 % and receives LIBOR for two years, and a swaption that gives it the right to preserve the 5 % rate for a further three years. If rates rise above 5 %, the purchaser can exercise the payer swaption so that it continues paying 5 % for an additional three years, effectively extending the original two-year swap into a synthetic five-year contract. If rates


⁷ In fact, IARs feature positive convexity, which makes them ideal for hedging negative convexity securities such as MBS.

⁸ Early extendible swap deals were associated with bonds with debt warrants (allowing the issuance of additional debt). For instance, investors could purchase a bond with warrant and associated extendible swap, thereby extending the swap terms of the original debt issue to cover new issuance upon exercise.

(a) Original two-year interest rate swap



+ (b) Three-year payer swaption



= (c) Synthetic five-year swap




Figure 10.17 Payer extensible swap

decline, the option expires unexercised. This example is highlighted in Figure 10.17. The receiver extensible swap, in contrast, is comprised of a fixed-receiver swap and a receiver swaption. In this case, the buyer can choose to extend the underlying swap and receive fixed rates for an additional period of time; it will do so if rates fall below the strike level of the swaption. As in all other option-based structures, the party purchasing the option to extend pays a premium, either upfront or as an adjustment to the rates payable/receivable under the swap.

Other structured derivatives

In addition to the structures described above, forwards, swaps, and options are combined regularly to generate other investment and risk management results; while the list of possible combinations is lengthy, we describe some of the most common in this section.

Collar

A combination of a long call option (or cap) and a short put option (or floor), or vice versa. This structure is also known as a range forward, as it gives one party a synthetic long position and the second party a synthetic short position over a particular range of values; it may also be referred to as a zero cost collar if the premiums paid/received are precisely equal. Figure 10.18 illustrates a sample long collar. Collars are commonly used in the liability management




Figure 10.18 Long collar

sector as a way of lowering and locking in funding costs. Consider, for example, a company that can issue a three-year FRN at $L + 100$ bps and wishes to optimize its funding. It arranges a three-year collar with a floor strike of 4 % and a cap strike of 6 %; since the floor strike is slightly closer to the money, it receives a net premium inflow of 24 bps (e.g. -50 bps on the purchased cap and $+74$ bps on the sold floor). Based on a straight-line amortization of the premium (ignoring any discounting for simplicity) over six semi-annual periods, the company uses the collar to lock in maximum funding levels of $L + 92$. Recall, also, that collars can be embedded in structured notes, e.g. the collared FRNs described in Chapter 5.

Break forward

A combination of a forward contract and an option. The break forward can be regarded as a form of long call option (or long put option) that gives the purchaser the ability to limit downside by “breaking” the contract. For instance, if a forward is set at a price of 150, a break forward might permit an institution to buy at a forward rate of 155, and unwind (“break”) the contract at 150; the difference between the two prices is the implicit premium paid for the option.

We illustrate the standard forward, long collar (range forward), and long break forward in Figure 10.19 to demonstrate the slight differences between the three structures.

Knock-out forward

A combination of a forward contract and a knock-out (barrier) option that remains in force until the barrier price is reached. This contract allows an investor to achieve a favorable rate on the underlying forward by selling the knock-out option to the intermediary. The benefits of the favorable rate can only be reaped as long as the barrier is not reached during the life of




Figure 10.19 Long forward, long break forward, long collar (range forward)

the contract. An extension of this structure is the double knock-out forward, which generates an even more favorable forward rate through the sale of knock-out options on both the upside and downside.

Semi-fixed swap

A combination of a swap and a digital option. This structure is akin to the range floaters we discussed in Chapter 5, where the cash flows paid or received in a given period depend on the movement of interest rates. Under a typical semi-fixed swap, one party is required to pay either an above-market rate or a below-market rate in exchange for LIBOR, depending on whether the digital trigger has been hit. This effectively allows the party to pay a reduced fixed rate in exchange for LIBOR, as long as the trigger level has not been reached. A multivariate version of this structure employs a non-interest rate trigger as a reference (e.g. oil prices, equity index level, currency rates).

Forward swap

A combination of a forward and a swap (also known as a delayed start swap), or two swaps with different maturity dates and pay/receive flows. As the name indicates, a forward swap is a swap that commences at some future date, with all transaction details contracted on trade date. For instance, two parties might enter into a three-year fixed/floating swap that begins in two years. Two years hence, one party will pay LIBOR flat and the other will pay the fixed rate contracted on the original trade date; these flows will continue until the underlying swap concludes three years after the swap start date (i.e. five years after the forward swap start date). Unlike the swaption, which conveys a contingent forward starting date (depending on when (if) exercise occurs by the buyer), the forward swap is not a contingent transaction, but a certain one – the swap will commence at a future date, with both parties to the transaction obligated to fulfill their roles. While the transaction can be viewed as a package of a forward and the underlying swap, it can also be viewed as a series of partially offsetting swaps. For instance, in the example above, the party paying fixed rates for three years in two years' time has booked a




Figure 10.20 Yield curves underpinning a differential swap

five-year swap, where it pays fixed and receives LIBOR, and a two-year swap, where it receives fixed and pays LIBOR; since the LIBOR legs cancel each other out, the structure condenses to a three-year stream of pay fixed cash flows commencing in year 2. Forward swaps are used regularly to lock in payments/receipts expected to arise from assets or liabilities appearing in the future; the ability to solidify future cash flows is obviously a useful investment and risk management tool.

Differential swap

A combination of a basis swap and a protective foreign exchange contract (also known as a quanto swap). In this transaction, parties exchange two currency basis indexes (e.g. LIBOR versus EURIBOR), but all payments are made in a single currency (e.g. all in US\$ or all in euros). The differential swap is used to take advantage of perceived views on implied forward rates in different markets without assuming foreign exchange risk, as illustrated in Figure 10.20.⁹ The risk/return to the end-user in this structure is based primarily on a compression or widening of the spread between the two yield curves, rather than absolute changes in either the yield curve or the currency rate. The results can be used to boost yield on an underlying asset position, or to lower funding costs.

A pure yield-curve swap functions in a similar manner, except that it is confined to a single market, and thus involves no foreign exchange conversion. An investor or end-user simply uses the swap market to express a view on implied forward rates by taking either side of a short-term market rate versus a long-term market rate, resettable every three or six-months (e.g. comparing the implied long-term forward rate in six months to the implied six-month LIBOR rate). The same swap can also be used to hedge specific liability flows.

Zero coupon/swap packages

A combination of a zero coupon bond and a swap. In some instances, it is economically rational for a company seeking funding to issue a zero coupon bond with an attached interest rate or currency swap; this occurs when rate levels and arbitrage opportunities in one or more markets

⁹ It should be noted that “quanto,” or currency-protected features, are included in other derivative-related packages. For instance, an investor seeking exposure to an international stock market but fearing exchange rate risk, can purchase an index contract with a quanto feature that eliminates any currency gains/losses.

are aligned properly. Thus, rather than simply tapping a domestic market, a firm might be able to access an offshore market through the issuance of offshore zero coupon debt coupled with a currency swap; this can produce synthetic home currency funding at favorable rates. For example, a US company may be able to issue a five-year C\$ zero coupon bond to investors and then swap the proceeds into US\$. In fact, the company pays the swap intermediary an upfront C\$ amount in exchange for US\$ and then receives the full zero coupon amount at maturity; the full amount is then repaid to investors. Though the company receives no intervening cash flows on the swap over the five-year period, it is not obligated to pay its investors any interim cash flows either. The end result can be a lower cost of funding on the combined zero coupon/swap structure.

Superfloater swap

A combination of an interest rate swap, the purchase of a cap, and sale of a floor, with the cap and floor notional amounts adjusted to reflect the desired level of leverage. Under this package of derivatives the borrower's cost of funds increases as rates fall below the band defined by the floor, and decreases as they rise above the band defined by the cap (e.g. as rates decline, the reduced floating-rate receipts increase the effective swap cost; as rates rise, the higher floating-rate receipts decrease the effective fixed swap cost). The introduction of leverage accelerates the degree to which costs can increase or decrease.

Dual coupon currency swap

A combination of an interest rate swap and a strip of currency options, which permits the borrower arranging the transaction to achieve a lower all-in funding cost payable under the swap (via the sale of the options). The expiry date of each option coincides with the interest coupon payments under a swap-based borrowing, and the pro rata share of the premium reduces the fixed rate payable by the borrower under the swap. This type of structure is useful for borrowers that have receipts in multiple currencies, as they may be required to deliver a foreign currency flow in the event of exercise.

Index amortizing rate (IAR) swap

A combination of an interest rate swap and a strip of digital options that triggers a decrease in the notional principal as particular interest rate thresholds are crossed on the downside. In exchange for selling the digital options via the swap, the IAR buyer receives a higher fixed-rate coupon (e.g. a de facto premium). In order to preserve as much of the premium as possible, the buyer seeks a calm market, with rates trading in a relatively narrow range. As volatility rises, the likelihood of rates moving through the downside barrier increases. Note that the IAR can also be structured as a swap and a short straddle on rates.

Contingent premium option

A combination of a standard option and a deferred, contingent cash flow that represents the premium of the option; in fact, the contingent cash flow can be likened to a digital option that will trigger a premium payment from the buyer to the seller of the structure only if the underlying option is exercised. In exchange for the uncertainty associated with the digital

component of the package (i.e. whether or not the seller of the contingent premium option will ever receive a cash inflow), the theoretical premium charged is higher than that of a standard option. The higher premium also incorporates compensation for cash flow deferral (e.g. time value of money). The contingent premium option can be written on a range of references, and can also be applied in the swaption market.

10.3.3 Asset swaps, liability swaps, and callable/puttable asset swap packages

Asset and liability swaps have emerged over the past two decades as an integral component of the investor- and issuer-driven capital markets. Growth has been quite rapid, and innovation has appeared at regular intervals. Standard asset swaps and callable and puttable asset swap packages have proven to be equally popular, giving investors the opportunity to achieve very specific investment results, and intermediaries or investors the ability to either retain future upside potential or protect against downside risk on a particular asset swap. Liability swaps, such as those outlined briefly in Chapter 2, are a major source of swap market volume, giving issuers the ability to lock in preferred funding levels on newly issued or existing debt. In this section we consider the basic mechanics and applications of standard asset swaps and liability swaps, which form the cornerstone of the synthetic swap sector, and then analyze applications involving callable and puttable structures.

There are times when an investor seeks to invest capital in a specific asset class, currency, or rate. If assets exist that can fulfill the requirement, then the investor simply buys the desired asset by funding it through normal sources. However, if the precise characteristics are not available, then the asset can be created synthetically by structuring or repackaging the host asset with an interest rate or currency swap. Similarly, there are times when asset swaps can be used to crystallize excess spreads by taking advantage of market inefficiencies, illiquidity, and mispricings. We recall from our discussion in Chapter 2 that there is an embedded quality differential within the corporate capital markets, where low-quality issuers have a comparative advantage in the floating-rate markets, and high-quality issuers have a comparative advantage in the fixed-rate markets (and an absolute advantage in both markets). Depending on market conditions and the size of the quality differential, investors can purchase high-rated-fixed-rate bonds directly, or high-rated FRNs swapped into fixed rates; similarly, they can purchase lower-rated FRNs directly, or lower-rated fixed-rate bonds swapped into floating.

This approach is, of course, the OTC derivative equivalent of the structured notes discussed in Chapter 5, with the investor executing both an asset purchase and a swap contract with the intermediary (e.g. the swap bank), rather than purchasing the package in a securitized form with the derivatives already embedded. By doing so, the investor assumes the credit risk of the intermediary on the swap flows, which adds an additional dimension of risk exposure.¹⁰ For instance, if an investor seeks to invest in Company XYZ's bonds on a floating-rate basis, but XYZ only has fixed-rate bonds outstanding, the intermediary can sell the investor the fixed-rate bonds and attach a swap that converts the fixed-rate coupons into LIBOR. In practice, the swap flows are often adjusted¹¹ to provide investors with a synthetic investment trading at par; this adds a dimension of transparency to the process. Though no specific floating-rate XYZ asset actually exists, the intermediary effectively synthesizes one by combining the fixed bond

¹⁰ The process is also manually more intensive, as the investor is required to measure and track the exposures on both the underlying bond and the swap, and may be required to account for the cash flows in different ways.

¹¹ The adjustments may relate to premium/discount on the underlying asset, as well as any differences in accrued interest.

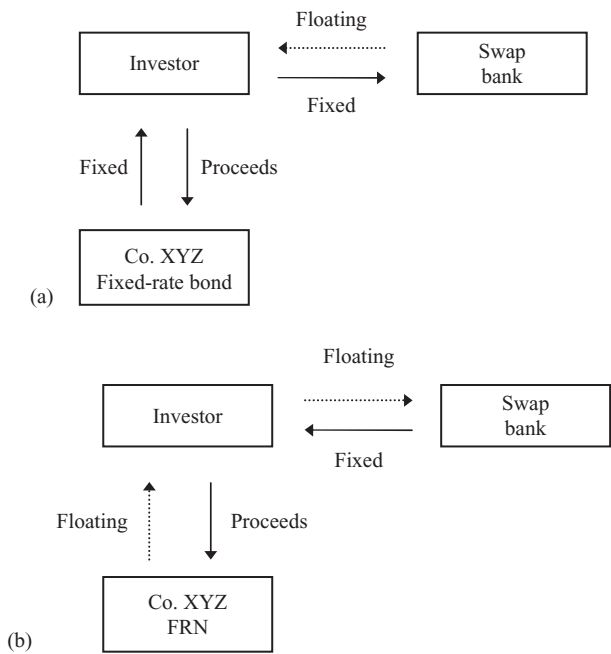


Figure 10.21 Asset swap packages. (a) Synthetic FRN investment package; (b) Synthetic fixed-rate investment package

and the interest rate swap. The reverse transaction can also be created: if XYZ only features LIBOR-based liabilities and an investor seeks a fixed-rate return, it can enter into an associated swap, where it receives fixed and pays LIBOR. The periodic flows on these fundamental asset swap packages (i.e. after the investor has purchased the asset from the swap bank and entered into the associated swap) are highlighted in Figure 10.21. Asset swaps centered on basis flows (e.g. CP versus LIBOR, EURIBOR versus EONIA) and nonstandard/amortizing/zero coupon cash flows can also be created.

Asset swap activity is driven by a number of factors, including arbitrage opportunities that permit the development of a synthetic investment at a higher yield, and restrictions that bar the creation of conventional fixed-income supply (or which prohibit investors from acquiring certain classes of assets). Arbitrage is a particularly important influence, much as the liability credit arbitrage drives the standard liability swap market. In fact, the arbitrage differential between the asset yield and swap levels must be large enough to provide economic benefits to all parties. The differential is a function of supply/demand for fixed versus floating issues, asset liquidity, new issue supply, and credit spread differentials between strong-investment-grade and weak-investment-grade/sub-investment-grade issuers. For instance, an investor seeking to invest in LIBOR + x bp assets that are funded at LIBOR can do so on an outright basis, assuming it can find the correct floating-rate assets and its funding levels are at LIBOR or sub-LIBOR. But it may also be able to take advantage of an arbitrage opportunity, and use an asset swap to achieve the same goals. If fixed payments made under a swap are less than the payments received from a fixed asset, then a LIBOR + spread is created; if this spread is greater than the spread on a direct FRN purchase (funded at LIBOR flat or sub-LIBOR), then

the asset swap arbitrage works. For example, if an investor can receive 5.5 % on a three-year, fixed-rate bond, and LIBOR + 15 bps on a three-year FRN, and can enter into a three-year pay fix/receive LIBOR at 5.25 %, the net return on the synthetic asset is LIBOR + 25 bps, or 10 bps more than the direct purchase. The arbitrage in this example generates the asset swap activity. There are, in fact, certain simple rules that allow for evaluation of fixed- and floating-rate asset swap opportunities:

- If the fixed rate bond spread over Treasuries¹² > swap spread + FRN spread over LIBOR, then an investor should purchase a conventional fixed-rate bond directly.
- If the fixed-rate bond spread over Treasuries < swap spread + FRN spread over LIBOR, then an investor should purchase an asset-swapped (synthetic) fixed-rate bond.
- If the FRN spread over LIBOR > fixed-rate bond spread over Treasuries – swap spread, then an investor should purchase a conventional FRN directly.
- If the FRN spread over LIBOR < fixed-rate bond spread over Treasuries – swap spread, then an investor should purchase an asset-swapped (synthetic) FRN.

Assets used in the asset swap process may include standard FRNs and fixed-rate bonds, illiquid/mispriced fixed-rate bonds, ex-warrant bonds with low coupons, ABS/MBS, and new issues floated expressly for asset swapping (e.g. opportunistic flotation). Illiquid bonds, in particular, are often repackaged via the asset swap market in order to improve marketability; this is generally true of bonds that have fallen out of favor,¹³ those that have lost a critical mass of interest post launch, and those hampered by poor new issue launch. Bond illiquidity of this type (which becomes very evident in comparison to liquid benchmark securities) can give rise to asset swap arbitrage opportunities. In fact, intermediaries that help arrange such issues, or repurchase them on a secondary basis, generally have dedicated teams that search for ways of repackaging, via swaps, securities held in inventory. Structured assets, including ABS/MBS or notes with a complex series of embedded derivatives,¹⁴ can also be a good source of arbitrage, though the opportunities tend to be more difficult to identify and analyze. New issues launched by lower-rated credits are also a significant source of supply: bonds issued on a fixed-rate basis are swapped immediately into floating rates by a bank for placement with investors (e.g. a direct link between new issue liability swaps and asset swaps). In fact, asset-swap-driven new issues can often be detected by the relatively small flotation size: deals of \$50–\$100 m are often targeted asset swap placements, while those of larger size typically are intended for general distribution. New issue asset swaps generally are intended for sophisticated institutional investors that are capable of managing the multiple cash flows, accounting issues, and credit risk exposures of the asset and swap components, and can help broaden an issuer's overall investor base. Investors lacking sophistication are often encouraged by intermediaries to obtain the desired synthetic results via a repackaged security or structured note (which we have already noted provide consolidated cash flows, feature greater liquidity, and eliminate the swap intermediary credit risk). An “intermediate” version of the structure also exists, where the investor pays the bank principal for the notes, and then receives note and swap cash flows from the bank; the investor faces the bank on the swap, instead of an SPE issuer, and must

¹² Or some other relevant risk-free government benchmark security.

¹³ Perpetual FRNs, for instance, were issued aggressively by banks in the 1980s, but ultimately fell out of favor as bank capital rules changed; the supply of outstanding assets has been used in the asset swap market since that time.

¹⁴ Indeed, some of the structured note variations we discussed in Chapter 5 have proven to be good asset swap candidates. Instruments such as leveraged inverse floaters, dual-currency bonds, reverse dual-currency bonds, and so forth, are often repackaged into standard synthetic FRNs through asset swap technologies.

still deal with credit risk issues related to the bank as intermediary, but receives all cash flows in a combined manner rather than from separate sources (e.g. coupon and principal from the bond issuer, swap payments from the bank). It is also worth noting that asset swap packages are designed to be held until maturity of the underlying bond/swap combination. If an investor chooses to exit early, it will do so at the current mark-to-market value and pay a second bid-offer spread (the first occurring, of course, at inception). The same is not necessarily true of structured notes, which have a greater amount of liquidity.

We summarized the mechanics of liability-driven credit arbitrage in Chapter 2, and expand on the process in this section. The asset swap mechanics described above can be used to create a liability swap package for a corporate or sovereign issuer seeking to raise funds in a marketplace that might not otherwise be readily accessible. In fact, this type of funding credit arbitrage has given rise to substantial growth in the global interest rate swap market over the past three decades. We have noted previously that weaker issuers (e.g. low-investment-grade/sub-investment-grade credits) have a relative advantage in raising floating-rate funds, while stronger issuers have a relative advantage in raising fixed-rate funds. This gives rise to liability swap opportunities: if a weaker issuer wishes to lock in fixed-rate funding but cannot do so through the fixed-rate capital markets (e.g. investors may refuse to lend on a fixed-rate basis, pricing for a fixed-rate security may be too high, and so on), then it can structure synthetic fixed-rate funding by combining its FRN issuance with a fixed/floating swap, as illustrated in Figure 10.22(a). A strong issuer, in turn, may combine a fixed-rate bond and a fixed/floating




Figure 10.22 Liability swap packages. (a) Synthetic fixed-rate funding package; (b) Synthetic floating-rate funding package

Table 10.5 Liability swap package

LIBOR = 4.75 %	Inflow/outflow
Pay LIBOR + 100 bps on loan	−5.75 %
Receive LIBOR on swap	+4.75 %
Pay fixed on swap	−6.00 %
Net all-in funding cost	−7.00 %
Hypothetical fixed rate bond	−7.25 %
Annual savings vs. fixed rate bond	+0.25 %

swap to generate synthetic floating-rate funds at price levels that are cheaper than a straight FRN issue; this is depicted in Figure 10.22(b).

Let us consider a simple example involving the creation of synthetic fixed-rate debt using a swap. Assume that a company can borrow on a floating-rate basis from a bank at LIBOR + 100 bps for three years, but wishes to lock in a fixed rate. If it issues a three-year fixed-rate bond, it will need to pay 7.25 %. However, the three-year swap market is quoted at 6 % versus LIBOR flat, so the company opts to borrow from the bank at LIBOR + 100 bps and enter into a swap where it pays 6 % fixed and receives LIBOR. Assuming current LIBOR is 4.75 %, the net cash flows on the financing, summarized in Table 10.5, reveal a net fixed financing cost of 7 %, which is 25 bps lower than a straight three-year fixed-rate bond issue.

Naturally, liability swaps can also be used to restructure cash flows on existing liabilities, rather than new or contemplated liabilities. This can be a benefit in a changing interest rate environment, allowing the debt obligor to alter its cash flow profile in response to actual or anticipated rate changes.

We now extend our discussion by considering callable and puttable asset swap packages, and asset swap switches. A callable asset swap (also known as a remarketable asset swap) is similar to the asset swap package described above, except that the swap bank selling the package to the investor retains a call option on the underlying fixed- or floating-rate asset, allowing it to repurchase the asset at a given spread at some future time. The total package can thus be seen as a combination of a callable swap (itself a package of a swap and an option) and an underlying fixed- or floating-rate bond. The option generally has a European or Bermudan exercise that is synchronized with the coupon dates of the asset, and the spread (i.e. the strike level) is generally equal to the investor's initial purchase level. If the spread on the asset tightens during the life of the transaction (e.g. the price of the asset rises as a result of specific or general market/credit conditions), the swap bank calls the package away from the investor; the investor receives proceeds equal to the strike spread plus invested principal from the swap bank. The swap bank can then sell the underlying asset in the marketplace at a profit, or enter into a new callable asset swap with a new investor at the tighter market spread; this process helps realize mark-to-market value on the deal. Naturally, if the spread widens, the swap bank will not exercise the option and the investor will preserve the package until the contracted maturity date. The investor, in exchange for giving the swap bank the right to call the package away, receives an incremental yield representing the premium from selling the option. This synthetic structure gives both parties specific benefits: the swap bank preserves the ability to crystallize value efficiently by liquidating or reselling the structure if asset spreads tighten, while the investor earns an incremental yield for granting the option. The callable asset swap structure (based on a fixed-rate bond) is summarized in Figure 10.23.

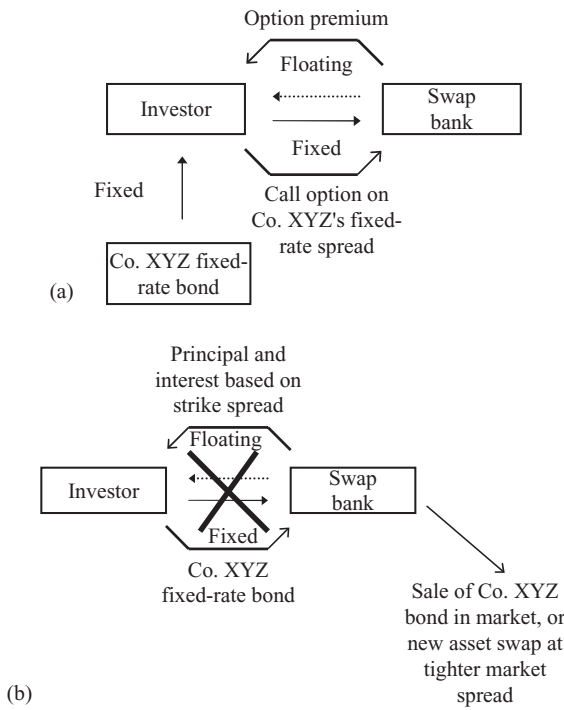


Figure 10.23 Callable asset swap. (a) Initial and ongoing flows assuming no exercise; (b) Terminal flows assuming swap bank exercises

A puttable asset swap functions in a similar manner, except that the investor, rather than the bank, enjoys the benefit of the option feature. Specifically, the investor acquires a package comprised of an asset (e.g. fixed rate bond or FRN) and a puttable swap, where the put feature allows the investor to sell the entire asset/swap package back to the bank at a predetermined strike spread. As with the callable structure above, the put option generally has a European or Bermudan exercise that is synchronized with the coupon dates of the asset, and the strike spread typically is set equal to the investor's purchase spread. If the spread widens during the life of the transaction (e.g. the price of the asset falls as a result of specific or general market/credit conditions), the investor puts the package to the bank, receiving principal and interest defined by the strike spread. The swap bank may then retain the asset in its portfolio, sell it in the marketplace, or attempt to arrange a new asset swap structure with another investor at the new (albeit wider) spread level. Depending on carrying value, the bank may or may not post a mark-to-market loss. If the spread tightens, the investor will not exercise the option and will thus continue to preserve its package until the contracted maturity date. The investor, in securing the right to put the package back to the bank, pays an option premium; this may be in the form of an upfront option payment or a lower yield on the asset swap coupon. As above, both parties obtain benefits through this synthetic structure: the swap bank receives incremental income from selling the put option (and either temporarily or permanently removes the underlying asset from its balance sheet), while the investor obtains de facto downside protection against spread widening on the underlying asset. Note that the puttable asset swap structure is also an essential

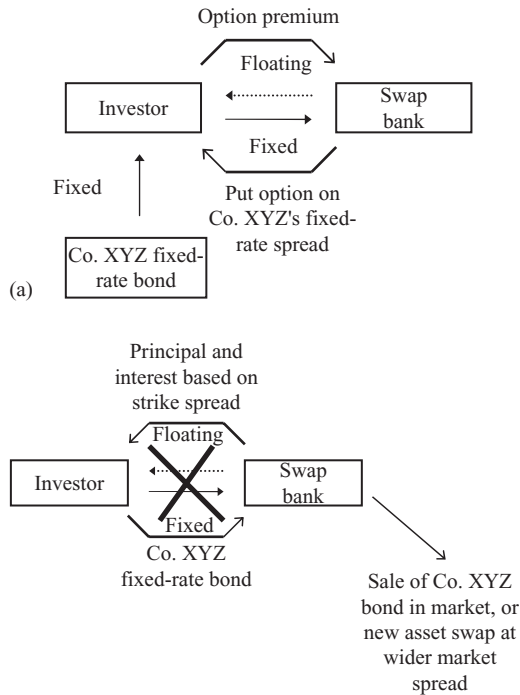


Figure 10.24 Puttable asset swap. (a) Initial and ongoing flows assuming no exercise; (b) Terminal flows assuming investor exercises

element of callable bond investing, as investors must have comfort that they can terminate the swap component of a callable/puttable asset swap strategy if the underlying bond is called back by the issuer. Absent this feature, investors would be left with naked swap positions upon exercise of the bond call. The puttable asset swap structure (based on a fixed-rate bond) is summarized in Figure 10.24.

Asset swap switches are another form of structured asset swap packaging. Asset swap switches, as the name implies, involve the exchange of two different asset swap packages. Under the most basic version, an investor acquires an asset swap package from a bank or financial intermediary, and simultaneously agrees to deliver the package and accept another one in return if the spread on the second package widens to a particular level; the two asset swap packages generally are uncorrelated. By agreeing to “swap” the two packages, the investor receives an enhanced yield on its investment. For instance, an investor owns an asset swap package on credit reference ABC at LIBOR + 30 bps, and agrees with Bank D to exchange it for an asset swap package on XYZ, currently trading at LIBOR + 50 bps – but only if XYZ’s trading spread widens to +70 bps. Bank D essentially has the right to put the XYZ package to the investor, while simultaneously calling the ABC package. In order for this transaction to function from an economic perspective, the investor and the bank must have different views on the current and future creditworthiness of ABC and XYZ. For instance, the investor must believe that XYZ represents good value at +70 bps, but is too rich at +50 bps. The bank clearly is exposed to 20 bps of spread widening before it can trigger the swap, but is able to acquire default protection through the structure.

While puttable swaps, callable asset swaps, and asset swap switches are the primary vehicles of the asset swap market, new products have been introduced in recent years. One popular asset-swap-related structure is the synthetic lending facility, commonly known as an asset swaption. Under this synthetic contract, the investor earns an upfront or periodic fee for agreeing to purchase a loan from an intermediary on a forward basis. This is equivalent to the investor selling a put option to the intermediary on a revolving credit line or note underwriting program arranged for a third party credit borrower/issuer; by accepting the payment, the investor essentially is participating in an unfunded revolver/note program until it is funded or terminated. If the intermediary exercises the option, the investor is obliged to purchase the reference asset on an outright or asset-swapped basis. An associated product is the synthetic credit facility, which functions in the same way, but can include any reference debt instrument rather than just revolving loans.

10.3.4 Credit derivatives and synthetic credit positions

In Chapter 5 we considered the development of credit linked notes, or securities embedded with one or more credit derivatives. In Chapter 6 we discussed the use of credit derivatives in the creation of certain synthetic CDOs. In this section we expand on the role of credit derivatives in the area of synthetic activity, by considering how the fundamental instruments can be used to replicate single credit risky positions or broader risky portfolios.

Credit swaps, options, and forwards

The credit default swap (CDS) has become the single most popular derivative instrument in the credit markets. Indeed, CDSs on major credit issuers have become extremely liquid over the years, to the point where many end-users and intermediaries prefer dealing in the derivatives rather than the underlying obligations; this is particularly true for those attempting to establish short positions in a reference credit, where it can be difficult to borrow and sell the physical security as a result of supply or regulatory issues. CDSs are also useful for portfolio diversification and rebalancing efforts, as they can reshape risk or investment exposures in an efficient and cost-effective manner.

We recall from earlier in the text that in a CDS, the credit-protection buyer pays the credit-protection seller an upfront or periodic fee in exchange for a compensatory payment if the reference credit defaults during the life of the contract;¹⁵ if no default occurs, no payment is made by the credit-protection seller. This is equivalent to the buyer purchasing a digital credit put option,¹⁶ with a payout equal to 0 (if no default occurs), or a post-default amount defined by market quotes (if default occurs); note that there is no continuum of payouts between the default/no-default scenarios as in the credit spread option considered below. This basic arrangement is reflected in Figure 10.25. A typical CDS can be settled in cash or physical terms: under cash settlement, the credit-protection seller pays the buyer a sum equal to the post-default price of the reference obligation (generally obtained from a panel of market dealers) times the notional amount; under physical settlement, the buyer delivers the notional reference obligation (which it may have in its portfolio, or which it sources from the market at post-default prices) and receives par value.

¹⁵ In fact, a quarterly payment of the fee, or premium, by the credit seller to the credit buyer is the market standard.

¹⁶ The reduced loss credit default put option is a variation on the theme, where the buyer of the option agrees to absorb a portion of the loss if a default occurs, in exchange for payment of a smaller premium. The first loss structure can be seen as a form of deductible.




Figure 10.25 Credit default swap

The credit spread option is the second major credit derivative contract, sharing similarities with the CDS (or equivalent credit default put option). Like the CDS, the credit spread option is a unilateral contract that provides the purchaser with a compensatory payment based on a credit event; unlike the CDS, however, the option provides a payment that is based on spread movements as well as default.¹⁷

In a standard spread option, the buyer pays the seller a risk premium in exchange for a compensatory payment if the reference credit spread widens beyond the strike level (for a put) or tightens inside the strike (for a call). A put option on the credit spread therefore gives the buyer the right to sell the spread at the strike, creating a gain as the reference credit deteriorates and its spread widens out; in the extreme, a defaulting credit will feature a very wide credit spread, so the credit spread put provides default protection. A call option on the spread, in contrast, gives the buyer the right to purchase the spread at the strike, meaning a gain as the reference credit improves and the spread tightens. The spread may be defined in terms of a floating-rate reference, such as LIBOR, or a fixed-rate risk-free benchmark, such as Treasuries. Assume an institution purchases a credit spread put option at a strike spread of Treasuries + 100 bps. If the reference credit is downgraded from A to BBB–, the spread on the bond may widen to $T + 200$ bps; though the credit is not in default, the institution is due the 100 bps spread over the strike spread (times the credit spread duration of the underlying bond reference¹⁸). Accordingly, the spread option can be viewed as a conventional credit put that provides the buyer with a full continuum of payoffs.

Credit spread options can be used to capture relative value changes in spreads independent of interest rates, or express a view on forward credit spreads¹⁹ or the term structure of spreads. For instance, an investor believing an asset is overvalued but that default is improbable can sell a credit spread put option on the asset, earning premium in the process. If the spread

¹⁷ The credit spread option has its foundation in MBS/US Treasury option strategies, where investors often sell calls on MBS and buy calls on US Treasuries, in order to protect against a general widening of MBS spreads. The original strategy ultimately expanded into other areas, such as spreads between Mexican government obligations and US Treasuries, and has now become an established part of the credit derivative market.

¹⁸ The credit spread duration, which measures the sensitivity of the security's price to a movement in spreads, captures the effects of price changes over the life of the security on a present value basis.

¹⁹ The forward credit spread can be computed by first determining the spot price of the risky bond and the associated risk-free benchmark; the forward prices of both instruments can then be computed and converted into forward yields; the forward spread is simply the differential between the two. It is worth noting that forward credit spreads may not reflect investor expectations and are notoriously volatile.

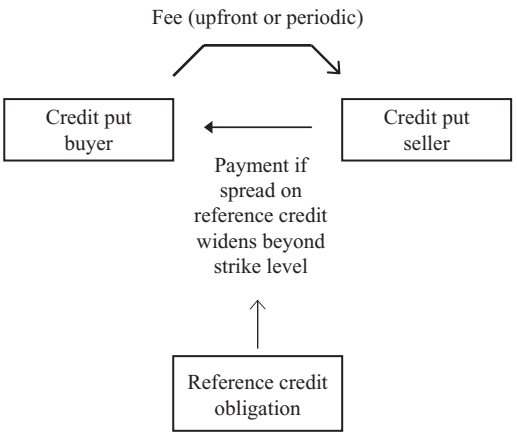


Figure 10.26 Credit spread option

widens to the strike, the investor will have to purchase the asset from the option buyer, but may be comfortable doing so as its target spread level has been reached; in addition, it will have generated premium income to supplement the position. Apart from these standard structures, exotic spread options, including knock-in and knock-out options, are available. For instance, in a knock-out credit spread put, the spread put is eliminated if the reference credit spread tightens by a certain number of basis points. A knock-in put, in turn, becomes active once spreads widen to a certain level, meaning the put seller must take delivery of the asset at the wider spread level, and the put buyer gains effective default protection. Since knock-ins and knock-outs do not provide for a full spectrum of payouts, they are cheaper than standard spread options.

Vanilla credit spread options are bought and sold on a range of reference credits, and can be structured with American or European exercise involving physical or cash settlement. Transaction maturities range from six months to over five years, and dealing size is typically \$10 m to \$100 m. Figure 10.26 highlights the cash-settled credit spread put structure.

The credit spread forward (or credit spread swap) is the third fundamental structure of the credit derivatives marketplace, and can be viewed much as any of the other forward contracts introduced in Chapter 2. The credit spread forward is a forward on a risky credit spread, computed in absolute or relative terms. Under the absolute spread version, the seller pays a fixed spread and receives a market spread (compared to a risk-free benchmark); the fixed spread is set at trade date and the evaluation versus the market spread occurs at transaction maturity. The buyer, by extension, pays the market spread in exchange for the fixed spread; this structure is illustrated in Figure 10.27. For instance, if a bank agrees to pay 100 bps fixed for the spread on a reference bond of Company ABC, and receive ABC’s market spread, it will generate a gain as ABC’s market spread widens beyond 100 bps, and will post a loss as it tightens. The bank’s counterparty will face the opposite scenario. The relative spread version functions in a similar manner, except that one party pays the spread over the risk-free benchmark on reference asset 1, while the second party pays the spread over the risk-free benchmark on reference asset 2; this removes the risk-free benchmark from the equation to focus solely on the spread differential between two credit-risky references. In practice, asset swap spreads are often used as references for credit spread forwards, even though LIBOR is not strictly a risk-free benchmark rate. While




Figure 10.27 Absolute credit spread forward

most forwards are structured for single periods, they can be extended into multi-period swap form.

Total return swaps

The total return swap (TRS) is similar to the credit forward, transferring the economics of a credit reference between two parties on a bilateral basis. The TRS receiver is entitled to the total return on the reference asset (appreciation plus coupons) in exchange for periodic floating payments (e.g. a LIBOR-based flow plus a spread); this is equivalent to creating a synthetic long position in the reference asset (indeed, the TRS is often considered a form of synthetic financing). The TRS payer, in turn, is entitled to any depreciation that occurs in the security, along with a LIBOR spread. Since the TRS payer often owns the underlying reference asset, it is transferring the economic risk of the asset and effectively is creating a synthetic short position (or first loss credit protection).²⁰ However, since the TRS payer still owns the position, it must continue to fund it through its own balance sheet sources.²¹ If the price of the reference asset rises, the receiver benefits from the appreciation; if it declines as a result of depreciation or default, the payer benefits. The TRS can be arranged on a funded or unfunded basis. A funded TRS requires the TRS receiver to purchase a risk-free or low-risk floating rate asset that yields the LIBOR stream payable to the TRS payer. No asset collateral is held in an unfunded TRS, so the receiver must source the LIBOR stream from cash flow; this also means that the position in the TRS is highly leveraged.

The standard TRS structure is defined in terms of a reference asset (e.g. a credit), notional amount, maturity, payment frequency, initial reference asset price, current asset price, quote method, and final settlement price determination. Though the structures generally are quite transparent, challenges arise when the underlying reference asset is illiquid and meaningful price quotes are difficult to obtain. The frequency and timing of payment calculations is negotiated between both parties at trade date. The “total return” referenced under terms of the transaction includes all interest payments on the reference asset, plus an amount based on the change in the asset’s market value. At maturity, the TRS receiver can buy the reference asset at

²⁰ It is worth noting that in transferring the price appreciation/depreciation, the TRS mechanism transfers economic changes that are due to both interest rate movements and credit spread movements; the two are not separated, meaning that the vehicle is as much an interest rate risk tool as a credit risk tool.

²¹ In fact, the net profit to the TRS payer is the basis point spread over the funding cost, times the notional amount of the transaction. Thus, if the payer receives LIBOR + 10 bps from the receiver, and can fund at LIBOR flat, it earns 10 bps running on the notional of the transaction.

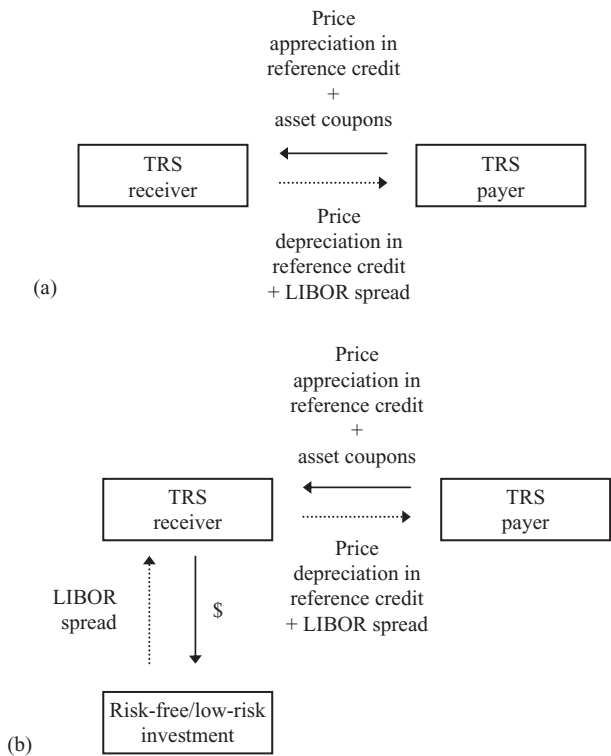


Figure 10.28 Total return swap. (a) Unfunded; (b) Funded

the prevailing market price. If a default occurs prior to maturity, the receiver makes the payer “whole” via a net payment equal to the difference between the asset price on trade date and post-default (cash settlement basis). Alternatively, it can take delivery of the asset following default by paying par value (physical settlement basis). Figures 10.28(a) and (b) summarize the flows of unfunded and funded TRSs.

Basket swaps

Basket swaps, which are derivative contracts comprised of a number of CDSs, are available in a number of different forms, including the standard basket swap, first (*n*th) to default basket swap, and subordinated basket swap.

Standard basket swaps

The standard basket swap, the most common of the synthetic credit portfolio structures, is created when the basket seller packages for the buyer a portfolio of several reference credits, and provides the buyer with a compensatory payment if any reference credit defaults; if every credit in the basket defaults, the buyer receives a payment for each one. For instance, if the buyer specifies a basket of four reference credits of \$10 m notional each, and three of the four

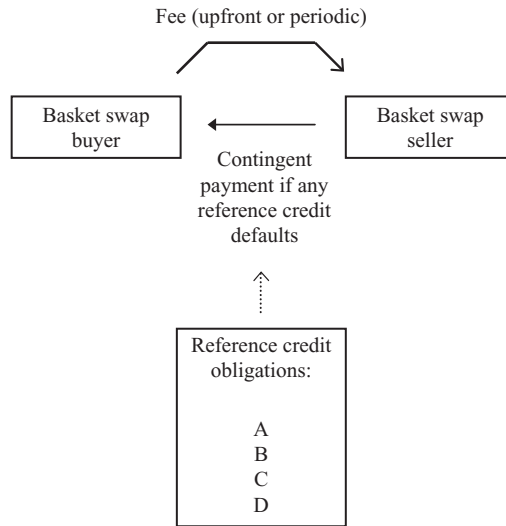


Figure 10.29 Basket default swap

default, it will receive a compensatory payment from the seller equal to the post-default price of the defaulting credits times \$10 m notional per credit. Note that the seller makes no payment if the reference credits deteriorate but do not default (e.g. credit spreads widen but the obligors remain current on their principal and interest payments); this helps reinforce the point that the basket swap is equal to a bundle of CDSs.

There is considerable flexibility in assembling the reference portfolio: a basket can include 2 to 20+ obligations, industries can be identical or different, country exposure can be single or multiple source, ratings can range from investment-grade to sub-investment-grade, and so forth. Maturities can range from six months to several years, and deal size can grow to several hundreds of millions of dollars. As with any individual CDS, the basket buyer pays the seller a fee; however, the inclusion of multiple reference credits can lower the overall fee payable – though this depends ultimately on the construction of the basket and the correlation between the reference credits. Baskets with credits that are highly correlated are unlikely to feature much/any premium reduction, while those that are uncorrelated may be somewhat cheaper. Figure 10.29 highlights a standard basket swap.

First and n th to default basket swaps

The first (or n th) to default swap is a more refined version of the basket swap. The structure again involves assembling a portfolio of related or unrelated reference credits. In this instance, however, the buyer of the basket only receives a compensatory payment when the first (or n th) credit in the portfolio defaults. Once the specified default event occurs, the structure unwinds and the seller faces no further liability if another default occurs. For instance, assume the basket contains ten reference credits and the buyer receives a payment based on the first credit to default (which of the ten actually defaults is irrelevant – they are all considered equally). Once the first default occurs, the buyer receives the notional amount of the basket times the

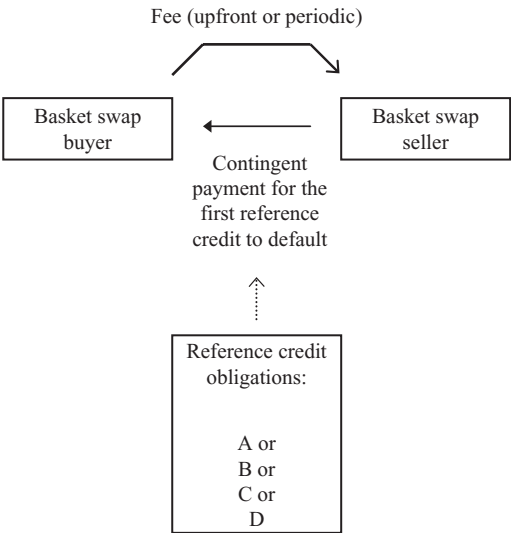


Figure 10.30 First to default swap

post-default price of the defaulting credit and, even though nine performing credits remain in the basket, the transaction terminates. Indeed, if a second default occurs one day after the first default, the seller is not obligated to make a compensatory payment. Though the first to default basket shares certain similarities with the standard basket swap (e.g. assembling a portfolio of desired reference credits) the fact that the payout is limited to a single default event means that the overall fee paid for the protection is much smaller. In some instances, the structure is set to accommodate the 2nd, 3rd, 4th, or n th, credit to default; this means $(n - 1)$ defaults must occur among the portfolio obligations before the buyer receives a compensatory payment. Since this is obviously a much lower probability event, n th to default structures are considerably cheaper than standard basket swaps and first to default swaps. Figure 10.30 illustrates the flows of a first to default swap.

It is worth noting that a first to default basket is similar to the standard CDO we discussed in Chapter 6. Specifically, the basket seller is akin to the investor in the equity tranche (or subordinated tranche, depending on how the transaction is structured), bearing the first loss if default occurs. However, the basket seller has an off-balance-sheet exposure, while the CDO investor has a funded balance sheet investment; in addition, the basket features a narrower reference portfolio than the CDO.

Senior and subordinated basket swaps

Senior and subordinated basket swaps are two additional variations on the basket structure. As above, these basket swaps essentially are customized synthetic credit portfolios featuring a number of reference credits. The payment characteristics are, however, unique. In the senior basket swap, the seller pays the buyer for any credits that default in the portfolio (and in that sense is identical to the standard basket swap), but will not do so until a minimum loss threshold has been reached – this creates a first loss/deductible structure. For instance, the reference portfolio might feature five credits with \$10 m notional apiece and a first loss level

of \$10 m. If all five credits default during the life of the structure, and the post-default price for each is 40 (again, for simplicity), then the compensatory payment on a standard basket is \$30 m. However, under the senior swap structure, the buyer absorbs the first \$10 m of losses as a form of deductible, meaning the compensatory payment nets to \$20 m.

The subordinated basket swap is also comprised of a customized portfolio of reference credits, providing a payout on every credit that defaults; however, it also features a cap, or aggregate limit, on the total amount of payout that can be made over the life of the contract. Thus, assuming we have a \$25 m subordinated basket contract based on the same five-credit, \$10 m basket, the total payout to the buyer will be \$25 m – even if all five credits default and the post-default price on each is 40. In this instance, the seller of the subordinated structure bears the first loss piece up to a maximum cap.

Risk, Legal, and Regulatory Issues

11.1 INTRODUCTION

The structured and synthetic products we have discussed in this book are risky instruments that must be controlled properly from both an internal and external perspective. It is essential for intermediaries creating the products, and end-users employing them, to have a solid grounding in the control-related dimensions of each product; failure to understand the risk and control dimensions of the products can lead to instances of unexpected loss.

We summarized in Chapter 2 the different elements of risk that can impact instruments in the marketplace, including market, credit, liquidity, legal, and operational risks. In this chapter we consider the control framework applicable for each one of these risk types by focusing on risk management, financial, audit, and legal control mechanisms. We also discuss the specific regulatory and accounting treatment relevant to products in the sector. As we shall note, these are often unique and can have a decided impact on an institution's regulatory capital and financial statement reporting.

11.2 RISK AND FINANCIAL CONTROLS

Institutions dealing with structures and derivative products on a regular basis must build and maintain a proper internal control framework that permits prudent management of risks. Internal controls can take various forms, but tend to relate primarily to credit and market risk management, independent financial and operational processing, and internal auditing. The combination of the three, working in a synchronous fashion, can create a more secure dealing environment and reduce the likelihood that an institution will sustain unexpected losses.

11.2.1 Market, liquidity, and credit risk management

Internal credit and market risk managers often serve as the “front line” of controls, enforcing a series of standards that are intended to keep an institution's risk-sensitive operations in balance. Credit and market risk managers generally are involved in establishing limits and other controls for synthetic and structured assets (as well as other risky transactions) that link directly to an institution's risk philosophy (i.e. the nature of the risk-related business it wants to engage in) and risk tolerance levels (i.e. the dollar amount of capital it wants to place at risk). The control framework may also include formal “new product” reviews that examine the structural nuances and risks of new instruments. Though the remit of most risk departments is often quite wide, we can summarize a series of minimum risk controls that should exist for those dealing in synthetic/structured products. The market risk framework requires:

- identifying all market risks impacting structured and synthetic products, including:
 - directional risk (equity, currency, commodity, interest rate);
 - basis risk;
 - spread risk;

- volatility risk;
- correlation risk;
- curve risk;
- liquidity risk;
- quantifying all risks arising from creating or purchasing structured and synthetic assets;
- establishing meaningful risk limits for the relevant risk exposure classes, with a direct link to the institution's stated risk-tolerance level and the potential returns that can be earned;
- monitoring exposures on a continuous basis to ensure that exposures generated remain within an institution's risk-tolerance limits, and making adjustments as necessary;
- considering new products proposed by external parties or internal originators and ensuring they meet the institution's risk criteria.

Similarly, the credit risk framework requires:

- identifying all credit risks related to synthetic and structured assets, including:
 - counterparty credit risk;
 - issuer default risk;
 - correlated credit risk;
- quantifying all risks arising from the products;
- establishing limits for the net credit exposures the institution is willing to assume, and ensuring some relationship to the stated risk tolerance and the potential returns that can be earned;
- monitoring exposures on a continuous basis to ensure that exposures remain within an institution's risk-tolerance limits.

These types of process should be reviewed for efficacy on a regular basis (e.g. every year); they must be related specifically, and explicitly, through an institution's governance process. Reporting must be frequent and sufficiently detailed to provide directors, executives, and external parties with enough information on the firm's risk profile and its risk trends.

11.2.2 Internal financial and operational controls

Internal financial and operational controls exist in order to track and verify transactions that impact an institution's balance sheet, income statement, and cash flow statement. The duties of financial and operational professionals cross important boundaries that affect the front-, mid-, and back-offices, as well as executive management and external regulators.

Minimum internal controls should include:

- ensuring pricing feeds for synthetic and structured assets come from an independent source that cannot be manipulated;
- establishing reserves for products that appear impaired (e.g. illiquid, close to default);
- interpreting and implementing accounting policies related to structured products;
- making certain that the technology platform/trade entry screens include the entire population of daily dealings, in order to avoid any breaks/fails/settlement problems or financial fraud;
- reconciling activity and positions in order to generate a link to the firm's profit and loss (P&L) account, and its official books and records;
- gathering independent pricing valuations of synthetic and structured assets in the marketplace to ensure that the valuation policy the institution follows is equitable;

- creating independent risk management reports reflecting activities; these may be applied to risk limits supplied by the market and credit risk management departments;
- preparing executive management/board level revenue and risk reporting to demonstrate trends.

11.2.3 Internal audit

Virtually every major institution has some form of internal audit function to ensure the integrity of its operations. The typical audit function examines business and control units on a regular cycle, testing activities against established policies and procedures to ensure proper compliance and control. Deficiencies, weaknesses, shortcomings, or other potential problems typically are highlighted and elevated for resolution.

Units that assume or transfer risk through synthetic and structured assets should form part of the regular audit cycle. In addition, risk management and credit management units that promote the use of financial products must be reviewed regularly. Auditors focusing on synthetic and structured business must ensure the market and credit risk limits indicated immediately above are in force and effective in controlling exposures. Equally, they must verify the nature, quality, and accuracy of the pricing values/marks that the independent financial control units derive for the computation of daily P&L, position tracking, and books and records reconciliation. Any discrepancies must, of course, be resolved as a matter of urgency.

It is worth stressing again that much of this discussion is applicable to both financial intermediaries creating products and end-users (investors, issuers) purchasing/using the products. Though an end-user's processes may not be as extensive (given what is likely to be a much smaller scope of business), the same rationale and end goals apply, i.e. operating in a secure fashion with only a minuscule possibility of experiencing surprises or unexpected losses.

11.3 LEGAL CONTROLS

Prudent dealing in synthetic and structured products requires an examination of legal risks and the creation of procedures for ensuring that proper legal documentation is utilized. Such documentation, which must have a legal basis in a local dealing jurisdiction, should reflect the terms of trade between two counterparties, and establish the rights of each party.

The nature of legal documentation will, of course, vary by product. Structured products that are bought and sold on a cash settlement basis within trade date + 1 to 5 days (e.g. MBS, CDOs, convertible bonds, equity hybrids, ETFs, and so forth), generally are documented by standard trading confirmations that evidence relevant details of a trade, including:

- product type;
- notional amount;
- market price;
- settlement procedures/accounts;
- accrued interest;
- counterparty (buyer/seller).

These confirmations are relatively standard across the industry, and generic templates are often used to document essential terms. Progress within the electronic sphere to convey confirmation details via standard protocols is making the process more efficient and secure.

Legal matters relating to synthetic products created through one or more derivative contracts are more complex. In such instances, one or both counterparties to the synthetic trade (e.g. the intermediary and end-investor) may face more intensive legal and risk exposures; this is particularly true since transactions are likely to last for several months to several years – a period during which market references may move and create significant market/credit exposures, and when counterparty creditworthiness may weaken.

As a result of these factors, we focus additional attention on derivatives-related legal documentation contracted under the ISDA framework. Though other forms of documentation are employed (including, for example, the German Rahmenverstrag and the Association Française de Banque forms), these constitute a minority; accordingly, our emphasis is on elements of the ISDA process.

ISDA, which was chartered in 1985, is the global trade association representing participants in the derivatives industry. The organization has routinely pioneered efforts to identify and reduce sources of risk in the derivatives and risk management business; notable accomplishments include:

- developing the ISDA Master Agreement;
- publishing a wide range of documentation materials covering a variety of derivative classes;
- obtaining legal opinions regarding the enforceability of netting and collateral arrangements and securing recognition of the risk-reducing effects of netting in determining capital requirements;
- promoting sound risk management practices;
- advancing the understanding and treatment of derivatives and risk management from both a public policy and regulatory capital perspective.

These endeavors are consistent with its overall mission, which is to encourage the prudent and effective development of the OTC derivatives business by:

- promoting practices conducive to the efficient conduct of the business, including the development and maintenance of derivatives documentation;
- encouraging the establishment of sound risk management practices;
- educating members and others on legislative, regulatory, legal, documentation, accounting, tax, operational, and technological issues.

As a result of these efforts, the ISDA documentation framework has become the accepted industry standard for those dealing in the OTC derivatives market – including those dealing in structured/synthetic assets that contain derivative elements.

11.3.1 ISDA documentation framework

The general ISDA framework, which has been refined and expanded over the past two decades, is comprised of several component parts, any (or all) of which an institution may utilize as part of its legal procedures:

- Master Agreement
 - Printed form
 - Schedule
- Credit Support Documents
 - Credit Support Annex, Credit Support Deed
 - Margin Provisions

- Confirmations
 - long form
 - short form
- Definitions

The Master Agreement is comprised of the Printed Form, a multi-page document of standard terms and conditions that is not intended to be altered by the two parties to the agreement, and the Schedule, which is the attachment that the two parties negotiate and can enhance, alter, or customize at will (including inserting changes that alter the Printed Form). The Master Agreement serves as a mechanism to avoid separate, and lengthy, documentation of individual transactions by establishing, and then repeatedly referencing, standard terms and conditions. The terms and conditions relate primarily to legal and credit issues. The 1992 Master Agreement, which replaced the 1987 agreement, was itself replaced by a new version in 2002. The 2002 version, which is 50 % longer than its predecessors, reflects the industry's continued efforts to clarify and standardize as much as possible within the legal environment, while retaining the flexibility to accommodate innovation.

The framework also allows for use of Credit Support Documents, which are intended primarily to establish credit arrangements between two parties (e.g. the need for one or both parties to post collateral, actions to be taken in the event certain credit exposure thresholds are breached or a credit rating downgrade occurs, and so forth). The primary credit documents include the Credit Support Annex, the Credit Support Deed (used primarily under English law), and the 2001 Margin Provisions.

Confirmations, the third element of the framework, serve as evidence of individual transactions, and can be negotiated in long or short form. Long-form confirmations, as the name suggests, are detailed documents that contain all relevant aspects of the legal/credit terms of a transaction (much as the Master Agreement does), and generally are used when the product/market is new and lacks dealing standards, or when a transaction is so customized that it requires specialized documentation. Short-form confirmations, in contrast, are simple forms that serve primarily to document the economics of a new transaction; terms and conditions can reference the Master Agreement and associated Definitions. Short forms are used when the market is standardized and mature, and when product definitions are well established. Not surprisingly, the mechanism is efficient and reduces costs and operational errors.

Definitions are the last major component of the framework, and serve to describe standard terms and provisions for different product/market segments. Definitions are introduced under ISDA sponsorship when a market has reached a state of maturity where standardization is possible; standard Definitions have been issued for bullion derivatives (1997), government bond options (1997), euro derivatives (1998), currency options (1998), commodity derivatives (1993, 2005), equity derivatives (1996, 2002), and credit derivatives (1999, 2003). Short-form confirmations generally rely on, and reference, Definitions as part of the documentary process.

The general ISDA framework¹ is illustrated in Figure 11.1.

In the section that follows we briefly review certain clauses of interest that are contained within the Master Agreement; those interested should consult the detailed agreement.

ISDA Master Agreement

The 2002 ISDA Master Agreement (Multicurrency Cross Border) updates the 1987 and 1992 documents, expanding the breadth of products covered (thus promoting a greater degree of

¹ Note that the ISDA framework also includes certain bridges (cross agreements) and protocols (master change documents); we shall not consider these in further detail.

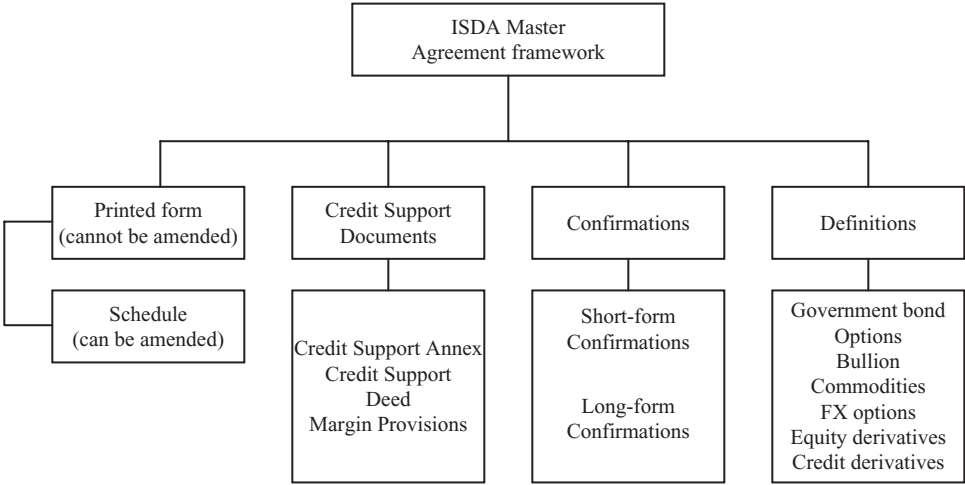


Figure 11.1 The general ISDA documentation framework

cross-product netting), addressing new legal developments that have appeared over the past 15 years, and clarifying important aspects of preceding versions. Key sections of note include the following.

Section 2: Obligations

Section 2(a) indicates that each transaction confirmation dictates payment terms for both parties, along with the timing and the method of payment. This section, which allows for physical deliveries, indicates that the payment obligations are subject to the conditions that no event of counterparty default has occurred, and that no early termination date has been designated.

Section 2(c) provides that payments may be netted, as long as they are due on the same date, are payable in the same currency, and pertain to the same transaction. Subparagraph (ii) of this section enables parties to elect the option of netting two or more transactions, provided that they occur on the same payment date and in the same currency.

Section 3: Representations

Section 3 allows both parties and any credit support provider (e.g. guarantor, letter of credit issuer) to make representations that are replicated automatically each time a new derivative transaction is arranged. Any misrepresentation comprises an event of default that enables the nondefaulting party to terminate the derivative agreement.

Subparagraph (a) contains basic representations that both parties declare, including those referencing:

- status;
- powers;
- lack of violation or conflict;

- consents;
- binding obligations.

Other subparagraphs contain representations about the absence of certain events of default and potential events of default, absence of litigation, and the accuracy of information identified in the schedule.

Section 4: Agreements

Section 4 contains a list of agreements that both parties are obligated to adhere to as long as either party has obligations under the agreement; this, in brief, includes the need for each party to:

- furnish specified information;
- maintain dealing authorizations;
- comply with applicable laws;
- comply with tax agreements;
- pay stamp taxes/duties as applicable.

Section 5: Events of default and termination events

Section 5 contains a list of events that constitute default or termination (note that these may be expanded or replaced by references to the 2003 credit derivative definitions).

Section 5(a): Events of default

The following constitute events of defaults:

- Failure to pay or deliver. This event applies to the failure of a party to make a scheduled payment or delivery under Section 2, three days after notice.
- Breach of agreement; repudiation of agreement. This event can exist as an event of default only in limited circumstances, and features a 30-day grace period following notice.
- Credit support default. This event provides that a credit support default can lead to an event of default under a derivative; it is obviously not applicable if the derivative counterparty's obligation under the transaction is not supported by another entity (guarantor).
- Misrepresentation. This event applies to misrepresentation made by either party or its credit support provider.
- Default under specified transaction. This event provides that if a party defaults on its obligation under a specified transaction, the nondefaulting party may elect to terminate the derivative.
- Cross default (obligation default). If specified in the schedule, this provision enables a party to declare the derivative to be in default if the other party or its credit support provider defaults on its obligation beyond an agreed threshold amount.
- Bankruptcy. This section applies to each party and any credit support provider, and specifies that if one party is subject to events associated with bankruptcy or insolvency, the other party may declare the derivative in default.
- Merger without assumption. If a party consolidates/amalgamates or merges with, or into, another entity, and the resulting entity fails to assume the obligations under the agreement, or no longer benefits from a credit support document (guarantee), then the other party can terminate any transaction under the agreement.

Section 5(b): Termination events

Section 5(b) indicates that the following events constitute termination events under the Master Agreement:

- illegality;
- force majeure event;
- tax event;
- tax event upon merger;
- additional termination events.

Each one of these events permits one of the parties to the transaction(s) to end the agreement and settle outstanding amounts due/payable.

Section 6: Early termination; close-out netting

Section 6 represents the right to end a transaction following an event of default or a termination event. If either occurs, the nondefaulting/affected party has the right to designate an early termination date (within 20 days of notice) on which the derivative(s) will end. Additionally, the parties have an opportunity to elect (in the Schedule) to have the derivative(s) terminate automatically and immediately upon the occurrence of certain events. This automatic termination election is designed to allow the nondefaulting party to exercise its termination rights outside of insolvency proceedings. In other words, it enables an early termination date to occur prior to the filing of insolvency proceedings. Note that it may be disadvantageous to select the automatic early termination election for derivatives that benefit from a credit support provider. The immediate termination of the derivative may not provide sufficient time to access and benefit from the credit support. The section also provides details on the effect of designation, determined either by referring to a market quotation source or by polling a group of dealers, and reflects changes in the credit profile of the reference obligor and reference asset. If it is not feasible to specify the method of valuation (such as default probabilities, or option pricing), an alternative might be to require the Calculation Agent to use a method of valuing the class of the referenced credit along with a detailed report outlining valuation methodology – formulas, assumptions, and models. If periodic valuations are required in the agreement, valuation methods should be consistent over the duration of the contract.

A Market Quotation is defined as a quotation from a leading dealer in the relevant market (selected by the nondefaulting party) for an amount that would be paid to the nondefaulting party or by the nondefaulting party in consideration of an agreement between the leading dealer and the nondefaulting party to enter into a transaction that would have the effect of preserving the economic equivalent of all payment or delivery obligations under the transactions. If more than three Market Quotations are provided, the Market Quotation is defined as the average of the remaining quotations after the highest and lowest are disregarded.

Section 6(e)(iv) of the ISDA Master Agreement supports the enforceability of the Market Quotation close-out valuation method. It states that the parties agree that, if Market Quotation applies, the amount recoverable under the ISDA Master Agreement on termination is a reasonable pre-estimate of loss, and not a penalty. Further, such an amount is payable for the loss of bargain and the loss of protection against future risks. Other than as specifically provided, neither party will be entitled to recover any additional damages as a consequence of such losses.

Loss is calculated by reference to the nondefaulting party's total losses and costs in connection with the ISDA Master Agreement as at the early termination date. It includes loss of

bargain, cost of funding, or, without duplication, loss or cost incurred as a result of its terminating, liquidating, obtaining or re-calculating on payment date, payments² on early termination (depending on whether early termination is via default or a termination event), and any set-off abilities. Close-out netting must be regarded as one of the chief benefits of the ISDA process, as it allows an entire portfolio of transactions between two parties to be condensed into a single net sum payable or receivable in the event of default. This eliminates the likelihood, at least in legal jurisdictions where netting is recognized, that a bankruptcy receiver will “cherry pick” only those contracts that have value to the defaulting party, while simultaneously dismissing those that are detrimental to the defaulting party.

Section 7: Transfer

Section 7 prevents both parties from transferring their obligations under the Agreement to another party without the prior written consent of the other party.

Section 8: Contractual currency

Section 8 indicates that payments under the Agreement will be made in the relevant currency specified in the Agreement for that payment (e.g. the “Contractual Currency”); in addition, it states that any shortfall or excess must be rectified immediately.

Other sections

Sections 9 to 14 provide for rights and obligations with regard to multi-branch dealings, expenses, and notices, and delineate governing law/jurisdiction and applicable definitions.

The ISDA framework is an essential legal control mechanism for institutions choosing to deal in synthetic instruments on an outright basis (e.g. as noted in Chapter 10), rather than those purchasing or selling derivative-type risks via structured notes. It is worth noting that in some instances, structured assets (e.g. a CLN) may cross reference the ISDA framework.

11.4 REGULATORY AND ACCOUNTING ISSUES

In this section we consider essential points related to regulatory capital, regulatory review, and formalized accounting treatment of derivatives and structured products.

11.4.1 Regulatory capital

We have noted at various points in the text that activity in some products is driven by attempts to reduce regulatory capital allocations (e.g. capital relief). Regulatory capital is the amount of qualifying capital that certain regulated institutions must hold in support of risky activities; the capital buffer is intended to protect against unexpected losses (much as reserves are intended to cover expected losses). Such capital typically is comprised of equity (retained earnings and paid-in capital), preferred stock, and certain types of qualifying long-term debt. Note that internal capital (or management capital), driven by an institution’s own assessment of its

² For example, an ISDA Master Agreement may prescribe the method that the Calculation Agent should use to value the asset when dealer quotations are not available or are not used. The dealer price is generally establishing any hedge or related trading position. It does not include legal fees. If Market Quotation has been specified, but cannot be determined or would not produce a commercially reasonable result, then loss is deemed to apply with respect to that particular transaction.

capital needs, operates on a “parallel” track: internal capital is intended to provide a similar level of support for unexpected losses, but the approach used to quantify the requirement is often different than the regulatory approach (primarily as a result of differences in modeling techniques, use of multiple legal entities and consolidation, and so forth).

An institution generally seeks capital relief through the use of a structured or synthetic mechanism when it transfers risk. In the most straightforward instances, favorable capital treatment can be obtained when the true sale mechanism transfers an originator’s assets (e.g. mortgages, risky bonds or loans), allowing a reduction in physical assets and the risk associated with such assets; this means capital need not be allocated against the transferred instruments. In synthetic transfers (e.g. those using CDSs or TRSs), the physical assets remain on the originator’s balance sheet but the risks are transferred out; this again results in capital reduction, as long as the transfer mechanism shifts the correct portfolio of exposures. It does not obviate the need to continue funding the physical balance sheet assets, but reduced regulatory capital is clearly in order as risk is reduced.

Apart from true sale or synthetic transfer of risky assets into a conduit, SPE, or trust, some institutions use the instruments we have discussed to hedge existing exposures. In such instances, regulatory capital reduction may again be in order – though this depends ultimately on the structure of the product and the nature/efficacy of the hedge. In order to consider how this process operates in practice, let us consider credit derivative instruments that are designed to transfer credit risk of a reference between two parties. In a TRS (or CDS), an institution that obtains downside protection against a reference credit has exposure to the underlying reference credit and the counterparty providing the protection. Counting both exposures is, of course, unnecessarily conservative: if the reference asset defaults, the worst that can happen is that the institution buying the protection will lose some amount associated with the default (if its counterparty fails to perform on the TRS). In practice, the relevant risk link is to the TRS counterparty that is providing the protection – indeed, that is the reason the institution has arranged the transaction. Accordingly, capital must be applied against the counterparty exposure, but not against the reference asset exposure. This is only true, however, if the reference asset in the TRS and the institution’s own trading book are the same, i.e. the TRS and reference are matched. In such cases, the institution can substitute the risk of the reference asset with the risk of the counterparty (which is almost certain to be of a higher quality credit). If the transaction is an offset but not a perfect match (e.g. the maturities are slightly different, or the seniority in the capital structure is different), then some capital relief is likely to be in order – but not the full amount. For instance, if the TRS has a shorter maturity than the reference asset, a full offset cannot be permitted.³ The reverse position can also be considered: the institution selling the protection needs to hold capital against the reference asset, but not against the counterparty.⁴ In instances where a high degree of correlation exists between the TRS seller and the reference asset, no offset may be permitted (e.g. a Turkish bank writing default protection on another Turkish corporate reference). The reasons for this are clear; default by the counterparty may occur at the same time as default of the reference asset if both parties are subject to the same local dislocation.

We can extend the discussion to structured notes. For instance, the issuer of a CLN (e.g. the buyer of credit protection) receives *de facto* cash collateralization from investors purchasing the

³ If a transaction is intended as a hedge but is mismatched with regard to maturity, the hedge exposure will become unhedged after the transaction has matured, suggesting that capital allocations will need to be increased once again.

⁴ It should be noted that amendments to certain regulatory frameworks, e.g. the Basle II approach, have sought to eliminate obvious discrepancies between regulatory capital treatment of contracts creating the same types of risk exposure, e.g. extending a direct loan with an 8 % capital charge versus creating a synthetic credit exposure to the same borrower via a CDS with a 1.6 % capital charge. Such large discrepancies generally have been eliminated, though some differences still exist.

note, and thus eliminates any capital requirement associated with either the reference credit or the counterparty risk elements. However, basis and maturity mismatches may still exist, suggesting that some amount of capital allocation may be required. Sellers of protection (e.g. the note investors) have exposure to the issuer of the CLN and the underlying reference asset; however, using the logic noted above, it need only allocate capital against the greater of the two exposures. Similar processes can be applied to other types of structured assets, e.g. CDOs, CMOs, credit card ABS, equity-linked notes, and so forth.

11.4.2 Regulatory review

Regulatory review of institutional dealings in synthetic and structured assets (including origination, distribution, trading) varies by local system, but typically involves one or more of the following: the national securities and/or banking regulator, the national derivatives regulator, and/or the local exchanges.

For instance, in the US, the SEC regulates synthetic and structured assets with a noncommodity linkage as standard securities transactions with embedded financial options. The SEC requires all public issues to be registered (individually or via shelf Rule 415), meaning that appropriate disclosure must be prepared and lodged by the issuer. Thus, a public CMO, CDO, CLN, or convertible bond must adhere to the SEC registration requirements. Transactions being structured and distributed via private placement are exempt from registration, but issuers must still prepare an information memorandum for distribution to qualified buyers; in addition, secondary trading restrictions (e.g. those limiting participation to a set number of qualified institutional buyers) under Rule 144A must be respected. Synthetic and structured assets with a commodity linkage are not classified as “prohibited” transactions, and are thus excluded from the statutory interpretation of the Commodity Futures Trading Commission (CFTC) regulations (i.e. CFTC Hybrid Statutory Interpretations). An excluded instrument is thus one where the present value of the index-independent payments is greater than the index-dependent price risk, the issuer receives full purchase price and the investor does not have to make any additional payments, the instrument is not marketed as a futures or options contract and is nondeliverable, and sales are restated to qualified persons/institutions.

In the UK, the Financial Services Authority (FSA) has jurisdiction over synthetic/structured products, and sets forth rules on how institutions must view, report, and treat any resulting activities. In Japan, the same scrutiny falls to the Ministry of Finance and Economy and the Japan Securities Dealers’ Association. Similar arrangements fall to regulators in various other countries.

11.4.3 Accounting treatment under FAS 133 and IAS 39

The accounting treatment of derivatives and structured/synthetic products with embedded derivatives has changed with the adoption of the Financial Accounting Standards 133 (FAS 133,⁵ used in the US) and International Accounting Standards 39 (IAS 39, used in

⁵ Prior to the introduction of FAS 133 in June 1999, US institutions dealing in derivatives and derivative-related instruments obtained guidance from FAS 115 (Account for Investment in certain Debt and Equity Securities), FAS 52 (Foreign Currency Translation), FAS 80 (Accounting for Futures Contracts), and general advice from the Emerging Issues Task Force. FAS 115 remains in force, and applies to debt and equity securities that generate income over the life of the asset, and that also feature one of the following characteristics: principal is at risk, the asset’s return is subject to variability, or the maturity of the asset is based on index/market circumstances beyond the control of the issuer and investor. Under FAS 115, assets are grouped into three classes: (1) hold to maturity assets, including structured assets (except those that are undated or those with conversion features); these securities are covered via the “retrospective interest method,” where income recognized for a particular reporting period is the difference between the amortized cost of the security at the end of the period and the amortized cost at the beginning of the period, plus any cash received during the period. The amortized cost at the end of the period is simply the PV of estimated future cash flows based on an interest rate assumption that equates all past actual and future cash flows to the initial investment. The effective yield is based on forward market rates or current spot rates as

regions/countries outside the US). FAS 133 and IAS 39 are accounting rules intended to clarify the treatment of financial derivatives and structured securities, by linking them directly to the corporate financial statements. Though FAS 133 and IAS 39 contain differences, efforts at harmonization are underway. Our focus in this section is on the more comprehensive FAS 133 framework, though comments remain quite applicable for institutions adhering to IAS 39.

Accounting requirements set forth by the Financial Accounting Standards Board (FASB) and the International Accounting Standards Board (IASB) require derivatives/structured asset users to account for their derivatives via the income statement/balance sheet, rather than simply in the financial footnotes. This adds a degree of transparency to an institution's use of derivatives, but the framework is not completely precise, and is still subject to some degree of interpretation. It has also proven expensive for firms to implement.⁶

FAS 133 requires an entity to reflect at fair value all derivatives and derivative-related instruments through the income statement/balance sheet. If certain conditions are met regarding the nature of hedge risks⁷ and the quality of the hedge,⁸ then a derivative can qualify for hedge accounting:

- **Fair value hedge.** A hedge against risk to changes in the fair value of an asset, liability, or commitment. If a derivative qualifies as a hedge of the exposure to changes in fair value of a recognized asset or liability, or an unrecognized commitment, then the gain or loss is recognized in earnings in the period of change, together with offsetting loss or gain on the hedge item. For example, the purchase of a put as a hedge for an underlying equity position held available for sale may be considered a fair value hedge.
- **Cash flow hedge.** A hedge against risk to variability of future cash flows of an asset, liability, or commitment. If the derivative qualifies as a hedge of the exposure to variable cash flows of a forecast transaction, the effective portion⁹ of gain/loss is reported in other comprehensive income (outside earnings), and reclassified as earnings when the forecast transaction affects earnings; the ineffective portion of gain/loss is reported in earnings immediately. For instance, the purchase of an interest rate forward as a hedge for a planned purchase of an underlying bond to be held available for sale may be considered a cash flow hedge.
- **Currency hedge.** A hedge against risk of change in the fair value of an asset, liability, or commitment in a foreign currency. If a derivative qualifies as a hedge of a foreign currency exposure of a net investment in a foreign operation or security held available for sale, the effective portion of gain/loss is reported in other comprehensive income (outside earnings), and reclassified as earnings when the foreign currency transaction affects earnings; the ineffective portion of gain/loss is reported in earnings immediately.

of the reporting period; (2) trading securities, including assets designed for short-term resale, which are carried at fair value at each reporting date, with changes reflected in the income statement; (3) assets not classed as (1) or (2), generally considered to be trading securities available for resale, which are carried at fair value at each reporting date, with unrealized changes reflected directly as other comprehensive income in the equity account. Any permanent impairment in the value of a security in (1)–(3) must be reflected in the current period income statement.

⁶ Some surveys have even suggested that the additional “onerous” reporting requirements, particularly for hedge accounting treatment, may discourage small/mid-sized companies from pursuing hedging activities. If true, this may dampen business – particularly for exotic derivatives, where it may be more difficult to justify hedge treatment; however, this may only be a short-term reaction to the framework.

⁷ Hedge risks are taken to include price risk, index risk, currency risk, interest rate risk, and default risk.

⁸ The actual conditions of hedge treatment can be complex to establish, and generally must be documented for ex-post audit or regulatory review.

⁹ The effective portion of a hedge is a function of cumulative change in fair value of the derivative to the cumulative change in the present value of expected cash flows of the forecast transaction.

In order to qualify for hedge accounting treatment, the strategy must be documented prior to the booking of a transaction (with objectives and strategy well delineated); in addition, the ongoing effectiveness of the hedge must be evaluated (i.e. at least every quarter). Short option positions (outright or embedded) generally do not qualify for hedge accounting treatment (apart from premium received), as the option seller assumes more risk than the amount given up.

If a derivative is not designated as a hedge the contract is carried at fair value, with any gain/loss recognized directly or indirectly in earnings in the period of change. This definition applies to most situations where an institution is dealing in synthetic and structured assets from an investment, speculation, or arbitrage perspective.

Conventional derivatives, structured assets, and synthetic derivatives we have discussed throughout the text generally must adhere to the FAS 133 guidelines above. This represents a change for structured products. Prior to FAS 133, instruments such as convertibles and structured notes were treated from an accounting perspective according to the predominant component of the package; the process began changing in the mid-1990s and was ultimately formalized via FAS 133. For instance, prior to FAS 133, convertibles and other debt with nondetachable warrants with an equity component featuring little value at issuance typically were treated as debt. Securities with a detachable option and/or considerable value were split between debt and equity, with the debt component issued at a discount and amortized as additional interest expense, and the equity classified as an addition to permanent equity. FAS 133 formalized this requirement, which is now applicable regardless of the value of the equity option.¹⁰

Under FAS 133, a derivative is defined in terms of its notional and underlying reference, and includes contracts that require no/small initial investment and which permit net settlement. This definition is applied to standalone derivatives and the embedded derivative component of a structured or synthetic asset. Thus, an institution dealing with a callable asset swap, convertible bond, CLN, or credit-based TRS must apply the rules summarized above in order to arrive at the correct accounting treatment. So, if any structured instrument is deemed to be a current, prospective, or currency hedge, it will receive hedge treatment; if any one is considered a speculative position, the institution will be required to post any gain/loss to earnings in the period of change.¹¹

Ultimately, financial and nonfinancial institutions that adhere to best practices established by regulators or their own boards of directors, and related to the entire range of internal controls, will be well positioned to deal in synthetic and structured products. Such activities should, in turn, lead to additional growth and innovation in what is already a very vibrant marketplace.

¹⁰ It is important to note that equity hybrids with physical conversion/redemption that is outside the explicit control of the issuer can be classified apart from equity – as a true hybrid between debt and equity. In effect, the option or warrant component is classed as “temporary equity” until it expires or is exercised; if exercise occurs, the shares become part of the issuer’s equity account. Equity instruments that may be settled in cash or other assets (but not an issuer’s new shares) cannot be treated in this fashion, and are thus subject to the FAS 133 accounting rules.

¹¹ The exception would arise if the entire package were already being carried at fair value for trading securities, in which case, gains/losses in the combined package would be recognized in earnings during the period of change. Thus, the embedded derivative should be recorded at fair value, and the host should be considered an investment in a discount bond with a yield that is distinct from that stated on the instrument (meaning that the host is likely to have a value that is different from the initial fair value). For instance, a \$1 m convertible bond with a coupon of 6 % and an embedded call option that is worth a theoretical \$100 000 suggests a \$900 000 host bond with a yield of 7.4 %, (recorded through coupon and discount amortization). Note that the yield is a function of the initial book value, and not a function of rates and credit spreads.

Selected References

- Acharya, V., Das, S. and Sundrama, R. (2002) "Pricing Credit Derivatives With Rating Transitions," *Financial Analysts Journal*, **58**, 28.
- Aldred, C. (2003) "Regulators Eyeing Credit Hedge Coverage," *Business Insurance*, **37**, 25.
- Banks, E. (2002) *The Credit Risk of Complex Derivatives*, 3rd edition, London: Palgrave.
- Banks, E. (2003) *Alternative Risk Transfer*, Chichester: John Wiley & Sons, Ltd.
- Banks, E. (2004) *Catastrophic Risk*, Chichester: John Wiley & Sons, Ltd.
- Banks, E. and Dunn, R. (2002) *Practical Risk Management*, Chichester: John Wiley & Sons, Ltd.
- Black, F. and Scholes, M. (1973) "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, **81**, May–June.
- Brennan, M. and Schwartz, E. (1980) "Analyzing Convertible Bonds," *Journal of Financial and Quantitative Analysis*, **15**, 907–929.
- Chen, R. and Sopranozzetti, B. (2003) "The Valuation of Default-Triggered Credit Derivatives," *Journal of Financial and Quantitative Analysis*, **38**(2), 359.
- Chew, D. (1992) "The Use of Hybrid Debt in Managing Corporate Risk," *Journal of Applied Corporate Finance*, **2**, Winter.
- Chorafas, D. (1999) *Credit Derivatives and The Management of Risk*, Chichester: John Wiley & Sons, Ltd.
- Choudhry, M. (2004) *Structured Credit Products*, Chichester: John Wiley & Sons, Ltd.
- Das, S. (1994) *Swaps and Financial Derivatives*, 2nd edition, Sydney: Law Book Company.
- Das, S. (1996) *Structured Notes and Derivative-Embedded Securities*, London: Euromoney.
- Das, S. (2002) *Credit Derivatives and Credit Linked Notes*, 2nd edition, Singapore: John Wiley & Sons.
- Fabozzi, F. (Ed.) (1993) *Bond Markets*, 2nd edition, Englewood Cliffs, NJ: Prentice Hall.
- Fabozzi, F. (Ed.) (1993) *The Handbook of Mortgage Backed Securities*, 2nd edition, Chicago: Probus.
- Francis, J., Joyce, A. and Whittaker, J. G. (1999) *The Handbook of Equity Derivatives*, New York: McGraw Hill.
- Goodman, L. and Fabozzi, F. (2002) *Collateralized Debt Obligations*, New York: John Wiley & Sons, Inc.
- McDonald, R. (2003) *Derivatives Markets*, Boston: Addison Wesley.
- Merton, R. (1974) "On the Pricing of Corporate Debt," *Journal of Finance*, **29**, 449–470.
- Nelken, I. (1999) *Implementing Credit Derivatives*, New Jersey: Irwin.
- Ong, M. (1999) *Internal Credit Risk Models*, London: Risk Books.
- Redmayne, J. (1993) *Convertibles*, London: Euromoney.
- Rizzi, J. (2003) "Risk Implications of Credit Derivative Instruments," *Commercial Lending Review*, **18**, 5.
- Smithson, C. (1999) *Managing Financial Risk*, 3rd edition, New York: McGraw Hill.
- Specht, B. (2001) "Synthetic Securitization Enters Next Generation," *Euromoney*, February, p. 28.
- Tavakoli, J. (1998) *Credit Derivatives*, New York: John Wiley & Sons, Inc.
- Zubulake, L. (1991) *Guide to Convertible Securities Worldwide*, New York: John Wiley & Sons, Inc.

Index

- ABS, see Asset-backed securities
- Accrual bonds, 84
- Adjustable convertible notes, 178
- Adjustable rate preferred stock, 177–178
- ARPS, see Adjustable rate preferred stock
- Asset-backed securities (ABS), 90–98
- Asset swaps, 231–234
- Asset swaptions, 238
- Asset switches, 237
- Backwardation, 23
- Balance sheet and arbitrage CDOs, 134–135
- Basis swaps, 28
- Bonds with debt warrants, 107
- Bonds with [equity] warrants, 172–173
- Break forwards, 227
- Butterflies, 217–218
- Buy/write securities, 173–177
- Calendar spreads, 219–220
- Call option, 29–30
- Callable, puttable, and stripped securities, 45–58
 - Development and market drivers, 45–47
 - Product mechanics and applications, 47–58
 - Callable bonds, 47–52
 - Puttable bonds, 52–53
 - Stripped securities, 53–58
- Callable asset swaps, 235–236
- Callable bonds, 47–52
- Callable credit default notes, 116
- Callable, puttable, and extendable swaps, 224–226
- Caps, 32–37
- Catastrophe bonds, 147–151
- CBO, see Collateralized bond obligations
- CDO, see Collateralized debt obligations
- CDOs-squared, 140
- CDS index tranches, 139–140
- CLN, see Credit-linked notes
- CLO, see Collateralized loan obligations
- Closed-end funds, 193–195
- CMBS, see Commercial mortgage backed securities
- CMO, see Collateralized mortgage obligations
- Collared FRNs, 105
- Collateralized debt obligations (CDOs), 121–140
 - Development and market drivers, 121–123
 - Product mechanics and applications, 123–140
 - Balance sheet and arbitrage CDOs, 134–135
 - Cash flow and market value CDOs, 133–140
 - Collateralized bond obligations (CBOs), 131–133
 - Collateralized loan obligations (CLOs), 130–131
 - Synthetic and structured CDOs, 135–140
- Collars, 226
- Collateralized bond obligations (CBOs), 131–133
- Collateralized loan obligations (CLOs), 130–131
- Collateralized mortgage obligations (CMOs), 60, 83–90
- Commercial mortgage backed securities (CMBS), 60, 79–83
- Committed capital facilities, 156–157
- Commodity-linked notes and loans, 109–110
- Condors, 218–219
- Constant prepayment rate (CPR), 65
- Contango, 22
- Contingency loans, 158–159
- Contingent capital structures, 153–162
- Contingent debt, 156–159
- Contingent equity, 159–162
- Contingent premium options, 230
- Contingent surplus notes, 157–158

- Convertible bonds and equity hybrids, 163–179
 - Development and market drivers, 163–165
 - Product mechanics and applications, 165–179
 - Bonds with [equity] warrants, 172–173
 - Buy/write securities, 173–177
 - Convertible bonds and variations, 166–172
 - Other equity hybrids, 177–179
- Convertible bonds, 166–172
- Convertible preferred stock, 178–179
- Convexity, 14–16
- CPR, see Constant prepayment rate
- Credit arbitrage, 25–27
- Credit default swaps, 113–114, 238
- Credit derivatives and synthetic credit positions, 238–245
- Credit-linked notes (CLNs), 113–119, 135–136
- Credit risk, 10
- Credit spread forwards, 114, 240–241
- Credit spread notes, 115
- Credit spread options, 114, 239–240
- Currency-linked notes, 107–109
- Currency swaps, 287
- Default notes, 115
- Derivative replication, repackaging, and structuring, 205–245
 - Development and market drivers, 205–207
 - Product mechanics and applications, 207–245
 - Asset swaps, liability swaps, asset swap packages, 231–238
 - Callable, puttable, and extendable swaps, 224–226
 - Credit derivatives and synthetic credit positions, 238–245
 - Multi-option strategies, 215–219
 - Multiple swap/option positions, 214–215
 - Other structured derivatives, 226–231
 - Swaptions, 220–224
 - Synthetic long and short options and swaps, 207–214
- Differential swaps, 229
- Dual coupon currency swaps, 230
- Dual currency bonds, 108–109
- Duration, 13
- Earthquake bonds, 148–149
- Equity-linked notes, 110–113
- ETFs, see Exchange-traded funds
- Exchange-traded funds (ETFs), 200–203
- Exchangeable bonds, 172
- FAS 133, See Financial Accounting Standards 133
- Financial Accounting Standards 133 (FAS 133), 257–259
- Financial building blocks, 9–43
 - Concepts, 9–16
 - Risks, 9–11
 - Time value of money and interest rates, 11–16
 - Derivative instruments, 16–37
 - Exchange-traded derivatives, 36–38
 - Futures, 37
 - Futures options, 37–38
 - Options, 37–38
 - OTC derivatives, 16–36
 - Options, 29–36
 - Forwards, 17–24
 - Swaps, 24–29
 - Financial engineering, 43
 - Host securities/liabilities, 38–39
 - Private placements, 38–39
 - Public notes, 38
 - Issuing/repackaging vehicles, 39–43
 - Investment companies, 42–43
 - Special purpose entities (SPEs), 40–41
 - Trusts, 41–42
- Federal Home Loan Mortgage Corporation (FHLMC), 59, 71–72
- Federal National Mortgage Association (FNMA), 59, 71–72, 88
- FHLMC, see Federal Home Loan Mortgage Corporation
- Floors, 32–37
- FNMA, see Federal National Mortgage Association
- Forward interest rates, 14–16
- Forward rate agreement (FRA), 19–21
- Forward swaps, 228–229
- Forwards, 17–24
- FRA, see Forward rate agreement
- Fund of funds, 197
- Futures, 37
- Futures options, 37–38
- GNMA, see Government National Mortgage Association
- Government National Mortgage Association (GNMA), 59, 71–72
- Guaranteed investment contract, 53
- Hedge funds, 195–200
- Hurricane bonds, 148
- IAR [swap], see Index amortizing rate swap
- IAS 39, See International Accounting Standards 39
- ILS, see Insurance-linked securities
- Index amortizing rate (IAR) swap, 225
- Insurance-linked securities (ILS), 144–153

- Insurance-linked securities and contingent capital, 141–162
 - Development and market drivers, 141–143
 - Product mechanics and applications, 143–162
 - Contingent capital structures, 153–162
 - Contingent debt, 156–159
 - Contingent equity, 159–162
 - Insurance-linked securities (ILS), 144–153
 - Catastrophe bonds, 147–151
 - Noncatastrophe bonds, 151–153
- Interest coverage, 129
- Interest only (IO) [strip], 53–58, 60, 88–89
- Interest rate-linked notes, 104–107
- Interest rate swaps, 24–27
- International Accounting Standards 39 (IAS 39), 257–259
- International Swaps and Derivatives Association (ISDA),
 - Framework, 250–255
 - Master Agreement, 251
- Inverse floaters, 104
- Investment companies, 42–43
- Investment funds, 181–203
 - Development and market drivers, 181–184
 - Product mechanics and applications, 184–203
 - Closed-end funds, 193–195
 - Exchange-traded funds (ETFs), 200–203
 - Hedge funds, 195–200
 - Open-end funds, 186–193
- IO [strip], see Interest only [strip]
- ISDA, see International Swaps and Derivatives Association
- Knock-out floater, 106
- Knock-out forwards, 227–228
- Leveraged inverse floaters, 104–105
- Liability swaps, 234–235
- Life insurance securitization, 152–153
- Liquidity risk, 10
- Loss equity puts, 159–161
- Mandatory convertible bonds, 172
- Market risk, 10
- Maturity shortening, 118
- MBS, see Mortgage-backed securities
- Mortgage- and asset backed securities, 59–98
 - Development and market drivers, 59–62
 - Product mechanics and applications, 62–98
 - Asset-backed securities, 90–98
 - Credit card securitizations, 91–94
 - Home equity loan securitizations, 94–95
 - Other asset-backed securitizations, 95–98
 - Mortgage-backed securities, 62–90
 - Collateralized mortgage obligations, 83–90
 - Pass-through securities, 69–83
- Mortgage-backed securities, 62–90
- Multiple peril bonds, 149–150
- Noncatastrophe bonds, 151–153
- OAS, see Option adjusted spread
- Open-end funds, 186–193
- Operating risk, 10
- Option adjusted convexity, 50–51
- Option adjusted duration, 50
- Option adjusted spread (OAS), 50, 67–68
- Options, 29–36
- Overcollateralization, 128–129
- PAC, see Planned amortization class [bond]
- Pass-through securities, 69–83
- PERCS and synthetic PERCS, 176–177
- Pfandbrief, 76–77
- Planned amortization class (PAC) [bond], 86–87
- PO [strip], see Principal only [strip]
- Prepayments, 62–67
- Price spreads, 216
- PRIMES and SCORES, 175–176
- Principal only (PO) [strip], 53–58, 60, 88–89
- Put-call parity, 36
- Put option, 30
- Put-protected equity, 161–162
- Puttable asset swaps, 236–237
- Puttable bonds, 52–53
- Range floaters, 106
- Ratchet collared FRNs, 105
- Real estate investment trust (REIT), 70
- REIT, see Real estate investment trust
- Repackaged bonds, 117
- Residual value bonds, 151–152
- Reverse convertible bonds, 172
- Risk, legal, and regulatory issues, 247–259
 - Legal controls, 249–255
 - ISDA documentation framework, 250–255
 - Regulatory and accounting issues, 255–259
 - Accounting treatment, 257–259
 - Regulatory capital, 255–257
 - Regulatory review, 257
- Risk and financial controls, 247–249
 - Internal audit, 249
 - Internal financial and operational controls, 248–249
- Market, liquidity and credit risk management, 247–248

-
- Semi-fixed swaps, 228
 - Separate Trading of Registered Interest and Principal Securities (STRIPS), 46
 - Single tranche CDOs, 138–139
 - SLMA, see Student Loan Marketing Association
 - Soft bullet, 93, 96
 - Special purpose entity (SPE), 40–41
 - Special purpose reinsurer (SPR), 143
 - SPE, see Special purpose entity
 - SPR, see Special purpose reinsurer
 - Step-up notes, 107
 - Straddles, 216
 - Strangles, 216–217
 - Stripped securities, 53–58
 - STRIPS, see Separate Trading of Registered Interest and Principal Securities
 - Structured notes and loans, 99–119
 - Development and market drivers, 99–102
 - Product mechanics and applications, 102–119
 - Commodity-linked notes and loans, 109–110
 - Credit-linked notes (CLNs), 113–119
 - Currency-linked notes, 107–109
 - Equity-linked notes, 110–113
 - Interest rate-linked notes, 104–107
 - Student Loan Marketing Association (SLMA), 97
 - Superfloater swaps, 230
 - Swaps, 24–29
 - Swaptions, 220–224
 - Synthetic and structured assets
 - Development of, 2–3
 - Drivers of activity, 3–5
 - New product design, 6–7
 - Synthetic and structured CDOs, 135–140
 - Synthetic asset positions, 212
 - Synthetic bonds, 119
 - Synthetic equity warrants, 113
 - Synthetic interest rate and currency swaps, 215
 - Synthetic option positions, 212
 - TAC, see Targeted amortization class [bond]
 - Targeted amortization class (TAC)[bond], 87–88
 - Time value of money, 11–16
 - Total return swap (TRS), 114
 - Total return swap credit-linked notes, 115
 - Trade credit securitizations, 152
 - TRS, see Total return swap
 - Trusts, 41–42
 - UIT, see Unit investment trust
 - Unit investment trust (UIT), 194, 200
 - VADM, see Very Accurately Defined Maturity [bond]
 - Very Accurately Defined Maturity (VADM) [bond], 88
 - Waterfall, 127
 - Weather bonds, 151
 - Windstorm bonds, 149
 - Zero coupon convertible bonds, 171