



THE LEHMAN BROTHERS

GUIDE TO EXOTIC CREDIT DERIVATIVES


THE LEHMAN BROTHERS GUIDE TO EXOTIC CREDIT DERIVATIVES

LEHMAN BROTHERS

risk waters group

Effective Structured Credit Solutions for our Clients

Structured Credit Solutions



With over seventy professionals worldwide, Lehman Brothers gives you access to top quality risk-management, structuring, research and legal expertise in structured credit. The team combines local market knowledge with global co-ordinated expertise.


Lehman Brothers has designed specific solutions to our clients' problems, including yield-enhancement, capital relief, portfolio optimisation, complex hedging and asset-liability management.

Product Innovation



- Credit Default Swaps
- Portfolio Swaps
- Credit Index Products
- Repackagings
- Default Baskets
- Secondary CDO trading
- Customised CDO tranches
- Default swaptions
- Credit hybrids

Leadership in Fixed Income Research



For further information please contact your local sales representative or call:

- London:** **Giancarlo Saronne**
+44 20 7260 2745 gsaronne@lehman.com
- New York:** **Mike Glover**
+1 212 526 7090 mgllover@lehman.com
- Tokyo:** **Jawahar Chirimar**
81-3-5571-7257 jchirima@lehman.com

LEHMAN BROTHERS
Where vision gets built.®

Foreword

The credit derivatives market has revolutionised the transfer of credit risk. Its impact has been borne out by its significant growth which has currently achieved a market notional close to \$2 trillion. While not directly comparable, it is worth noting that the total notional outstanding of global investment grade corporate bond issuance currently stands at \$3.1 trillion.

This growth in the credit derivatives market has been driven by an increasing realisation of the advantages credit derivatives possess over the cash alternative, plus the many new possibilities they present to both credit investors and hedgers. Those investors seeking diversification, yield pickup or new ways to take an exposure to credit are increasingly turning towards the credit derivatives market.

The primary purpose of credit derivatives is to enable the efficient transfer and repackaging of credit risk. In their simplest form, credit derivatives provide a more efficient way to replicate in a derivative format the credit risks that would otherwise exist in a standard cash instrument.

More exotic credit derivatives such as syn-

thetic loss tranches and default baskets create new risk-return profiles to appeal to the differing risk appetites of investors based on the tranching of portfolio credit risk. In doing so they create an exposure to default correlation. CDS options allow investors to express a view on credit spread volatility, and hybrid products allow investors to mix credit risk views with interest rate and FX risk.

More recently, we have seen a stepped increase in the liquidity of these exotic credit derivative products. This includes the development of very liquid portfolio credit vehicles, the arrival of a two-way correlation market in customised CDO tranches, and the development of a more liquid default swaptions market. To enable this growth, the market has developed new approaches to the pricing and risk-management of these products.

As a result, this book is divided into two parts. In the first half, we describe how exotic structured credit products work, their rationale, risks and uses. In the second half, we review the models for pricing and risk managing these various credit derivatives, focusing on implementation and calibration issues.

Authors

Dominic O'Kane

T. +44 207 260 2628

E. dokane@lehman.com

Marco Naldi

T. +1 212 526 1728

E. mnaldi@lehman.com

Sunita Ganapati

T. +1 415 274 5485

E. sganapati@lehman.com

Arthur Berd

T. +1 212 526 2629

E. arthur.berd@lehman.com

Claus Pedersen

T. +1 212 526 7775

E. cmpeders@lehman.com

Lutz Schloegl

T. +44 207 260 2113

E. luschloe@lehman.com

Roy Mashal

T. +1 212 526 7931

E. rmashal@lehman.com

Contents

Foreword	1
Credit Derivatives Products	
Market overview	3
The credit default swap	4
Basket default swaps	8
Synthetic CDOs	12
Credit options	23
Hybrid products	28
Credit Derivatives Modelling	
Single credit modelling	31
Modelling default correlation	33
Valuation of correlation products	39
Estimating the dependency structure	43
Modelling credit options	47
Modelling hybrids	51
References	53

Credit Derivatives Products

Market overview


The credit derivatives market has changed substantially since its early days in the late 1990s, moving from a small and highly esoteric market to a more mainstream market with standardised products. Initially driven by the hedging needs of bank loan managers, it has since broadened its base of users to include insurance companies, hedge funds and asset managers.

The latest snapshot of the credit derivatives market was provided in the 2003 *Risk Magazine* credit derivatives survey. This survey polled 12 dealers at the end of 2002, composed of all the major players in the credit derivatives market. Although the reported numbers cannot be considered 'hard', they can be used to draw fairly firm conclusions about the recent direction of the market.

According to the survey, the total market outstanding notional across all credit derivatives products was calculated to be \$2,306 billion, up more than 50% on the previous year. **Single name CDS remain the most used instrument in the credit derivatives world with 73% of market outstanding notional, as shown in Figure 1. This supports our observation that the credit default market has become more mainstream, focusing on the liquid standard contracts. We believe that this growth in CDS has been driven by hedging demand generated by synthetic CDO positions, and by hedge funds using credit derivatives as a way to exploit capital structure arbitrage opportunities and to go outright short the credit markets.**

An interesting statistic from the survey is the relatively equal representation of North American and European credits. The survey

Figure 1. Market breakdown by instrument type




Source: *Risk Magazine 2003 Credit Derivatives Survey*

showed that 40.1% of all reference entities originate in Europe, compared with 43.8% from North America. This is in stark contrast to the global credit market which has a significantly smaller proportion of European originated bonds relative to North America.

The base of credit derivatives users has been broadening steadily over the last few years. We show a breakdown of the market by end-users in Figure 2 (overleaf). Banks still remain the largest users with nearly 50% share. This is mainly because of their substantial use of CDS as hedging tools for their loan books, and their active participation in synthetic securitisations. The hedging activity driven by the issuance of synthetic CDOs (discussed later) has for the first time satisfied the demand to buy protection coming from bank loan hedgers. Readers are referred to Ganapati et al (2003) for a full discussion of the market impact.

Insurance companies have also become an important player, mainly by investing in investment-grade CDO tranches. As a result,

Figure 2. Breakdown by end users



Source: Risk Magazine 2003 Credit Derivatives Survey.

the insurance share of credit derivatives usage has increased to 14% from 9% the previous year.

More recently, the growth in the usage of credit derivatives by hedge funds has had a marked impact on the overall credit derivatives market itself, where their share has increased to 13% over the year. Hedge funds have been regular users of CDS especially around the convertible arbitrage strategy. They have also been involved in many of the 'fallen angel' credits where they have been significant buyers of protection. Given their ability to leverage, they have substantially increased their volume of CDS contracts traded, which in many cases has been disproportionate to their absolute size.

Finally, in portfolio products, by which we mean synthetic CDOs and default baskets, the total notional for all types of credit derivatives portfolio products was \$449.4 billion. Their share has kept pace with the growth of the credit derivatives market at about 22% over the last two years. This is not a surprise, since there is a fundamentally symbiotic relationship between the synthetic CDO and single name CDS markets, caused by dealers originating synthet-

ic tranches either by issuing the full capital structure or hedging bespoke tranches.

Since this survey was published, the credit derivatives market has continued to consolidate and innovate. The ISDA 2003 Credit Derivative Definitions were another milestone on the road towards CDS standardisation. The year 2003 has also seen a significant increase in the usage of CDS portfolio products. There has been a stepped increase in liquidity for correlation products, with daily two-way markets for synthetic tranches now being quoted. The credit options market, in particular the market for those written on CDS, has grown substantially.

A number of issues still remain to be resolved. First, there is a need for the generation of a proper term structure for credit default swaps. The market needs to build greater liquidity at the long end and, especially, the short end of the credit curve. Greater transparency is also needed around the calibration of recovery rates. Finally, the issue of the treatment of restructuring events still needs to be resolved. Currently, the market is segregated along regional lines in tackling this issue, but it is hoped that a global standard will eventually emerge.

The credit default swap

The credit default swap is the basic building block for most 'exotic' credit derivatives and hence, for the sake of completeness, we set out a short description before we explore more exotic products.

A credit default swap (CDS) is used to transfer the credit risk of a reference entity (corporate or sovereign) from one party to another. In a standard CDS contract one party purchases credit protection from the other party, to cover the loss of the face value of an asset following a credit event. A credit event is a legally defined event that typically includes

bankruptcy, failure to pay and restructuring. Buying credit protection is economically equivalent to shorting the credit risk. Equally, selling credit protection is economically equivalent to going long the credit risk.

This protection lasts until some specified maturity date. For this protection, the protection buyer makes quarterly payments, to the protection seller, as shown in Figure 3, until a credit event or maturity, whichever occurs first. This is known as the premium leg. The actual payment amounts on the premium leg are determined by the CDS spread adjusted for the frequency using a basis convention, usually Actual 360.

If a credit event does occur before the maturity date of the contract, there is a payment by the protection seller to the protection buyer. We call this leg of the CDS the protection leg. This payment equals the difference between par and the price of the assets of the reference entity on the face value of the protection, and compensates the protection buyer for the loss. There are two ways to settle the payment of the protection leg, the choice being made at the initiation of the contract. They are:

Physical settlement – This is the most widely used settlement procedure. It requires the protection buyer to deliver the notional


amount of deliverable obligations of the reference entity to the protection seller in return for the notional amount paid in cash. In general there are several deliverable obligations from which the protection buyer can choose which satisfy a number of characteristics. Typically they include restrictions on the maturity and the requirement that they be *pari passu* – most CDS are linked to senior unsecured debt.

If the deliverable obligations trade with different prices following a credit event, which they are most likely to do if the credit event is a restructuring, the protection buyer can take advantage of this situation by buying and delivering the cheapest asset. The protection buyer is therefore long a cheapest to deliver option.

Cash settlement – This is the alternative to physical settlement, and is used less frequently in standard CDS but overwhelmingly in tranching CDOs, as discussed later. In cash settlement, a cash payment is made by the protection seller to the protection buyer equal to par minus the recovery rate of the reference asset. The recovery rate is calculated by referencing dealer quotes or observable market prices over some period after default has occurred.

Suppose a protection buyer purchases five-year protection on a company at a CDS spread of 300bp. The face value of the protection is \$10m. The protection buyer therefore makes quarterly payments approximately (we ignore calendars and day count conventions) equal to $\$10\text{m} \times 0.03 \times 0.25 = \$75,000$. After a short period the reference entity suffers a credit event. Assuming that the cheapest deliverable asset of the reference entity has a recovery price of \$45 per \$100 of face value, the payments are as follows:

Figure 3. Mechanics of a CDS



- ☐ The protection seller compensates the protection buyer for the loss on the face value of the asset received by the protection buyer and this is equal to \$5.5m.
- ☐ The protection buyer pays the accrued premium from the previous premium payment date to the time of the credit event. For example, if the credit event occurs after a month then the protection buyer pays approximately $\$10\text{m} \times 300\text{bp} \times 1/12 = \$25,000$ of premium accrued. Note that this is the standard for corporate reference entity linked CDS.

For severely distressed reference entities, the CDS contract trades in an up-front format where the protection buyer makes a cash payment at trade initiation which purchases protection to some specified maturity – there are no subsequent payments unless there is a credit event in which the protection leg is settled as in a standard CDS. For a full description of up-front CDS see O’Kane and Sen (2003).

Liquidity in the CDS market differs from the cash credit market in a number of ways. For a start, a wider range of credits trade in the CDS market than in cash. In terms of maturity, the most liquid CDS is the five-year contract, followed by the three-year, seven-year and 10-year. The fact that a physical asset does not need to be sourced means that it is generally easier to transact in large round sizes with CDS.

Uses of a CDS

The CDS can do almost everything that cash can do and more. We list some of the main applications of CDS below.

- ☐ The CDS has revolutionised the credit markets by making it easy to short credit.

This can be done for long periods without assuming any repo risk. This is very useful for those wishing to hedge current credit exposures or those wishing to take a bearish credit view.

- ☐ CDS are unfunded so leverage is possible. This is also an advantage for those who have high funding costs, because CDS implicitly lock in Libor funding to maturity.
- ☐ CDS are customisable, although deviation from the standard may incur a liquidity cost.
- ☐ CDS can be used to take a spread view on a credit, as with a bond.
- ☐ Dislocations between cash and CDS present new relative value opportunities. This is known as trading the default swap basis.

Evolution of CDS documentation

The CDS is a contract traded within the legal framework of the International Swaps and Derivatives Association (ISDA) master agreement. The definitions used by the market for credit events and other contractual details have been set out in the ISDA 1999 document and recently amended and enhanced by the ISDA 2003 document. The advantage of this standardisation of a unique set of definitions is that it reduces legal risk, speeds up the confirmation process and so enhances liquidity.

Despite this standardisation of definitions, the CDS market does not have a universal standard contract. Instead, there is a US, European and an Asian market standard, differentiated by the way they treat a restructuring credit event. This is the consequence of a desire to enhance the posi-

tion of protection sellers by limiting the value of the protection buyer's delivery option following a restructuring credit event. A full discussion and analysis of these different standards can be found in O'Kane, Pedersen and Turnbull (2003).

Determining the CDS spread

The premium payments in a CDS are defined in terms of a CDS spread, paid periodically on the protected notional until maturity or a credit event. It is possible to show that the CDS spread can, to a first approximation, be proxied by either (i) a par floater bond spread (the spread to Libor at which the reference entity can issue a floating rate note of the same maturity at a price of par) or (ii) the asset swap spread of a bond of the same maturity provided it trades close to par.

Demonstrating these relationships relies on several assumptions that break down in practice. For example, we assume a common market-wide funding level of Libor, we ignore accrued coupons on default, we ignore the delivery option in the CDS, and we ignore counterparty risk. Despite these assumptions, cash market spreads usually provide the starting point for where CDS spreads should trade. The difference between where CDS spreads and cash LIBOR spreads trade is known as the Default Swap Basis, defined as:

Basis = CDS Spread – Cash Libor Spread.

A full discussion of the drivers behind the CDS basis is provided in O'Kane and McAdie (2001). A large number of investors now exploit the basis as a relative value play.

Determining the CDS spread is not the same as valuing an existing CDS contract.

For that we need a model and a discussion of the valuation of CDS is provided on page 32.

Funded versus unfunded

Credit derivatives, including CDS, can be traded in a number of formats. The most commonly used is known as swap format, and this is the standard for CDS. This format is also termed 'unfunded' format because the investor makes no upfront payment. Subsequent payments are simply payments of spread and there is no principal payment at maturity. Losses require payments to be made by the protection seller to the protection buyer, and this has counterparty risk implications.

The other format is to trade the risk in the form of a credit linked note. This format is known as 'funded' because the investor has to fund an initial payment, typically par. This par is used by the protection buyer to purchase high quality collateral. In return the protection seller receives a coupon, which may be floating rate, ie, Libor plus a spread, or may be fixed at a rate above the same maturity swap rate. At maturity, if no default has occurred the collateral matures and the investor is returned par. Any default before maturity results in the collateral being sold, the protection buyer covering his loss and the investor receiving par minus the loss. The protection buyer is exposed to the default risk of the collateral rather than the counterparty.

Traded CDS portfolio products

CDS portfolio products are products that enable the investor to go long or short the credit risk associated with a portfolio of CDS in one transaction.

In recent months, we have seen the emergence of a number of very liquid portfolio products, whose aim is to offer investors a diverse, liquid vehicle for assuming or hedg-

ing exposure to different credit markets, one example being the TRAC-XSM vehicle. These have added liquidity to the CDS market and also created a standard which can be used to develop portfolio credit derivatives such as options on TRAC-X.

The move of the CDS market from banks towards traditional credit investors has greatly increased the need for a performance benchmark linked directly to the CDS market. As a consequence, Lehman Brothers has introduced a family of global investment grade CDS indices which are discussed in Munves (2003).

These consist of three sub-indices, a US 250 name index, a European 150 name index and a Japanese 40 name index. All names are corporates and the maturity of the index is maintained close to five years. Daily pricing of all 440 names is available on our LehmanLive website.

Basket default swaps

Correlation products are based on redistributing the credit risk of a portfolio of single-name credits across a number of different securities. The portfolio may be as small as five credits or as large as 200 or more credits. The redistribution mechanism is based on the idea of assigning losses on the credit portfolio to the different securities in a specified priority, with some securities taking the first losses and others taking later losses. This exposes the investor to the tendency of assets in the portfolio to default together, ie, default correlation. The simplest correlation product is the basket default swap.

A basket default swap is similar to a CDS, the difference being that the trigger is the nth credit event in a specified basket of reference entities. Typical baskets contain five to 10 reference entities. In the particular case

of a first-to-default (FTD) basket, $n=1$, and it is the first credit in a basket of reference credits whose default triggers a payment to the protection buyer. As with a CDS, the contingent payment typically involves physical delivery of the defaulted asset in return for a payment of the par amount in cash. In return for assuming the nth-to-default risk, the protection seller receives a spread paid on the notional of the position as a series of regular cash flows until maturity or the nth credit event, whichever is sooner.

The advantage of an FTD basket is that it enables an investor to earn a higher yield than any of the credits in the basket. This is because the seller of FTD protection is leveraging their credit risk.


To see this, consider that the fair-value spread paid by a credit risky asset is determined by the probability of a default, times the size of the loss given default. FTD baskets leverage the credit risk by increasing the probability of a loss by conditioning the payoff on the first default among several credits. The size of the potential loss does not increase relative to buying any of the assets in the basket. The most that the investor can lose is par minus the recovery value of the FTD asset on the face value of the basket.

The advantage is that the basket spread paid can be a multiple of the spread paid by the individual assets in the basket. This is shown in Figure 4 where we have a basket of five investment grade credits paying an average spread of about 28bp. The FTD basket pays a spread of 120bp.

More risk-averse investors can use default baskets to construct lower risk assets: second-to-default (STD) baskets, where $n=2$, trigger a credit event after two or more assets have defaulted. As such they are lower risk second-loss exposure products which will pay a lower spread than an FTD basket.

TRAC-X is a service mark of JPMorgan and Morgan Stanley

Figure 4. Five-year first to default (FTD) basket on five credits. We show the five year CDS spreads of the individual credits



The basket spread

One way to view an FTD basket is as a trade in which the investor sells protection on all of the credits in the basket with the condition that all the other CDS cancel at no cost following a credit event. Such a trade cannot be replicated using existing instruments. Valuation therefore requires a pricing model. The model inputs in order to determine the nth-to-default basket spread are:

- **Value of n:** An FTD ($n=1$) is riskier than an STD ($n=2$) and so commands a higher spread.
- **Number of credits:** The greater the number of credits in the basket, the greater the likelihood of a credit event, and so the higher the spread.
- **Credit quality:** The lower the credit quality of the credits in the basket, in terms of spread and rating, the higher the spread.

- **Maturity:** The effect of maturity depends on the shape of the individual credit curves.
- **Recovery rate:** This is the expected recovery rate of the nth-to-default asset following its credit event. This has only a small effect on pricing since a higher expected recovery rate is offset by a higher implied default probability for a given spread. However, if there is a default the investor will certainly prefer a higher realised recovery rate.
- **Default correlation:** Increasing default correlation increases the likelihood of assets to default or survive together. The effect of default correlation is subtle and significant in terms of pricing. We now discuss this in more detail.


Baskets and default correlation

Baskets are essentially a default correlation product. This means that the basket spread depends on the tendency of the reference assets in the basket to default together.

It is natural to assume that assets issued by companies within the same country and industrial sector should have a higher default correlation than those within different industrial sectors. After all, they share the same market, the same interest rates and are exposed to the same costs. At a global level, all companies are affected by the performance of the world economy. We believe that these systemic sector risks far outweigh idiosyncratic effects so we expect that default correlation is usually positive.

There are a number of ways to explain how default correlation affects the pricing of default baskets. Confusion is usually caused by the term 'default correlation'. The fact is

Figure 5. Loss distribution for a five-credit basket with 0%, 20% and 50% correlation



that if two assets are correlated, they will not only tend to default together, they will also tend to survive together.

There are two correlation limits in which a FTD basket can be priced without resorting to a model – independence and maximum correlation.

□ **Independence:** Consider a five-credit basket where all of the underlying credits have flat credit curves. If the credits are all independent and never become correlated during the life of the trade, the natural hedge is for the basket investor to buy CDS protection on each of the individual names to the full notional. If a credit event occurs, the CDS hedge covers the loss on the basket and all of the other CDS hedges can be unwound at no cost, since they should on average have rolled down their flat credit curves. This implies that the basket spread for independent assets should be equal to the sum of the spreads of the names in the basket.

□ **Maximum correlation:** Consider the same FTD basket but this time where the

default correlation is at its maximum. In practice, this means that when any asset defaults, the asset with the widest spread will always default too. As a result, the risk of one default is the same as the risk of the widest spread asset defaulting. Because an FTD is triggered by only one credit event, it will be as risky as the riskiest asset and the FTD basket spread should be the widest spread of the credits in the basket.

The best way to understand the behaviour of default baskets between these two correlation limits is to study the loss distribution for the basket portfolio. See page 33 for a discussion of how to model the loss distribution.

We consider a basket of five credits with spreads of 100bp and an assumed recovery rate for all of 40%. We have plotted the loss distribution for correlations of 0%, 20%, and 50% in Figure 5. The spread for an FTD basket depends on the probability of one or more defaults which equals one minus the probability of no defaults. We see that the probability of no defaults increases with increasing correlation – the probability of credits surviving together increases – and the FTD spread should fall.

The risk of an STD basket depends on the probability of two or more defaults. As correlation goes up from 0–20%, the probability of two, three, four and five defaults increases. This makes the STD spread increase.

The process for translating these loss distributions into a fair value spread requires a model of the type described on page 39. Essentially we have to find the basket spread for which the present value of the protection payments equals the present value of the premium payments.

We should not forget that in addition to the


protection leg, the premium leg of the default basket also has correlation sensitivity because it is only paid for as long as the n th default does not occur.

Using a model we have calculated the correlation sensitivity of the FTD and STD spread for the five-credit basket shown in Figure 6. At low correlation, the FTD spread is close to 146bp, which is the sum of the spreads. At high correlation, the basket has the risk of the widest spread asset and so is at 30bp. The STD spread is lowest at zero correlation since the probability of two assets defaulting is low if the assets are independent. At maximum correlation the STD spread tends towards the spread of the second widest asset in the basket which is also 30bp.

Applications

- ☐ Baskets have a range of applications. Investors can use default baskets to leverage their credit exposure and so earn a higher yield without increasing their notional at risk.
- ☐ The reference entities in the basket are all typically investment grade and so are familiar to most credit analysts.

Figure 6. Correlation dependence of spread for FTD and STD basket



- ☐ The basket can be customised to the investors' exact view regarding size, maturity, number of credits, credit selection, FTD or STD.
- ☐ Buy and hold investors can enjoy the leveraging of the spread premium. This is discussed in more detail later.
- ☐ Credit investors can use default baskets to hedge a blow-up in a portfolio of credits more cheaply than buying protection on the individual credits.
- ☐ Default baskets can be used to express a view on default correlation. If the investor's view is that the implied correlation is too low then the investor should sell FTD protection. If implied correlation is too high they should sell STD protection.

Hedging default baskets

The issuers of default baskets need to hedge their risks. Spread risk is hedged by selling protection dynamically in the CDS market on all of the credits in the default basket. Determining how much to sell, known as the delta, requires a pricing model to calculate the sensitivity of the basket value to changes in the spread curve of the underlying credit.

Although this delta hedging should immunise the dealer's portfolio against small changes in spreads, it is not guaranteed to be a full hedge against a sudden default. For instance, a dealer hedging an FTD basket where a credit defaults with a recovery rate of R would receive a payment of $(1-R)F$ from the protection seller, and will pay $D(1-R)F$ on the hedged protection, where F is the basket face value and D is the delta in terms of percentage of face value. The net payment to the protection buyer is therefore $(1-D)(1-R)F$.

There will also probably be a loss on the other CDS hedges. The expected spread widening on default on the other credits in the basket due to their positive correlation with the defaulted asset will result in a loss when they are unwound. The greater unwind losses for baskets with higher correlations will be factored into the basket spread.

One way for a default basket dealer to reduce his correlation risk is by selling protection on the same or similar default baskets. However this is difficult as it is usually difficult to find protection buyers who select the exact same basket as an investor.

The alternative hedging approach is for the dealer to buy protection using default baskets on other orders of protection. This is based on the observation that a dealer who is long first, second, third up to Mth order protection on an M-credit basket has almost no correlation risk, since this position is almost economically equivalent to buying full face value protection using CDS on all M credits in the basket.

Figure 7 shows an example basket with the delta and spread for each of the five credits. Note that the deltas are all very similar. This reflects the fact that all of the assets have a similar spread. Differences are mainly due to our different correlation assumptions.

Figure 7. Default basket deltas for a €10m notional five-year FTD basket on five credits. The FTD spread is 246bp.

Reference entity	CDS Spread	Delta
Walt Disney	62bp	6.26m
Rolls Royce	60bp	6.55m
Sun Microsystems	60bp	6.87m
Eastman Chemical	60bp	7.16m
France Telecom	64bp	7.57m

Hedgers of long protection FTD baskets are also long gamma. This means that as the spread of an asset widens, the delta will increase and so the hedger will be selling protection at a wider spread. If the spread tightens, then the delta will fall and the hedger will be buying back hedges at a tighter level. So spread volatility can be beneficial. This effect helps to offset the negative carry associated with hedged FTD baskets. This is clear in the previous example where the income from the hedges is 211bp, lower than the 246bp paid to the FTD basket investor.

Different rating agencies have developed their own model-based approaches for the rating of default baskets. We discuss these on page 39.


Synthetic CDOs

Synthetic collateralised debt obligations (Synthetic CDOs) were conceived in 1997 as a flexible and low-cost mechanism for transferring credit risk off bank balance sheets. The primary motivation was the banks' reduction of regulatory capital.

More recently, however, the fusion of credit derivatives modelling techniques and derivatives trading have led to the creation of a new type of synthetic CDO, which we call a customised CDO, which can be tailored to the exact risk appetites of different classes of investors. As a result, the synthetic CDO has become an investor-driven product.

Overall, these different types of synthetic CDO have a total market size estimated by the *Risk* 2003 survey to be close to \$500 billion. What is also of interest is that the dealer-hedging of these products in the CDS market has generated a substantial demand to sell protection, balancing the traditional protection-buying demand coming from bank loan book managers.

Figure 8. A standard synthetic CDO



The performance of a synthetic CDO is linked to the incidence of default in a portfolio of CDS. The CDO redistributes this risk by allowing different tranches to take these default losses in a specific order. To see this, consider the synthetic CDO shown in Figure 8. It is based on a reference pool of 100 CDS, each with a €10m notional. This risk is redistributed into three tranches; (i) an equity tranche, which assumes the first €50m of losses, (ii) a mezzanine tranche, which take the next €100m of losses, and (iii) the senior tranche with a notional of €850m takes all remaining losses.

The equity tranche has the greatest risk and is paid the widest spread. It is typically unrated. Next is the mezzanine tranche which is lower risk and so is paid a lower spread. Finally we have the senior tranche which is protected by €150m of subordination. To get a sense of the risk of the senior tranche, note that it would require more than 25 of the assets in the 100 credit portfolio to default with a recovery rate of 40% before the senior tranche would take a principal loss. Consequently the senior tranche is typically paid a very low spread.

The advantage of CDOs is that by changing the details of the tranche in terms of its attachment point (this is the amount of sub-

ordination below the tranche) and width, it is possible to customise the risk profile of a tranche to the investor's specific profile.

Full capital structure synthetics

In the typical synthetic CDO structured using securitisation technology, the sponsoring institution, typically a bank, enters into a portfolio default swap with a Special Purpose Vehicle (SPV). This is shown in Figure 9 (overleaf).

The SPV typically provides credit protection for 10% or less of the losses on the reference portfolio. The SPV in turn issues notes in the capital markets to cash collateralise the portfolio default swap with the originating entity. The notes issued can include a non-rated 'equity' piece, mezzanine debt and senior debt, creating cash liabilities. The remainder of the risk, 90% or more, is generally distributed via a senior swap to a highly rated counterparty in an unfunded format.


Reinsurers, who typically have AAA/AA ratings, have traditionally had a healthy appetite for this type of senior risk, and are the largest participants in this part of the capital structure – often referred to as super-senior AAAs or super-senior swaps. The initial proceeds from the sale of the equity and notes are invested in highly rated, liquid assets.

If an obligor in the reference pool defaults, the trust liquidates investments in the trust and makes payments to the originating entity to cover default losses. This payment is offset by a successive reduction in the equity tranche, then the mezzanine and finally the super-seniors are called to make up losses. See Ganapati et al (2001) for more details.

Mechanics of a synthetic CDO

When nothing defaults in the reference portfolio of the CDO, the investor simply

Figure 9. The full capital structure synthetic CDO



receives the Libor spread until maturity and nothing else changes. Using the synthetic CDO described earlier and shown in Figure 8, consider what happens if one of the reference entities in the reference portfolio undergoes the first credit event with a 30% recovery, causing a €7m loss.

The equity investor takes the first loss of €7m, which is immediately paid to the originator. The tranche notional falls from €50m to €43m and the equity coupon, set at 1500bp, is now paid on this smaller notional. These coupon payments therefore fall from €7.5m to 15% times €43m = €6.45m.

If traded in a funded format, the €3m recovered on the defaulted asset is either reinvested in the portfolio or used to reduce the exposure of the senior-most tranche (similar to early amortisation of senior tranches in cash flow CDOs).

The senior tranche notional is decreased by €3m to €847m, so that the sum of protected notional equals the sum of the collateral notional which is now €990m. This has no effect on the other tranches.

This process repeats following each credit event. If the losses exceed €50m then the

mezzanine investor must bear the subsequent losses with the corresponding reduction in the mezzanine notional. If the losses exceed €150m, then it is the senior investor who takes the principal losses.

The mechanics of a standard synthetic CDO are therefore very simple, especially compared with traditional cash flow CDO waterfalls. This also makes them more easily modelled and priced.

The CDO tranche spread

The synthetic CDO spread depends on a number of factors. We list the main ones and describe their effects on the tranche spread.

- **Attachment point:** This is the amount of subordination below the tranche. The higher the attachment point, the more defaults are required to cause tranche principal losses and the lower the tranche spread.
- **Tranche width:** The wider the tranche for a fixed attachment point, the more losses to which the tranche is exposed. However, the incremental risk ascending

the capital structure is usually declining and so the spread falls.

- **Portfolio credit quality:** The lower the quality of the asset portfolio, measured by spread or rating, the greater the risk of all tranches due to the higher default probability and the higher the spread.
- **Portfolio recovery rates:** The expected recovery rate assumptions have only a secondary effect on tranche pricing. This is because higher recovery rates imply higher default probabilities if we keep the spread fixed. These effects offset each other to first order.
- **Swap maturity:** This depends on the shapes of the credit curves. For upward sloping credit curves, the tranche curve will generally be upward sloping and so the longer the maturity, the higher the tranche spread.
- **Default correlation:** If default correlation is high, assets tend to default together and this makes senior tranches more risky. Assets also tend to survive together making the equity safer. To understand this more fully we need to better understand the portfolio loss distribution.

The portfolio loss distribution

No matter what approach we use to generate it, the loss distribution of the reference portfolio is crucial for understanding the risk and value of correlation products. The portfolio loss is clearly not symmetrically distributed: it is therefore informative to look at the entire loss distribution, rather than summarising it in terms of expected value and standard deviation. We can use models of the type discussed on page 33 to calculate

the portfolio loss distribution. We can expect to observe one of the two shapes shown in Figure 10. They are (i) a skewed bell curve; (ii) a monotonically decreasing curve.

The skewed bell curve applies to the case when the correlation is at or close to zero. In this limit the distribution is binomial and the peak is at a loss only slightly less than the expected loss.

As correlation increases, the peak of the distribution falls and the high quantiles increase: the curves become monotonically decreasing. We see that the probability of larger losses increases and, at the same time, the probability of smaller losses also increases, thereby preserving the expected loss which is correlation independent (for further discussion see Mashal, Naldi and Pedersen (2003)).

For very high levels of asset correlations (hardly ever observed in practice), the distribution becomes U-shaped. At maximum default correlation all the probability mass is located at the two ends of the distribution. The portfolio either all survives or it all defaults. It resembles the loss distribution of a single asset.

Figure 10. Portfolio loss distribution for a large portfolio at 0%, 20% and 95% correlation




Figure 11. Equity tranche loss distribution for correlations of 20% and 50%





Figure 12. Mezzanine tranche loss distribution for correlation of 20% and 50%. We have eliminated the zero loss peak, which is about 86% in both cases



How then does the shape of the portfolio loss distribution affect the pricing of tranches? To see this we must study the tranche loss distribution.

The tranche loss distribution

We have plotted in Figures 11–13 the loss distributions for a CDO with a 5% equity, 10% mezzanine and 85% senior tranche for correlation values of 20% and 50%. At 20% corre-

lation, we see that most of the portfolio loss distribution is inside the equity tranche, with about 14% beyond, as represented by the peak at 100% loss. As correlation goes to 50% the probability of small losses increases while the probability of 100% losses increases only marginally. Clearly equity investors benefit from increasing correlation.

The mezzanine tranche becomes more risky at 50% correlation. As we see in Figure 12, the 100% loss probability jumps from 0.50% to 3.5%. In most cases mezzanine investors benefit from falling correlation – they are short correlation. However, the correlation directionality of a mezzanine tranche depends upon the collateral and the tranche. In certain cases a mezzanine tranche with a very low attachment point may be a long correlation position.

Senior investors also see the risk of their tranche increase with correlation as more joint defaults push out the loss tail. This is clear in Figure 13. Senior investors are short correlation.

In Figure 14 we plot the dependence of the value of different CDO tranches on correlation. As expected, we clearly see that:

- ☐ Senior investors are short correlation. If correlation increases, senior tranches fall in value.
- ☐ Mezzanine investors are typically short correlation, although this very much depends upon the details of the tranche and the collateral.
- ☐ Equity investors are long correlation. When correlations go up, equity tranches go up in value.

In the process of rating CDO tranches, rating agencies need to consider all of these

risk parameters and so have adopted model based approaches. These are discussed on page 43.

Customised synthetic CDO tranches

Customisation of synthetic tranches has become possible with the fusion of derivatives technology and credit derivatives. Unlike full capital structure synthetics, which issue the equity, mezzanine and senior parts of the capital structure, customised synthetics may issue only one tranche. There are a number of other names for customised CDO tranches, including bespoke tranches, and single tranche CDOs.

The advantage of customised tranches is that they can be designed to match exactly the risk appetite and credit expertise of the investor. The investor can choose the credits in the collateral, the trade maturity, the attachment point, the tranche width, the rating, the rating agency and the format (funded or unfunded). Execution of the trade can take days rather than the months that full capital structure CDOs require.

The basic paradigm has already been discussed in the context of default baskets. It is to use CDS to dynamically delta-hedge the first order risks of a synthetic tranche and to use a trading book approach to hedge the higher order risks. This is shown in Figure 15 (overleaf).

For example, consider an investor who buys a customised mezzanine tranche from Lehman Brothers. We will then hedge it by selling protection in an amount equal to the delta of each credit in the portfolio via the CDS market. The delta is the amount of protection to be sold in order to immunise the portfolio against small changes in the CDS spread curve for that credit. Each credit in the portfolio will have its own delta.

Figure 13. Senior tranche loss distribution for correlations of 20% and 50%. We have eliminated the zero loss peak, which is greater than 96% in both cases





Figure 14. Correlation dependence of CDO tranches




Understanding delta for CDOs

For a specific credit in a CDO portfolio, the delta is defined as the notional of CDS for that credit which has the same mark-to-market change as the tranche for a small movement in the credit's CDS spread curve. Although the definition may be straightforward, the behaviour of the delta is less so.

One way to start thinking about delta is to

Figure 15. Delta hedging a synthetic CDO



imagine a queue of all of the credits sorted in the order in which they should default. This ordering will depend mostly on the spread of the asset relative to the other credits in the portfolio and its correlation relative to the other assets in the portfolio. If the asset whose delta you are calculating is at the front of this queue, it will be most likely to cause losses to the equity tranche and so will have a high delta for the equity tranche. If it is at the back of the queue then its equity delta will be low. As it is most likely to default after all the other asset, it will be most likely to hit the senior tranche. As a result the senior tranche delta will rise. This framework helps us understand the directionality of delta.

The actual magnitude of delta is more difficult to quantify because it depends on the tranche notional and the contractual tranche spread, as well as the features of the asset whose delta we are examining. For example the delta for a senior tranche to a credit whose CDS spread has widened will fall due to the fact that it is more likely

to default early and hit the equity tranche, and also because the CDS will have a higher spread sensitivity and so require a smaller notional.

To show this we take an example CDO with 100 credits, each \$10m notional. It has three tranches: a 5% equity, a 10% mezzanine and an 85% senior tranche. The asset spreads are all 150bp and the correlation between all the assets is the same.

The sensitivity of the delta to changing the spread of the asset whose delta we are calculating is shown in Figure 16. If the single asset spread is less than the portfolio average of 150bp, then it is the least risky asset. As a result, it would be expected to be the last to default and so most likely to impact the senior-most tranche. As the spread of the asset increases above 150bp, it becomes more likely to default before the others and so impacts the equity or mezzanine tranche. The senior delta drops and the equity delta increases.

In Figure 17 we plot the delta of the asset versus its correlation with all of the other

assets in the portfolio. These all have a correlation of 20% with each other. If the asset is highly correlated with the other assets it is more likely to default or survive with the other assets. As a result, it is more likely to default en masse, and so senior and mezzanine tranches are more exposed. For low correlations, if it defaults it will tend to do so by itself while the rest of the portfolio tends to default together. As a result, the equity tranche is most exposed.

There is also a time effect. Through time, senior and mezzanine tranches become safer relative to equity tranches since less time remains during which the subordination can be reduced resulting in principal losses. This causes the equity tranche delta to rise through time while the mezzanine and senior tranche deltas fall to zero.


Building intuition about the delta is not trivial. There are many further dependencies to be explored and we intend to describe these in a forthcoming paper.

Higher order risks

If properly hedged, the dealer should be insensitive to small spread movements. However, this is not a completely risk-free position for the dealer since there are a number of other risk dimensions that have not been immunised. These include correlation sensitivity, recovery rate sensitivity, time decay and spread gamma. There is also a risk to a sudden default which we call the value-on-default risk (VOD).

For this reason, dealers are motivated to do trades that reduce these higher order risks. The goal is to flatten the risk of the correlation book with respect to these higher order risks either by doing the offsetting trade or by placing different parts of the capital structure with other buyers of customised tranches.


Figure 16. Dependency of tranche delta on the spread of the asset



Idiosyncratic versus systemic risk

In terms of how they are exposed to credit, there is a fundamental difference between equity and senior tranches. Equity tranches are more exposed to idiosyncratic risk – they incur a loss as soon as one asset defaults. The portfolio effect of the CDO is only expressed through the fact that it may take several defaults to completely reduce the equity notional. This implies that equity investors should focus less on the overall properties of the collateral, and more on trying to choose assets which they believe will

Figure 17. Dependency of tranche delta on the asset's correlation with the rest of the portfolio



not default. As a result we would expect equity tranche buyers to be skilled credit investors, able to pick the right credits for the portfolio, or at least be able to hedge the credits they do not like.

On the other hand, the senior investor has a significant cushion of subordination to insulate them from principal losses until maybe 20 or more of the assets in the collateral have defaulted. As a consequence, the senior investor is truly taking a portfolio view and so should be more concerned about the average properties of the collateral than the quality of any specific asset. The senior tranche is really a deleveraged macro credit trade.

Evolution of structures

Initially full capital structure synthetic CDOs had almost none of the structural features typically found in other securitised asset classes and cash flow CDOs. It was only in 1999 that features that diverted cash flows from equity to debt holders in case of certain covenant failures began entering the landscape. The intention was to provide some defensive mechanism for mezzanine holders fearing that the credit cycle would affect tranche performance. Broadly, these features fit into two categories – ones that build extra subordination using excess spread, and others that use excess spread to provide upside participation to mezzanine debt holders.

The most common example of structural ways to build additional subordination is the reserve account funding feature. Excess spread (the difference between premium received from the CDS portfolio and the tranche liabilities) is paid into a reserve account. This may continue throughout the life of the deal or until the balance reaches a predetermined amount. If structured to accumulate to maturity, the equity tranche

will usually receive a fixed coupon throughout the life of the transaction and any upside or remainder in the reserve account at maturity. If structured to build to a predetermined level, the equity tranche will usually receive excess interest only after the reserve account is fully funded. More details are provided in Ganapati and Ha (2002).

Other structures incorporated features to share some of the excess spread with mezzanine holders or to provide a step up coupon to mezzanines if losses exceeded a certain level or if the tranche was downgraded. Finally, over-collateralisation trigger concepts were adopted from cash flow CDOs.

Principal protected structures

Investors who prefer to hold highly rated assets can do so by purchasing CDO tranches within a principal protected structure. This is designed to guarantee to return the investor's initial investment of par. One particular variation on this theme is the Lehman Brothers High Interest Principal Protection with Extendible Redemption (HIPER). This is typically a 10-year note which pays a fixed coupon to the investor linked to the risk of a CDO equity tranche.

This risk is embedded within the coupons of the note such that each default causes a reduction in the coupon size. However the investor is only exposed to this credit risk for a first period, typically five years, and the coupon paid for the remaining period is frozen at the end of year five. The coupon is typically of the form:

$$Coupon = 8\% \times \max \left[1 - \frac{Portfolio\ loss}{Tranche\ size}, 0 \right]$$

In Figure 18 we show the cash flows

assuming two credit events over the lifetime of the trade. The realised return is dependent on the timing of credit events. For a given number of defaults over the trade maturity, the later they occur, the higher the final return.

Managed synthetics


The standard synthetic has been based on a static CDO, ie, the reference assets in the portfolio do not change. However, recently, Lehman Brothers and a number of other dealers have managed to combine the customised tranche with the ability for an asset manager or the issuer of the tranche to manage the portfolio of reference entities. This enables investors to enjoy all the benefits of customised tranches and the benefits of a skilled asset manager. The customisable characteristics include rating, rating agency, spread, subordination, issuance format plus others.

The problem with this type of structure is that the originator of the tranche has to factor into the spread the cost of substituting assets in the collateral. Initially this was based on the asset manager being told the cost of substituting an asset using some black-box approach.

More recently the format has evolved to one where the manager can change the portfolio subject to some constraints. One example of such technology is Lehman Brothers' DYNAMO structure. The advantage of this approach is that it frees the manager to focus on the credits without having to worry about the cost of substitution.

The other advantages of such a structure for the asset manager are fees earned and an increase in assets under management. For investors the incentive is to leverage the management capabilities of a credit asset manager in order to avoid blow-ups in the

Figure 18. The HIPER structure



portfolio and so better manage downturns in the credit cycle.

The CDO of CDOs


A recent extension of the CDO paradigm has been the CDO of CDOs, also known as 'CDO squared'. Typically this is a mezzanine 'super' tranche CDO in which the collateral is made up of a mixture of asset-backed securities and several 'sub' tranches of synthetic CDOs. Principal losses are incurred if the sum of the principal losses on the underlying portfolio of synthetic tranches exceeds the attachment point of the super-tranche. Looking forward, we see growing interest in synthetic-only portfolios.

Leveraging the spread premium

Market spreads paid on securities bearing credit risk are typically larger than the levels implied by the historical default rates for the same rating. This difference, which we call the spread premium, arises because investors demand compensation for being exposed to default uncertainty, as well as other sources of risk, such as spread movements, lack of liquidity or ratings downgrades.

Portfolio credit derivatives, such as basket default swaps and synthetic CDO tranches, offer a way for investors to take advantage of this spread premium. When an investor

Figure 19. Spread premium for an FTD compared with a Ba3 single-name asset



sells protection via a default basket or a CDO tranche, the note issuer passes this on by selling protection in the CDS market. This hedging activity makes it possible to pass this spread premium to the buyer of the structured credit asset. For buy and hold credit investors the spread premium paid can be significant and it is possible to show, see O’Kane and Schloegl (2003) for details of the method, that under certain criteria, these assets may be superior to single-name credit investments.

Our results show that an FTD basket leverages the spread premium such that the size of the spread premium is much higher for an FTD basket than it is for a single-credit asset paying a comparable spread. This is shown in Figure 19 where we see that an FTD basket paying a spread of 350bp has around 290bp of spread premium. Compare this with a single-credit Ba3 asset also paying a spread close to 340bp. This has only 70bp of spread premium.

For an STD basket we find that the spread premium is not leveraged. Instead, it is the ratio of spread premium to the whole spread which goes up. There are therefore two conclusions:

1. FTD baskets leverage spread premium. This makes them suitable for buy and hold yield-hungry investors who wish to be paid a high spread but also wish to minimise their default risk.
2. STD baskets leverage the ratio of spread premium to the market spread. This is suitable for more risk-averse investors who wish to maximise return per unit of default risk.

We therefore see that default baskets can appeal to a range of investor risk preferences.

CDO tranches exhibit a similar leveraging of the premium embedded in CDS spreads. The advantage of CDO of CDOs is that they provide an additional layer of leverage to the traditional CDO. This can make leveraging the spread premium arguments even more compelling.

The conclusion is that buy-and-hold correlation investors are overcompensated for their default risk compared with single-name investors.

CDO strategies

Investors in correlation products should primarily view them as buy and hold investments which allow them to enjoy the spread premium. This is a very straightforward strategy for mezzanine and senior investors. However, for equity investors, there are a number of strategies that can be employed in order to dynamically manage the idiosyncratic risk. We list some strategies below.

1. The investor buys CDO equity and hedges the full notional of the 10 or so worst names. The investor enjoys a significant positive carry and at the same time reduces his idiosyncratic default risk. The investor may also sell CDS protection on the tightest names, using the

income to offset some of the cost of protection on the widest names.

2. The investor may buy CDO equity and delta hedge. The net positive gamma makes this trade perform well in high spread volatility scenarios. By dynamically re-hedging, the investor can lock in this convexity. The low liquidity of CDOs means that this hedging must continue to maturity.
3. The investor may use the carry from CDO equity to over-hedge the whole portfolio, creating a cheap macro short position. While this is a negative carry trade, it can be very profitable if the market widens dramatically or if a large number of defaults occur.

For more details see Isla (2003).

Credit options

Activity in credit options has grown substantially in 2003. From a sporadic market driven mostly by one-off repackaging deals, it has extended to an increasingly vibrant market in both bond and spread options, options on CDS and more recently options on portfolios and even on CDO tranches.

This growth of the credit options market has been boosted by declines in both spread levels and spread volatility. The reduction in perceived default risk has made hedge funds, asset managers, insurers and proprietary dealer trading desks more comfortable with the spread volatility risks of trading options and more willing to exploit their advantages in terms of leverage and asymmetric payoff.

The more recent growth in the market for options on CDS has also been driven by the increased liquidity of the CDS market, enabling investors to go long or short the option delta amount.

Hedge funds have been the main growth user of credit options, using them for credit arbitrage and also for debt-equity strategies. They are typically buyers of volatility, hedging in the CDS market and exploiting the positive convexity. Asset managers seeking to maximise risk-adjusted returns are involved in yield-enhancing strategies such as covered call writing. Bank loan portfolio managers are beginning to explore default swaptions as a cheaper alternative to buying outright credit protection via CDS.


One source of credit optionality is the cash market. Measured by market value weight, 5.6% of Lehman Brothers US Credit Index and 54.7% of Lehman Brothers US High Yield Index have embedded call or put options. Hence, two strategies which have been, and continue to be important in the bond options market are the repack trade and put bond stripping.

The repack trade

The first active market in credit bond options was developed in the form of the repack trade, spearheaded by Lehman Brothers and several other dealers. Figure 20 (overleaf) shows the schematic of one such transaction.

In a typical repack trade, Lehman Brothers purchased \$32,875,000 of the Motorola (MOT) 6.5% 2028 debentures and placed them into a Lehman Trust called CBTC. The Trust then issued \$25 par class A-1 Certificates to retail investors with a coupon set at 8.375% – the prevailing rate for MOT in the retail market at the time. Since the 8.375% coupon on the CBTC trust is higher than the coupon on the MOT Bond, the CBTC trust must be over-collateralised with enough face value of MOT bonds to pay the 8.375% coupon. An A-2 Principal Only (PO) tranche captures the excess principal. Both class A-1

Figure 20. Mechanics of MOT 6.5% 15/11/28 repack transaction



and class A-2 certificates are issued with an embedded call.

This call option was sold separately to investors in a form of a long-term warrant. The holder of the MOT call warrant has the right but not the obligation to purchase the MOT bonds from the CBTC trust beginning on 19/7/07 and thereafter at the preset call strike schedule. This strike is determined by the proceeds needed to pay off the A-1 certificate at par plus the A-2 certificate at the accreted value of the PO. Because retail investors are willing to pay a premium for the par-valued low-notional bonds of well-known high quality issuers, the buyers of the call warrant can use this structure to source volatility at attractive levels.

Put bond stripping

According to Lehman Brothers' US Credit Index, bonds with embedded puts constitute approximately 2.3% of the US credit bond market, by market value. These bonds grant the holder the right, but not the obligation to sell the bond back to the issuer at a predetermined price (usually par) at one or more future dates.

This option can be viewed as an extension option since by failing to exercise it, the bond maturity is extended. In the past several years, the market has priced these bonds as though they matured on the first put date and has not given much value to the extension option. Recently, credit investors have realised a way to extract this extension risk premium via a put bond stripping strategy.

Essentially, an investor can buy the put bond, and sell the call option to the first put date at a strike price of par. Thus, the investor has a long position in the bond coupled with a synthetic short forward (long put plus short call) with a maturity coinciding with the first put date. He then hedges this position by asset swapping the bond to the put date, effectively eliminating all of the interest rate risk and locking in the cheap volatility. Given the small amount of outstanding put bonds, this strategy has led to more efficient pricing of the optionality in these securities.

Bond options

There are a variety of bond options traded in the market. The two most important ones for investors are:

Price-based options: at exercise, the option holder pays a fixed amount (strike price) and receives the underlying bond – the payoff is proportional to the difference between the price of the bond and the strike price. Examples of actively trading price options are Brady bond options, corporate bond options, CBTC call warrants and calls on put bonds. See the illustration in Figure 21.


Spread-based options: at exercise, the option holder pays an amount equal to the value of the underlying bond calculated using the strike spread and receives the underlying bond – the payoff is proportional to difference between the underlying spread and the strike spread (to a first order of approximation). Spread options can be structured using spreads to benchmark Treasury bonds, default swap spreads or asset swap spreads.

The exercise schedule can be a single date (European), multiple pre-specified dates (Bermudan) or any date in a given range (American). Currently, the most active trading occurs in short-term (less than 12 months to expiry) European-style price options on bonds.

Two of the most common strategies using price options on bonds are covered calls and naked puts. They can be considered respectively as limit orders to sell or buy the underlying bond at a predetermined price (the option strike) on a predetermined day (the option expiry date).

Covered call strategy: an investor who owns the underlying bond sells an out-of-the-money call on the same face value, receiving an upfront premium. If the bond price on the expiry date is greater than the strike, the investor delivers the bonds and receives the strike price. The option premium offsets the

Figure 21. Three month price call option on F 7.25% 25/10/11, struck at 100.76% price.



investor's loss of upside on the price. If the price is less than the strike the investor keeps the bonds and the premium.

Naked put strategy: an investor writes an out-of-the-money put on a bond which he does not own but would like to buy at a lower price. If the bond price on the expiry date is lower than the strike price, it is delivered to the investor. The option premium compensates him for not being able to buy the bond more cheaply in the market. If the bond price is above the option strike price, the investor earns the premium.

In both of these strategies, the main objective for the investor is to find a strike price at which he is willing to buy or sell the underlying bond and which provides sufficient premium to compensate for the potential upside that he forgoes.

Default swaptions and callable CDS

An exciting development in the credit derivatives markets in the past 12 months has been the emergence of default swaptions. These are options on credit default swaps.

The emerging terminology from this mar-

ket is that protection calls (option to buy protection) are called payer default swaptions. Protection puts (option to sell protection) are called receiver default swaptions.

Unlike price options on bonds, the exercise decision for default swaptions is based on credit spread alone. As a result, they are essentially a 'pure' credit product, with pricing being mostly driven by CDS spread volatility.

Default swaptions give investors the opportunity to express views on the future level and variability of default swap spreads for a given issuer. They can be traded outright or embedded in callable CDS. The typical maturity of the underlying CDS is five years but can range from one–10 years, and the time to option expiry is typically three months to one year.

Payer default swaption

The option buyer pays a premium to the option seller for the right but not the obligation to buy CDS protection on a reference entity at a predetermined spread on a future date. Payer default swaptions can be structured with or without a provision for knock out at no cost if there is a credit event between trade date and expiry date. If the knock out provision is included in the swaption, the option buyer who wishes to maintain protection over the entire maturity range can separately buy protection on the underlying name until expiry of the swaption.

The relevant scenarios for this investment are complementary to the ones in the case of the protection put. If spreads tighten by the expiry date, the option buyer will not exercise the right to buy protection at the strike and the option seller will keep the option premium.

Receiver default swaption

In a receiver default swaption, the option buyer pays a premium to the option seller

for the right, but not the obligation, to sell CDS protection on a reference entity at a predetermined spread on a future date. This spread is the option strike.

We do not need to consider what happens if the reference entity experiences a credit event between trade date and expiry date as they would never exercise the option in this case. As a result, there is no need for a knockout feature for receiver default swaptions. Consider the following example.

Lehman Brothers pays 1.20% for an at-the-money receiver default swaption on five-year GMAC, struck at the current five year spread of 265bp and with three months to expiry. The investor is short the option. From the investor's perspective, the relevant scenarios are:

- If five-year GMAC trades above 265bp in three months, Lehman does not exercise, as they can sell protection for a higher spread in the market. The investor has realised an option premium of 1.20% in a quarter of a year.
- If five-year GMAC trades at 238bp in three months, the trade breaks even. (1.20% upfront option premium equals the payoff of $(265\text{bp} - 238\text{bp}) = 27\text{bp}$ times the five-year PV01 of 4.39). If five-year GMAC trades below 238bp in three months, the loss on the exercise of the option will be greater than the upfront premium and the investor will underperform on this trade.

Hedging default swaptions

Dealers hedge these default swaptions using a model of the type discussed on page 49. The underlying in a default swaption is the forward CDS spread from the option expiry date to the maturity date of the CDS. Theoretically, a knock-out payer swaption should be delta

Figure 22. Default swaption types

Product	Payer default swaption	Receiver default swaption
Description	Option to buy protection	Option to sell protection
Exercised if	CDS spread at expiry > strike	CDS spread at expiry < strike
Credit view	Short credit forward	Long credit forward
Knockout	May trade with or without	Not relevant

hedged with a short protection CDS to the final maturity of the underlying CDS and a long protection CDS to the default swaption expiry date. This combination will produce a synthetic forward CDS that knocks out at default before the forward date. In practice, swaptions with expiry of 1 year and less are hedged only with CDS to final maturity due to a lack of liquidity in CDS with short maturities. We have summarised the key features of these different swaption types in Figure 22.

Callable default swaps

In addition to default swaptions, there is a growing interest in callable default swaps. These are a combination of plain vanilla CDS with an embedded short receiver swaption position. The seller of a callable default swap is long credit exposure but this exposure can be terminated by the option buyer at some strike spread on a future date. Consider an example.

Lehman Brothers buys five-year GMAC protection, callable in one year, for 315bp from an investor. The assumed current mid-market spread for five-year GMAC protection is 265bp.

If the four-year GMAC spread in one year is less than the strike spread of 315bp, then Lehman Brothers will exercise the option and so cancel the protection, enabling us to buy protection at the lower market spread. The investor therefore has earned 315bp for selling five-year protection on GMAC for one year.

If the four-year GMAC spread in one year is greater than 315bp, the contract continues and the investor continues to earn 315bp annually.

From the perspective of the option seller, the callable default swap has a limited MTM upside compared with plain vanilla CDS. The additional spread of 315bp–260bp=55bp in this example compensates the option seller for the lost potential upside.

Selling protection in callable default swap is equivalent to a covered call strategy on underlying issuer spreads and is particularly suitable as a yield-enhancement technique for asset managers and insurers.

Credit portfolio options

Starting in mid-2003 market participants have been able to trade in portfolio options whose underlying asset is the TRAC-X North America portfolio with 100 credits. Liquidity is also growing in the European version.

The rationale for options based on TRAC-X is that the portfolio effect will reduce the option volatility and make it easier for dealers to hedge. From an investor perspective it presents a way to take a macro view on spread volatility.

We are now seeing investors trading both at-the-money and out-of-the-money puts and calls to maturities extending from three to nine months. The contracts are typically traded with physical delivery. If the TRAC-X portfolio spread is wider than the strike level on the expiry date, the holder of the

payer default swaption will exercise the option and lock in the portfolio protection at more favourable levels. Conversely, if the TRAC-X spread is tighter than the strike, the holder of the receiver swaption will benefit from exercising the option and realising the MTM gain.

Investors can monetise a view on the future range of market spreads by trading bearish spread (buying at-the-money receiver swaption and selling farther out-of-money receiver swaption) or bullish spread (buying ATM payer swaption and selling farther out-of-money payer swaption) strategies. Other strategies include expressing views on spread changes over a given time horizon by trading calendar spreads (buying near maturity options and selling farther maturity options).

Finally, because the TRAC-X spread is less subject to idiosyncratic spread spikes, and because of the existing two-way markets with varying strikes, investors can express their views on the direction of changes in the macro level of spread volatility by trading straddles, ie, simultaneously buying payer and receiver default swaptions as a way to go long volatility while being neutral to the direction of spread changes.

Hybrid products

Hybrid credit derivatives are those which combine credit risk with other market risks such as interest rate or currency risk. Typically, these are credit event contingent instruments linked to the value of a derivatives payout, such as an interest rate swap or an FX option.

There are various motivations for entering into trades which have these hybrid risks. Below, we give an overview of the economic rationale for different types of structures. We discuss the modelling of hybrid credit derivatives in more detail on page 51.

Clean and perfect asset swaps

One important theme is the isolation of the pure credit risk component in a given instrument. For example, a European CDO investor may wish to access USD collateral without incurring any of the associated currency risks.

Cross-currency asset swaps are the traditional mechanism by which credit investors transform foreign currency fixed-rate bonds into local currency Libor floaters. This has the benefit that it substantially reduces the currency and interest rate risk, converting the bond from an FX, interest rate and credit play into an almost pure credit play.

However, the currency risk has not been completely removed. First, note that a cross currency asset swap is really two trades: (i) purchase of a foreign currency asset; and (ii) entry into a cross-currency swap. In the case of a European investor purchasing a dollar asset, the investor receives Euribor plus a spread paid in euros.

As long as the underlying dollar asset does not default during the life of the asset swap there is no currency risk to the investor. However, if the asset does default, the investor loses the future dollar coupons and principal of the asset, just receiving some recovery amount which is paid in dollars on the dollar face value. As the cross-currency swap is not contingent, meaning that the payments on the swap contract are unaffected by any default of the asset, the investor is therefore obliged to either continue the swap or to unwind it at the market value with a swap counterparty. This unwind value can be positive or negative – the investor can make a gain or loss – depending on the direction of movements in FX and interest rates since the trade was initiated.

The risk is significant. We have modelled

the distribution of the swap MTM at default, shown in Figure 23, for a five year euro-dollar cross-currency swap using market calibrated parameters. The downside risk is significant with possible swap related losses comparable in size to the loss on the defaulted credit.

As a result, the basic cross-currency asset swap has a default contingent interest rate and currency risk. This can be removed through the use of a hybrid. Two variations exist.


1. In the clean asset swap, Lehman Brothers takes on the default contingent swap unwind value risk.
2. In the perfect asset swap, Lehman Brothers takes on the default contingent swap unwind value risk and guarantees (quantoes) the recovery rate of the defaulted asset in the investor's base currency.

Both structures have featured widely in the CDO market where they have been used to immunise mixed-currency high-yield bond portfolios against currency risk in order to allow the structure to qualify for the desired rating from the rating agency.

However, they can also be used by investors to convert foreign currency assets into their base currency. For example, European investors can use the perfect asset swap to take advantage of the often higher spread levels which exist for the same credits when denominated in US dollars.

The cost of removing this default contingent swap MTM risk and paying the recovery rate in the investors domestic currency depends upon a number of factors including the volatility of the FX rate, the credit quality of the reference credit, the shape of the Libor curves in both currencies, interest rate volatilities in both currencies and the

Figure 23. Modelled distribution of the swap MTM at default as a percentage of face value for a five-year euro-dollar swap



correlations between rates, FX and credit.

The cost can be amortised over the life of the trade as a reduction in the asset swap spread paid. It is interesting to note that the reduction in the perfect asset swap spread may not be significant given that the two FX risks to the swap MTM and the recovery rate are actually partially offsetting.

Counterparty risk hybrids

The mitigation of counterparty risk gives rise to another type of hybrid. Consider an investor who enters into an interest rate swap with a credit risky counterparty. Suppose also that this counterparty defaults with a recovery rate of R . If the MTM of the swap is negative from the viewpoint of the investor, the swap is unwound at market value. This means that in an MTM framework the investor incurs no loss from the default event; the swap could be replaced by one with more advantageous terms at zero cost. However, if the value is positive, the investor loses a fraction of this MTM.

A way to mitigate the counterparty risk is therefore to buy protection on the counter-

party, where the contingent payout is linked to the replacement cost of the swap. As we discuss on page 51, the investor is effectively purchasing an interest rate swaption which is automatically exercised upon a default event. Clearly, similar structures can be constructed to provide credit protection for other payouts.

Yield enhancement

In the current interest rate environment with very low short term rates and steep yield curves, coupled with the more mature credit derivatives market, there is increasing investor interest in credit-linked notes with more exotic coupons. These include paying a spread over a Constant Maturity Swap (CMS) rate or a particular inflation index.

Hedge cost mitigation

The final class of application concerns the pure hedging of credit contingent FX or interest rate risks, such as those faced by an international corporation in the course of its business activities. For example, one way to reduce the cost of credit hedging could be to purchase default protection linked to an FX rate being above or below a specified threshold.

Companies looking to hedge interest rate or FX risk may be concerned with the cost of outright hedging using vanilla derivatives. In this case, a hedge which knocks out on the default of a reference credit can provide an adequate hedge while significantly decreasing costs. Clearly, the hedger is implicitly taking a bullish view on the reference credit.

Credit Derivatives Modelling

To be able to price and risk-manage credit derivatives, we need a framework for valuing credit risk at a single issuer level, and at a multi-issuer level. The growth of the credit derivatives market has created a need for more powerful models and for a better understanding of the empirical evidence needed to calibrate these models. In this section we will present a detailed overview of modelling approaches from a practical perspective, ie, we will discuss models, implementation and calibration.

Single credit modelling

The world of credit modelling is divided into two main approaches, called structural and reduced-form. In the structural approach, the default is characterised as the consequence of some event such as a company's asset value being insufficient to cover a repayment of debt. Such models are usually extensions of Merton's 1974 model that used a contingent claims analysis for modelling default.

Structural models are generally used to say at what spread corporate bonds should trade based on the internal structure of the company. They therefore require information about the balance sheet of the company and can be used to establish a link between pricing in the equity and debt markets. However, they are limited in a number of ways including the fact that they generally lack the flexibility to fit exactly a given term structure of spreads; and they cannot be easily extended to price complex credit derivatives.

In the reduced-form approach, the credit process is modelled directly via the probability of the credit event itself. Reduced-form models also generally have the flexibility to refit the prices of a variety of credit instru-

ments of different maturities. They can also be extended to price more exotic credit derivatives. It is for these reasons that they are used for credit derivatives pricing.

The hazard rate approach

The most widely used reduced-form approach is based on the work of Jarrow and Turnbull (1995), who characterise a credit event as the first event of a Poisson counting process which occurs at some time t with a probability defined as

$$\Pr[\tau \leq t + dt | \tau > t] = \lambda(t)dt$$

ie, the probability of a default occurring within the time interval $[t, t+dt)$ conditional on surviving to time t , is proportional to some time dependent function $\lambda(t)$, known as the hazard rate, and the length of the time interval dt . Over a finite time period T , it is possible to show that the probability of surviving is given by

$$Q(0, T) = E^Q \left[\exp \left(- \int_0^T \lambda(s) ds \right) \right]$$

The expectation is taken under the risk-neutral measure. A common assumption is that the hazard rate process is deterministic. By extension, this assumption also implies that the hazard rate is independent of interest rates and recovery rates.

Pricing model for CDS

The breakeven spread in a CDS is the spread at which the present values (PV) of premium and protection legs are equal, ie

$$\text{Premium PV} = \text{Protection PV.}$$

To determine the spread we therefore need to be able to value the protection and premium legs. It is important to take into account the timing of the credit event because this can have a significant effect on the present value of the protection leg and also the amount of premium paid on the premium leg. Within the hazard rate approach we can solve this timing problem by conditioning on the probability of defaulting within each small time interval $[s, s+ds]$, given by $Q(0, s)\lambda(s)ds$, then paying $(1-R)$ and discounting this back to today at the risk-free rate. Assuming that the hazard rate and risk free rate term structures are flat, we can write the value for the protection leg as

$$(1-R) \int_0^T \lambda e^{-(r+\lambda)s} ds = \frac{\lambda(1-R)(1-e^{-(r+\lambda)T})}{r+\lambda}$$

The value of the premium leg is the PV of the spread payments which are made to default or maturity. If we assume that the spread S on the premium leg is paid continuously, we can write the present value of the premium leg as

$$S \int_0^T e^{-(r+\lambda)s} ds = \frac{S(1-e^{-(r+\lambda)T})}{r+\lambda}$$

Equating the protection and premium legs and solving for the breakeven spread gives

$$S = \lambda(1-R)$$

This relationship is known as the credit triangle because it is a relationship between three variables where knowledge of any two is sufficient to calculate the third. It basically states that the spread paid per small time interval exactly compensates the investor

for the risk of default per small time interval. Within this model the interest rate dependency drops out.

Given a CDS which has a flat spread curve at 150bp, and assuming a 50% recovery rate, the implied hazard rate is 0.015 divided by 0.5, which implies a 3% hazard rate. The implied one-year survival probability is therefore $\exp(-0.03)=97.04\%$. For two years it is $\exp(-0.06)=94.18\%$, and so on.

Valuation of a CDS position

The value of a CDS position at time t following initiation at time 0 is the difference between the market implied value of the protection and the cost of the premium payments, which have been set contractually at S_C . We therefore write

$$MTM(t) = \pm (\text{Protection PV} - \text{Premium PV}),$$

where the sign is positive for a long protection position and negative for a short protection position. If the current market spread is given by $S(t)$ then the MTM can be written as

$$MTM(t) = (S(t) - S(0)) \times RPV01$$

where the $RPV01$ is the risky $PV01$ which is given by

$$RPV01(t) = \frac{(1 - e^{-(r+\lambda)(T-t)})}{(r+\lambda)},$$

and where

$$\lambda = \frac{S(t)}{1-R}.$$

An investor buys \$10m of five-year protection at 100bp. One year later, the credit trades at 250bp. Assuming a recovery rate

of 40%, the value is given by substituting, $r=3.0\%$, $R=40\%$, $S(t)=0.025$, $S(0)=0.01$ and $t=4$ into the above equation to give $\lambda=4.17\%$ and an MTM value of \$521,661.

This is a simple yet fairly accurate model which works quite well when the interest rate and credit curves are flat. When this is not the case, it becomes necessary to use bootstrapping techniques to build a full term structure of hazard rates. This may be assumed to be piecewise flat or piecewise linear. For a description of such a model see O'Kane and Turnbull (2003).

Default probabilities

The default probabilities calculated for pricing purposes can be quite different from those calculated from historical default rates of assets with the same rating. These *real-world* default probabilities are generally much lower. The reason for this is that the credit spread of an asset contains not just a compensation for pure default risk; it also depends on the market's risk aversion expressed through a risk premium, as well as on supply-and-demand imbalances.

One should also comment on the market's use of Libor as a risk-free rate in pricing. Pricing theory shows that the price of a derivative is the cost of replicating it in a risk-free portfolio using other securities. Since most market dealers are banks which fund close to Libor, the cost of funding these other securities is also close to Libor. As a consequence it is the effective risk-free rate for the derivatives market.

Calibrating recovery rates

The calibration of recovery rates presents a number of complications for credit derivatives. Strictly speaking, the recovery rate used in the pricing of credit derivatives is the

expected recovery rate following a credit event where the expectation is under the risk-neutral measure.

Such expectations are only available from price information, and the problem in credit is that given one price, it is difficult to separate the probability of default from the recovery rate expectation.

The market standard is therefore to revert to rating agency default studies for estimates of recovery rates. These typically show the average recovery rate by seniority and type of credit instrument, and usually focus on a US corporate bond universe. Adjustments may be made for non-US corporate credits and for certain industrial sectors.


Problems with rating agency recovery statistics include the fact that they are backward looking and that they only include the default and bankruptcy credit events – restructuring is not included. In their favour, one should say that, as they represent the price of the defaulted asset as a fraction of par some 30 days after the default event, they are similar to the definition of the recovery value in a CDS.

Recent work (Altman et al 2001) shows that there is a significant negative correlation between default and recovery rates. One way to incorporate this effect is to assume that recovery rates are stochastic. The standard approach is to use a beta distribution.

Modelling default correlation

By modelling correlation products, we mean modelling products whose pricing depends upon the joint behaviour of a set of credit assets. These include default baskets and synthetic CDOs. As a result of the growth in usage of these products, this is an area of pricing which has recently gained a lot of attention, and in which we have seen a number of significant modelling developments.

Figure 24. Scatterplot of 1,000 simulated asset returns with 0% and 90% correlation. Default thresholds are also shown



In this section we will describe some of these current models, show how they can be applied to the valuation of baskets and CDOs, and towards the end discuss model calibration issues.

Modelling joint defaults

The valuation of default-contingent instruments calls for the modelling of default mechanisms. As discussed earlier, a classical dichotomy in credit models distinguishes between a 'structural approach', where default is triggered by the market value of the borrower's assets (in terms of debt plus equity) falling below its liabilities, and a 'reduced-form approach', where the default event is directly modelled as an unexpected arrival. Although both the structural and the reduced-form approaches can in principle be extended to the multivariate case, structural models calibrated to market-implied default probabilities (often called 'hybrid' models) have gained favour among practitioners because of their tractability in high dimensions.

If we think of defaults as generated by asset values falling below a given boundary, then the probabilities of joint defaults over a specified horizon must follow from the joint

dynamics of asset values: consistent with their descriptive approach of the default mechanism, multivariate structural models rely on the dependence of asset returns in order to generate dependent default events.

This is shown in Figure 24 where we have simulated 1,000 pairs of asset returns modelled as normally distributed random variables, for two firms i and j for two different asset return correlations of 10% and 90%. The vertical and horizontal lines represent default thresholds for firms i and j respectively. Clearly we see that the probability of both i and j defaulting, represented by the number of points in the bottom left quadrant defined by the default thresholds, increases as the asset return correlation increases. Therefore asset correlation leads to default correlation.

Although the pay-offs of multi-credit default-contingent instruments such as nth-to-default baskets and synthetic loss tranches cannot be statically replicated by trading in a set of single-credit contracts, the current market practice is to value correlation products using standard no-arbitrage arguments. It follows that the valuation of these multi-credit exposures boils down to the computation of (risk-neutral) expectations over all possible default scenarios.

A number of different hybrid frameworks have been proposed in the literature for modelling correlated defaults and pricing multi-name credit derivatives. Hull and White (2001) generate dependent default times by diffusing correlated asset values and calibrating default thresholds to replicate a set of given marginal default probabilities. Multi-period extensions of the one-period CreditMetrics paradigm are also commonly used, even if they produce the undesirable serial independence of the realised default rate.¹ While most multi-credit models require simulation, the

need for accurate and fast computation of greeks has pushed researchers to look for modelling alternatives. Finger (1999) and, more recently, Gregory and Laurent (2003) show how to exploit a low-dimensional factor structure and conditional independence to obtain semi-analytical solutions.

Asset and default event correlation

Default event correlation (DEC) measures the tendency of two credits to default jointly within a specified horizon. Formally, it is defined as the correlation between two binary random variables that indicate defaults, ie

$$DEC = \frac{P_{AB} - P_A P_B}{\sqrt{P_A(1 - P_A)} \sqrt{P_B(1 - P_B)}}$$

where p_A and p_B are the marginal default probabilities for credits A and B , and P_{AB} is the joint default probability. Of course, p_A , p_B and P_{AB} all refer to a specific horizon. Notice that default event correlation increases linearly with the joint probability of default and is equal to zero if and only if the two default events are independent. Its limits are not -100% to $+100\%$ but are actually a function of the marginal probabilities themselves.

Default event correlations are the fundamental drivers in the valuation of multi-name credit derivatives. Unfortunately, the scarcity of default data makes joint default probabilities, and thus default event correlations, very hard to estimate directly. As a result, market participants rely on alternative methods to calibrate the frequency of joint defaults within their models. Hybrid models,

such as the simulation approach and the semi-analytical framework described below, use the dependence among asset returns to generate joint defaults, therefore avoiding the need for a direct estimation of joint default probabilities.

Default-time simulation


Monte Carlo models generally aim at the generation of default paths, where each path is simply a list of default times for each of the credits in the reference portfolio drawn at random from the joint default distribution. Knowing the time and identity of each default event allows for a precise valuation of any multi-credit product, no matter how complex the contractual specification of the payoff.

In an influential paper, Li (2000) presents a simple and computationally inexpensive algorithm for simulating correlated defaults. His methodology builds on the implicit assumption that the multivariate distribution of default times and the multivariate distribution of asset returns share the same dependence structure, which he assumes to be Gaussian and is therefore fully characterised by a correlation matrix.

For valuation purposes, we need to sample from the multivariate distribution of default times under the risk-neutral probability measure. In this case, it is common practice to back out the marginal distributions of default times, which we will denote with $F_1, F_2, \dots, F_d$, from single-credit defaultable instruments (such as CDS). We then join these marginal distributions with a correlation matrix, which according to the stated assumptions represents the correlation matrix of the asset returns of the reference credits. Since asset returns are not directly observable, it is common practice to proxy asset correlations using equity correlations. Towards the end of

¹ Finger (2000) offers an excellent comparison of several multivariate models in terms of the default distributions that they generate over time when calibrated to the same marginals and first-period joint default probabilities.

Figure 25. Mapping a normal random variable to a default time



this section we will discuss whether this seems to be a reasonable approximation.

The high-dimensionality of multi-credit instruments means that it is not possible to use market prices to obtain the full dependence structure. Instead, practitioners generally estimate the necessary correlations using historical returns, implicitly relying on the extra assumption that the correlations among asset returns remain unchanged when we move from the objective to the pricing probability measure.

In the bivariate case, the joint distribution function of default times and is simply given by

$$P(\tau_1 < x, \tau_2 < y) = N_{2,\rho}(N^{-1}(F_1(x)), N^{-1}(F_2(y))),$$

where $N_{2,\rho}$ denotes the bivariate cumulative standard Normal with correlation ρ , and N^{-1} indicates the inverse of a cumulative standard Normal. The extension to the d -dimensional case is immediate.

Simulating default times from this distribution is straightforward. With d reference credits and correlation matrix Σ we have the following algorithm:

1. Choleski decomposition of the correlation matrix to simulate a multivariate Normal random vector $X \in R^d$ with correlation Σ .
2. Transform the vector into the unit hypercube using

$$U = (N(x_1), N(x_2), \dots, N(x_d))$$

3. Translate U into the corresponding default times vector t using the inverse of the marginal distributions:

$$\tau = (\tau_1, \tau_2, \dots, \tau_d) = (F_1^{-1}(u_1), F_2^{-1}(u_2), \dots, F_d^{-1}(u_d))$$

The simulation algorithm is illustrated in Figure 25. It is easy to verify that τ has the given marginals and a Normal dependence structure fully characterised by the correlation matrix Σ .

Once we know how to sample from the risk-neutral distribution of default times, it is straightforward to price a correlation trade. Generally speaking, the valuation of a given correlation product always boils down to the computation of an expectation of the form,

$E[f(\tau)]$ where $\tau = (\tau_1, \tau_2, \dots, \tau_d)$ represents the vector of default times and f is a function describing the discounted cash flows (both positive and negative) associated with the instrument under consideration (see, eg, Mashal and Naldi (2002b)).

In summary, Monte Carlo simulation allows for accurate valuations and risk measurements of multi-credit payoffs, even when complex path-dependencies such as reserve accounts, interest coverage or collateralisation tests are involved. This is because both the exact default times and the identity of the defaulters are known on every simulated path. Moreover, we can easily expand the set of variables to be treated as random. For example, Frye (Risk, 2003) argues that modelling stochastic recoveries and allowing for negative correlation between recovery and default rates is essential for a proper valuation of credit derivatives.

The precision and flexibility of this approach, however, come at the cost of computational speed. The basic problem of using simulation is that defaults are rare events, and a large number of simulation paths are usually required to achieve a sufficient sampling of the probability space. There are ways to improve the situation. Useful techniques include antithetic sampling, importance sampling and the use of low-discrepancy sequences. This problem becomes particularly significant when we turn our attention to the calculation of sensitivity measures, since for a reasonable number of paths the simulation noise can be similar to or greater than the price change due to the perturbation of the input parameter. Once again, techniques exist to alleviate the problem, but it is hard to achieve precise hedge ratios in a reasonable amount of time. The question then

arises whether some other numerical approach can be used.

A semi-analytical approach

The recent development of correlation trading and the associated need for fast computation of sensitivities have generated a great deal of interest in semi-analytical models. To enjoy the advantages of fast pricing, one needs to impose more structure on the problem. One way to do this is to rely on two basic simplifications:

1. assume a one-factor correlation structure for asset returns, and
2. discretise the time-line to allow for a finite set of dates at which defaults can happen. These are chosen with a resolution to provide a sufficient level of accuracy.

In particular, 1) makes it possible to compute the risk-neutral loss distribution of the reference portfolio for any given horizon, while 2) is needed to price an instrument knowing only the loss distribution of the reference portfolio at a finite set of dates.

While these two assumptions are sufficient to price plain vanilla portfolio swaps, derivatives structures with more complex path-dependencies may also require that we approximate the payoff function $f(\tau)$. Even in this case, it is usually possible to obtain reasonable approximations, so that the benefit of precise sensitivities can be retained at a relatively low cost. Moreover, the error can be controlled by comparing the analytical solution to a Monte Carlo implementation.

We now discuss how to exploit a one-factor model to construct the risk neutral loss distribution. Later on, we will show how to use a sequence of loss distributions at different horizons to price synthetic loss tranches.

A one-factor model

Let us start by assuming that the asset return of the i th issuer between today and a given horizon is described by a standard Normal random variable A_i with mean 0 and standard deviation of 1, and is of the form

$$A_i = \beta_i Z_M + \sqrt{1 - \beta_i^2} Z_i$$

where $Z_1, Z_2, \dots, Z_D$ are $N(0,1)$ distributed independent random variables. The variable Z_M describes the asset returns due to a common market factor, while Z_i models the idiosyncratic risk of the i th issuer, and β_i stands for the correlation of the asset return of issuer i with the market. The asset return correlation between asset i and asset j is given by $\beta_i \beta_j$.

Default will occur if the realised asset return falls below a given threshold. Mathematically, the i th issuer defaults in the event that $A_i < C_i$. Given that A_i is $N(0,1)$ distributed, it is possible to write

$$p_i = N(C_i) \Rightarrow C_i = N^{-1}(p_i)$$

and calibration of the marginal probabilities is no more complicated than inverting the cumulative normal function.

In the single-issuer case, this is merely an exercise in calibration, as the default thresholds are chosen to reproduce default probabilities which re-price market instruments. The concept of asset returns becomes meaningful when studying the joint behaviour of more than one credit. In this case, the returns are assumed to be described by a multivariate Normal distribution. Using the threshold levels determined before, it is possible to obtain the probabilities of joint defaults.

With a general correlation structure, the calculation of joint default probabilities

becomes computationally intensive, making it necessary to resort to Monte Carlo simulation. However, substantial simplifications can be achieved by imposing more structure on the model.

The advantage of this one-factor setup is that, conditional on Z_M , the asset returns are independent. This makes it easy to compute conditional default probabilities. Conditional on the market factor Z_M , an asset defaults if

$$Z_i \leq \frac{C_i - \beta_i Z_M}{\sqrt{1 - \beta_i^2}}$$

The conditional default probability $p_i(Z_M)$ of an individual issuer i is therefore given by

$$p_i(Z_M) = N\left(\frac{C_i - \beta_i Z_M}{\sqrt{1 - \beta_i^2}}\right)$$

If we assume that the loss on default for each issuer is the same unit (consistent with all assets having the same seniority), u , then building up the portfolio loss distribution can be done iteratively by adding assets to the portfolio. The process is as follows:

1. Beginning with asset one, there are two outcomes on the loss distribution: a loss of zero with a probability $1 - p_1(Z_M)$ and a loss of u with a probability $p_1(Z_M)$.
2. Adding a second asset, we adjust each of the previous losses. The zero loss peak requires that the new asset survives and so has a probability $(1 - p_1(Z_M))(1 - p_2(Z_M))$. A loss u corresponds with the previous zero loss probability times the probability that asset two defaults plus the previous u loss probability times the probability asset two does not default. We arrive at a

probability $(p_1(Z_M) + p_2(Z_M) - 2p_1(Z_M)p_2(Z_M))$. A loss of $2u$ can only occur if the previous loss of u is multiplied by the probability that asset two also defaults to give a probability of $p_1(Z_M)p_2(Z_M)$.

3. We continue adding on to the portfolio until all of the assets have been included. We then repeat with a different value of Z_M and integrate the loss distributions over the $N(0,1)$ distributed market factor Z_M .

The advantage of this approach is that there is no simulation noise. Other algorithms exist, including Fast Fourier techniques which may be more efficient.

Both the times-to-default model and the one-factor model described above are Gaussian copula models. This means that as long as both are calibrated to the same marginals, and as long as both use the same (one-factor) correlation matrix, both models should generate exactly the same prices. This is an important point – it means that models can be categorised by their copula specification alone. It also means that different products may be priced consistently using different default mechanisms, as long as they use the same copula.

Valuation of correlation products

We now explain in detail how these models may be applied to the pricing of correlation products. We start with the default basket.

Pricing default baskets

A default basket is inherently an idiosyncratic product in the sense that the identity of the defaulted asset must be known. One approach is therefore to use the times-to-default model via a Monte Carlo implementation. The generation of default times

within this framework has been described in detail above. Given that we know when each asset in the portfolio defaults in each simulation path, to calculate the fair-value spread we proceed as follows:

1. Sort default times in ascending order and denote the n th time to default as $\tau_n(i)$, where i is the label of the defaulted asset.
2. Calculate the PV of a 1bp coupon stream paid to time $\tau^* = \min[\tau_n(i), T]$ where T is the basket maturity.
3. If $\tau_n(i) < T$ then calculate the PV = $B(\tau_n(i))(1-R(i))$, where B is the Libor discount factor and $R(i)$ is the recovery rate for asset i . Otherwise the protection PV = 0.
4. Average both the premium leg PV of 1bp, known as the basket PV01, and the protection leg PV over all paths.
5. Divide the protection leg by the basket PV01 to get the fair-value spread.

This approach is simple to implement and the size of default baskets, typically $m=5-10$ assets means that pricing is quite fast in Monte Carlo. Monte Carlo is also flexible enough to enable you to introduce stochastic recovery rates, perhaps drawn from a beta distribution. It is also quite straightforward to introduce alternative copulas, see Mashal and Naldi (2002a).

An analytical approach is also possible here. The main constraint is to build the basket n -to-default probabilities while retaining the identity of the defaulted assets. See Gregory and Laurent (2003) for a full discussion.

Rating models for default baskets

Rating agencies have developed a number of models for rating default baskets. For example Moody's has adopted a Monte Carlo-based extension of the asset value approach, in which an asset value is simu-

lated to multiple periods for each of the assets in the basket. These asset values may be correlated internally in order to induce a default correlation. If the asset value falls below the default threshold at the future period, the asset defaults and a recovery rate is drawn from a beta distribution. The recovery amount and asset value can be correlated. The model is calibrated using historical default statistics and the assigned rating is linked to the expected loss of the basket to the trade maturity.

S&P has recently switched its approach from a weakest-link approach which assigns an FTD the rating of the lowest-rated entity in the basket, to one based on a Monte Carlo model, which uses the same framework as its CDO Evaluator model.

Pricing synthetic loss tranches

We would now like to consider in detail how we would set about valuing a loss tranche. One approach is to use the times to default model via a Monte Carlo implementation as discussed earlier. Given that we know when each asset in the portfolio defaults in each simulation path, we would proceed as follows:

1. Calculate the present value of all principal losses on the protection leg in each path.
2. Calculate the PV of 1bp on the premium leg taking care to reduce the notional of the tranche following defaults which cause principal losses in that path.
3. Average both the 1bp premium leg PV, known as the tranche *PV01*, and the protection leg PV over all paths.
4. Divide the protection PV by the tranche *PV01* to compute the breakeven tranche spread.

This approach is simple to implement but

can be computationally slow if there are a large number of assets in the portfolio. Another way is to use a semi-analytical approach which relies on the fact that a standard synthetic loss tranche can be priced directly from its loss distribution.

To see this, consider a portfolio CDS on a tranche defined by the attachment point K_d and the upper boundary K_u , expressed as percentages of the reference portfolio notional. This is N_{port} so that the tranche notional is given by

$$N_{tranche} = N_{port}(K_u - K_d)$$

The maturity of the portfolio swap is T , and for each time $t \leq T$ we denote by $L(t)$ the cumulative portfolio loss up to time t . The tranche loss is therefore given by

$$L_{tranche}(t) = \max[L(t) - K_d, 0] - \max[L(t) - K_u, 0]$$

Note that this is similar to an option style payoff. Indeed it is possible to think of CDO tranches as options on the portfolio loss amount.

The two legs of the swap can be priced in the same way as a CDS if we introduce a tranche 'default probability' $P(t)$ defined as

$$P(t) = \frac{E_0^Q[L_{tranche}(t)]}{N_{tranche}}$$

where we use the risk-neutral (pricing) measure for taking the expectation.

Assuming that the credit, recovery rate and interest rates processes are independent, the contingent leg is therefore given by

Protection Leg PV =

$$N_{tranche} \int_0^T B(0, u) dP(u)$$

Where $B(0, u)$ is the Libor discount factor from today, time 0 to time u . We can discretise the integral by introducing the grid points $0 = t_0 < t_1 < \dots < t_K = T$ where greater accuracy is obtained by having a higher number of grid points. In practice we would generally assume that monthly intervals would be sufficient. There are a number of ways of evaluating an integral on this interval, the simplest being first order differences:

Protection Leg PV =

$$N_{tranche} \sum_{i=1}^K B(0, t_i) (P(t_i) - P(t_{i-1}))$$

To value the premium leg, we denote the contractual spread on the tranche by s , and the coupon payment dates as $0 = T_0 < T_1 < \dots < T_k = T$. The accrual factor for the i th payment is denoted by Δ_i . The PV of the premium leg is then given by

Premium Leg PV =

$$s N_{tranche} \sum_{j=1}^n \Delta_j (1 - P(T_j)) B(0, T_j)$$

The PV of the tranche from the perspective of the investor who is receiving the spread is therefore given by

Tranche PV = Premium Leg PV – Protection Leg PV.

As a result we can value a standard syn-

thetic CDO if we have the loss distribution at the set of future dates t and T . One way to do this is to use the one-factor model described above to calculate a loss distribution at each of the future dates required. Care must be taken to ensure that the default threshold is re-calibrated at each horizon date such that the marginal distribution is correctly recovered.

This approach can be more efficient than Monte Carlo since it uses the loss distribution directly and there are fast ways of calculating this. To emphasise this point, one can actually exploit the large number of assets in a synthetic CDO to derive a closed-form analytical solution to calculating the loss distribution of a portfolio.

An asymptotic approximation

It is possible to exploit the high dimensionality of the CDO to derive a closed form model for analysing CDO tranches. In fact we can use this approach to obtain results which differ from the exact approach by 1–5% for a real CDO.

To begin with, let us assume that the portfolio is homogeneous, ie, that each asset's β and default probability are the same. Hence, all assets have the same default threshold C . As a result, the conditional default probability of any individual issuer in the reference portfolio is given by

$$P(Z_M) = N \left(\frac{C - \beta Z_M}{\sqrt{1 - \beta^2}} \right).$$

If we also assume that the loss exposure to each issuer is of the same notional amount u , the probability that the percentage loss L of the portfolio is ku is equal to the probability that exactly k of the m issuers default,

which is given by the simple binomial

$$P[L = ku | Z] = \binom{m}{k} N\left(\frac{C - \beta Z_M}{\sqrt{1 - \beta^2}}\right)^k \left(1 - N\left(\frac{C - \beta Z_M}{\sqrt{1 - \beta^2}}\right)\right)^{m-k}$$

This distribution becomes computationally intensive for large values of m , but we can use methods of varying sophistication to approximate it. One very simple and surprisingly accurate method is the so-called 'large homogeneous portfolio' (LHP) approximation, originally due to Vasicek (1987). Since the asset returns, conditional on Z_M are independent and identically distributed, by the law of large numbers, the fraction of issuers defaulting will tend to the conditional probability of an asset defaulting $p(Z_M)$. As a result, the conditional loss is directly linked to the value of the market factor Z_M which itself is normally distributed. We can then write the probability of the portfolio loss being less than or equal to some loss threshold K as

$$P[L \leq K] = N[-p^{-1}(K)]$$

where $N(x)$ is the cumulative normal function. More involved calculations show that the distribution of L actually converges to this limit as m tends to infinity.

We can use the portfolio distribution to derive the loss distributions of individual tranches. If $L(K_1, K_2)$ is the percentage loss of the mezzanine tranche with attachment point K_1 and upper loss threshold K_2 , then this can be written as a function of the loss on the reference portfolio. As a fraction of

the tranche notional this is given by

$$L(K_1, K_2) = \frac{\max(L - K_1, 0) - \max(L - K_2, 0)}{K_2 - K_1}$$

If $K < 1$, then

$$L(K_1, K_2) \leq K \Leftrightarrow L \leq K_1 + K(K_2 - K_1)$$

This equation shows that we can easily derive the loss distribution of the tranche from that of the reference portfolio. This is discontinuous at the edges owing to the probability of the portfolio loss falling outside the interval, and this discontinuity becomes more pronounced when the tranche is narrowed. We can further compute the expected loss of the tranche analytically. For more details see O'Kane and Schloegl (2001). We then arrive at

$$E[L(K_1, K_2)] = \frac{N_2(-N^{-1}(K_1), C, -\sqrt{1-\beta^2}) - N_2(-N^{-1}(K_2), C, -\sqrt{1-\beta^2})}{K_2 - K_1}$$

This is only a one-period approach. However, it can easily be extended to multiple periods as described earlier.

Correlation skew

It is now possible to observe tradable tranche spreads for different levels of seniority. If we attempt to imply out the market correlation using a simple Gaussian copula model fitted to observed market tranche spreads, we can observe a skew in the average correlation as a function of the width and attachment point of the tranche. This skew may be the consequence of a number of factors such as the assumption of independence of default and recovery rates. It may also be due to our incorrect specification of the dependence structure, as

explained later. Supply and demand imbalances also play a role.

Rating agency models for CDOs

Different rating agencies have their own models for rating CDO tranches. All of them attempt to capture the risks of CDOs in terms of asset quality, recovery rates, default correlation and structural features.

Moody's standard rating model for CDO tranches is a multinomial extension of the Binomial Expansion Technique (BET) model. To capture default correlation, the portfolio is represented by a lower number of independent assets, known as the diversity score. Roughly speaking, this is the number of independent assets which have the same width of loss distribution as the actual CDO reference portfolio of correlated assets. The diversity score is determined using a lookup table and is based on the incremental effect of having groups of assets in the same industry classification.

After calibrating to historical default data, the model is able to generate an expected loss for a tranche which takes into account the subordination. This expected loss is then mapped to a rating category.

In S&P's ratings methodology, a Monte Carlo simulation is used to derive a probability loss distribution for the underlying collateral pool based on the total principal balance of the portfolio. Each corporate asset is assigned a default probability based on S&P's historical default studies, dependent on its rating and maturity. Corporate sectors are assumed to have a correlation of 30% within a given industry and 0% between industry sectors. From these inputs and the par amounts of each asset, a default probability distribution is created. Unlike Moody's which rates on the basis of expected loss, S&P rates on the basis of the probability of incurring a loss.

Estimating the dependency structure

Several well-known multivariate models, including the ones described earlier in this chapter, rely on the assumption that the dependence structure (or 'copula') of asset returns is Normal. The widespread use of the Normal dependence structure, which is fully characterised by a correlation matrix, is certainly related to its simplicity. It remains to be seen, however, whether this assumption is supported by empirical evidence. A number of recent studies have shown that the joint behaviour of equity returns is better described by a 'fat-tailed' Student-*t* copula than by a Normal copula, and that correlations are therefore not sufficient to appropriately characterise their dependence structure.² The first goal of this section is to apply the same kind of analysis to asset returns, and test the null hypothesis of Gaussian dependence versus the alternative of 'joint fat tails'.

Of course, we face a major obstacle when attempting to estimate the dependence structure of asset returns: asset values are not directly observable. In fact, the use of unobservable underlying processes is one of several criticisms that the structural approach has received over the years. Given the lack of observable asset returns, it has become customary to proxy the asset dependence with equity dependence, and to estimate the parameters governing the joint behaviour of asset returns from equity return series.³ However, the use of equity returns to infer the joint behaviour of asset

² See, for example, Mashal and Naldi (2002a) and Mashal and Zeevi (2002).

³ Fitch Ratings (2003) have recently published a special report describing their methodology for constructing portfolio loss distributions: it is based on a Gaussian copula parameterised by equity correlations.

returns is often criticised on the grounds of the different leverage of assets and equity. The second goal of this section is to shed some light on the magnitude of the error induced by using equity data as a proxy for asset returns.


To provide a plausible answer to these questions, we first need to 'back out' asset values from observable data. One way to estimate the market value of a company's assets is to implement a univariate structural model. Such a procedure is at the heart of KMV's CreditEdge™, a popular credit tool that first computes a measure of distance-to-default and then maps it into a default probability (EDF™) by means of a historical analysis of default frequencies.⁴ In what follows, we summarise recent work by Mashal, Naldi and Zeevi (2002), who use the asset value series generated by KMV's model to study the dependence properties of asset returns.

Methodology

A key observation in modelling and testing dependencies is that any d -dimensional multivariate distribution can be specified via a set of d marginal distributions that are 'knitted' together using a copula function. Alternatively, a copula function can be viewed as 'distilling' the dependencies that a multivariate distribution attempts to capture, by factoring out the effect of the marginals. Copulas have many important characteristics that make them a central concept in the study of joint dependencies, see, eg, the recent survey paper by Embrechts et al. (2003).

A particular copula that plays a crucial role

Figure 26. DJIA portfolio: asset and equity returns test statistics as functions of null hypothesis for DoF



in our study is given by the dependence structure underlying the multivariate Student- t distribution. The Gaussian distribution lies at the heart of most financial models and builds on the concept of correlation; the Student- t retains the notion of correlation but adds an extra parameter into the mix, namely, the degrees-of-free-

⁴ Copyright © 2000-2002 KMV LLC. All rights reserved. KMV and the KMV logo are registered trademarks of KMV LLC. CreditEdge and EDF are trademarks of KMV LLC.

dom (DoF). The latter plays a crucial role in modelling and explaining extreme co-movements in the underlyings.

Moreover, it is well known that the Student- t distribution is very 'close' to the Gaussian when the DoF is sufficiently large (say, greater than 30); thus, the Gaussian model is nested within the t -family. The same statement holds for the underlying dependence structures, and the DoF parameter effectively serves to distinguish the two models. This suggests how empirical studies might test whether the ubiquitous Gaussian hypothesis is valid or not. In particular, these studies would target the dependence structure rather than the distributions themselves, thus eliminating the effect of marginal returns that would 'contaminate' the estimation problem in the latter case. To summarise, the t -dependence structure constitutes an important and quite plausible generalisation of the Gaussian modelling paradigm, which is our main motivation for focusing on it.

With this in mind, the key question that we now face is how to estimate the parameters of the dependence structure. Mashal, Naldi and Zeevi (2002) describe a methodology which can be used to estimate the parameters of a t -copula without imposing any parametric restriction on the marginal distributions of returns. They also construct a likelihood ratio statistic to test the hypothesis of Gaussian dependence, and compare the dependence structures of asset and equity returns to evaluate the common practice of proxying the former with the latter. In the remainder of this section we summarise their empirical findings.

Empirical results

For the purpose of this study, asset and equity values are both obtained from KMV's

database. The reader should keep in mind, however, that equity values are observable, while asset values have been 'backed out' by means of KMV's implementation of a univariate Merton model. We apply our analysis to a portfolio of 30 credits included in the Dow Jones Industrial Average and use daily data covering the period from 31/12/00 to 8/11/02. The reader is referred to Mashal, Naldi and Zeevi (2003) for more examples using high yield credits and different sampling frequencies.

Following the methodology mentioned above, we estimate the number of degrees-of-freedom (DoF) of a t -copula without imposing any structure on the marginal distributions of returns. Then, using a likelihood ratio test statistic, we perform a sensitivity analysis for various null hypotheses of the underlying tail dependence, as captured by the DoF parameter. The two horizontal lines in Figure 26 represent significance levels of 99% and 99.99%; a value of the test statistic falling below these lines corresponds to a value of DoF that is not rejected at the respective significance levels.

The minimal value of the test statistic is achieved at 12 DoF for asset returns and at 13 DoF for equity returns. In both cases, we can reject any value of the DoF parameter outside the range [10,16] with 99% confidence; in particular, the null assumption of a Gaussian copula (DoF= ∞) can be rejected with an infinitesimal probability of error. Also, the point estimates of the asset returns' DoF lies within the non-rejected interval for the equity returns' DoF, and vice versa, indicating that the two are essentially indistinguishable from a statistical significance viewpoint. Moreover, the difference between the joint tail behaviour of a 12- and a 13-DoF t -copula is negligible in terms of any practical application.

Figure 27. Maximum likelihood estimates of DoF for DJIA portfolios

Portfolio	Asset returns	Equity returns
	DoF	DoF
30-credit DJIA	12	13
First 10 credits	8	9
Middle 10 credits	10	10
Last 10 credits	9	9

Figure 27 reports the point estimates of the DoF for asset and equity returns in the DJIA basket, as well as for three subsets consisting of the first, middle, and last 10 credits (in alphabetical order). The similarities between the joint tail dependence (as measured by the DoF) of asset and equity returns are quite striking.⁵

Next, we compare the remaining parameters that define a t -copula, ie, the correlation coefficients. Using a robust estimator based on Kendall's rank statistic⁶, we compute the two 30x30 correlation matrices from asset and equity returns. The maximum absolute difference (element-by-element) is 4.6%, and the mean absolute difference is 1.1%, providing further evidence of the similarity of the two dependence structures.

Example: synthetic loss tranche

The models described earlier in this chapter can be modified to account for a fat-tailed dependence structure of asset returns. Here we analyse the impact of a non-Normal assumption on the expected discounted loss (EDL) of a portfolio loss tranche. We focus on

EDL because this measure relates both to the agency rating (when computed under real-world probabilities) and to the fair compensation for the credit exposure (when computed under risk-neutral probabilities).

We consider a five-year deal with a reference portfolio of 100 credits, each with \$1m notional. We assume i) uniform recovery rates of 35%, for every credit in the reference portfolio; ii) 1% yearly hazard rate for each reference credit; iii) 20% asset correlation between every pair of credits; iv) a risk-free curve flat at 2%. Using a default-time simulation, Figure 28 compares the expected discounted losses for several tranches under the two alternative assumptions of Gaussian dependence and t dependence with 12 DoF. The results show the significant impact that the (empirically motivated) consideration of tail dependence has on the distribution of losses across the capital structure: expected losses are clearly redistributed from the junior to the senior tranches, as a consequence of the increased volatility of the overall portfolio loss distribution. Notice that even larger differences can be observed if one compares higher moments or tail measures of the tranches' loss distributions.

The LHP with tail dependence

It is possible to incorporate a Student- t copula into the LHP model discussed above. To do so, we must change the distribution for the asset returns to be multivariate Student- t distributed where we denote the DoF parameter with v and retain the one-factor correlation structure. This gives

$$A_i = \frac{\beta_i Z + \sqrt{1 - \beta_i^2} \varepsilon_i}{\sqrt{W/v}}$$

where W is an independent random variable

⁵ The range of accepted DoF is very narrow in each case, exhibiting similar behaviour to that displayed in Figure 26.

⁶ See Lindskog (2000).

following a chi-square distribution with v degrees of freedom. This is the simplest possible way to introduce tail dependence via a Student- t copula function.

Note that this is no longer a 'factor model' in the sense that the asset return is composed of two independent and random factors. The fact that both the market and idiosyncratic terms 'see' the same value of W means that they are no longer independent.

For this model, we have been able to calculate a closed-form solution for the density of the portfolio loss distribution, and show in O'Kane and Schloegl (2003) that it possesses the same tail dependent properties as described above, ie, a widening of senior spreads and a reduction in equity tranche spreads.

Summary

Our empirical investigation of the dependence structure of asset returns sheds some light on the two main issues that were raised at the beginning of this section. First, the assumption of Gaussian dependence between asset returns can be rejected with extremely high confidence in favour of an alternative 'fat-tailed dependence.' Multivariate structural models that rely on the normality of asset returns will generally underestimate default correlations, and thus undervalue junior tranches and overvalue senior tranches of multi-name credit products. A fat-tailed dependence of asset returns will produce more accurate joint default scenarios and more accurate valuations. Second, the dependence structures of asset and equity returns appear to be strikingly similar. The KMV algorithm that produces the asset values used in our analysis is nothing other than a sophisticated way of de-leveraging the equity to get to the value of a company's assets. Therefore, the

Figure 28. Expected discounted loss, 100K paths, standard errors in parenthesis

Tranche (%)	Normal copula EDL (std err %)	t copula DoF=12 EDL (std err %)	Pctg diff
0-5	\$2,256,300 (0.14)	\$2,012,200 (0.23)	-11
5-10	\$533,020 (0.63)	\$601,630 (0.66)	13
10-15	\$146,160 (1.37)	\$221,120 (1.06)	51
15-20	\$41,645 (1.70)	\$90,231 (1.62)	117

popular conjecture that the different leverage of assets and equity will necessarily create significant differences in their joint dynamics seems to be empirically unfounded. From a practical point of view, these results represent good news for practitioners who only have access to equity data for the estimation of the dependence parameters of their models.

Modelling credit options

We separate credit options into options on bonds and options on default swaps.

Pricing options on bonds

Options on corporate bonds are naturally divided into three groups according to how the exercise price is specified. The option can be struck on price, yield or credit spread. The exercise price is constant for options struck on price, but for options struck on yield it depends on the time to maturity of the underlying bond and is found from a standard yield-to-maturity calculation. Obviously, for European options there is no difference between specifying a strike price and a strike yield.

Bond options struck on spread are different. For credit spread options the exercise price depends both on the time to maturity


of the bond and on the term structure of interest rates at the exercise time. A credit spread strike is commonly specified as a yield spread to a Treasury bond or interest rate swap, or as an asset swap spread.

For short-dated European options with price (or yield) strike on long-term bonds, the well-known Black-Scholes formula goes a long way but is not recommended beyond this limited universe. An important problem with Black-Scholes is that it does not properly account for the pull-to-par of the bond price. This problem can be solved by a yield diffusion model where the bond's yield is the underlying stochastic process. It is relatively easy to build a lattice for a lognormal yield, say, and price the option by standard backwards induction techniques. The lattice approach also allows for easy valuation of American/Bermudan exercise. Yield diffusion models can be constructed to fit forward bond prices but would usually assume constant yield volatility which is inconsistent with empirical evidence.

The problems of pull-to-par of price and non-constant volatility can be solved by a full term structure model such as Black-Derman-Toy, Black-Karasinski or Heath-Jarrow-Morton. These models were designed for default-free interest rates but can be applied analogously to credit risky issuers. Instead of calibrating the model to Libor rates, the model should be calibrated to an issuer-specific credit curve.

It is rarely possible to calibrate volatility parameters because of the lack of liquid bond option prices. Usually it is more appropriate to base volatility parameters on historical estimates. Naldi, Chu and Wang (2002) and Berd and Naldi (2002) present a multi-factor framework for modelling the stochastic components of corporate bond returns.

Figure 29. Receiver default swaption



The framework can be used to derive price/yield volatilities from return volatilities and tend to give more robust estimates than direct estimation.

Extending interest rate models to modelling risky rates only works if we assume either zero recovery on default, or assume that recovery is paid as a fraction of the market value at default. However, creditors have a claim for return of full face value in bankruptcy, so it is necessary to specifically model the recovery at default as a fraction of face value. This is especially important for lower credit quality issuers where the bonds trade on price rather than yield.

For high quality issuers, the above approaches are less controversial for pricing options struck on price or yield. For these issuers, interest rate volatility is the main driver of price action. Volatility parameters should therefore be related to implied volatilities from interest rate swaption markets but must also incorporate the negative correlation between credit spreads and interest rates (see Berd and Rangelova (2003)) which can cause yield volatility on corporate bonds to be significantly lower than comparable interest rate swaption volatility.

When pricing spread options, it is important to specifically take into account the default risk. In the next section we discuss

how to price options to buy or sell protection through CDS. Bond options struck on credit spread can be priced in similar fashion.

Pricing default swaptions

The growing market in default swaptions has led to a demand for models to price these products. First, let us clarify terminology. An option to buy protection is a payer swaption and an option to sell protection is a receiver swaption. This terminology is analogous to that used for interest rate swaptions.

Black's formulas for interest rate swaptions can be modified to price European default swaptions. Consider a European payer swaption with option expiry date T and strike spread K . The contract is to enter into a long protection CDS from time T to T_M , and it knocks out if default occurs before T (see Figure 29). Conditional on surviving to time T , the option payoff is

$$PS_T = PV01_T \cdot \max\{S_T - K, 0\}$$

where the $PV01_T$ is the value at T of a risky 1bp annuity to time T_M or default, and S_T is the market spread observed at T on a CDS with maturity T_M .

Ignoring the maximum, this is the standard payoff calculation for a forward starting CDS. See page 32 for a discussion of CDS pricing.

From finance theory we know that for any given security, say B , that only makes a payment at T , there is a probability distribution of the spread S_T such that for any other security, say A , that also only makes a payment at T , the ratio A_0/B_0 of today's values of the securities is equal to the expectation of the ratio A_T/B_T of the security payments at T . The result is valid even if B_T can be 0 as long as $A_T = 0$ when $B_T = 0$. The states where B_T

$= 0$ are simply ignored in that case and the distribution of S_T will be such that the probability that $B_T = 0$ is 0. (See Harrison and Kreps (1979) for details.)

To use this result we first let A be a security that at T pays 0 if default has occurred and otherwise pays the upfront cost (as of T) of a zero-premium CDS with the same maturity as the CDS underlying the swaption. We let B be a security that at T pays 0 if default has occurred and otherwise pays $PV01_T$. With these definitions the ratio A_T/B_T is equal to the spread S_T if default has not occurred at T , otherwise A_T/B_T and S_T are not defined. The distribution of S_T should then be such that $E[S_T] = A_0/B_0$ where the probability of default before T is put to 0. A_0/B_0 is the T -forward spread, denoted F_0 , for the underlying CDS where the forward contract knocks out if default occurs before T .

The next step is to let A be the swaption, in which case A_T is 0 if default happens before T and PS_T otherwise. The value of the swaption today is then

$$PS_0 = B_0 E \left[\frac{PS_T}{B_T} \right] = PV01_0 \cdot E[\max\{S_T - K, 0\}]$$

If we make the assumption that $\log(S_T)$ is normally distributed with variance $\sigma^2 T$, corresponding to the spread following a log-normal process with constant volatility σ , then with the requirement $E(S_T) = F_0$ (the forward spread), we have determined the distribution of S_T to be used to find $E[\max\{S_T - K, 0\}]$. It is easy to calculate this expectation and we arrive at the Black formula

$$PS_0 = PV01_0 \cdot (F_0 \cdot N(d_1) - K \cdot N(d_2))$$

$$d_1 = \frac{\log(F_0 / K) + \sigma^2 T / 2}{\sigma \sqrt{T}} \quad d_2 = d_1 - \sigma \sqrt{T}$$

where N is the standard normal distribution function. The Black formula for a receiver swaption is found analogously. (See Hull and White (2003) and Schonbucher (2003) for more details.)

It is important that the forward spread, F_0 , and $PV01_0$ are values under the knockout assumption. Calculation of F_0 and $PV01_0$ should be done using a credit curve that has been calibrated to the current term structure of CDS spreads.

To value a payer swaption that does not knock out at default, we must add the value of credit protection from today until option maturity. Under a non-knockout payer swaption, a protection payment is not made until option maturity. The value of the credit protection is therefore less than the upfront cost of a zero-premium CDS that matures with the option. The knockout feature is not relevant for receiver swaptions, as they will never be exercised after default.

The last step in valuing the swaption is to find the volatility σ to be used in the Black formula. An estimate of σ can be made from a time series of CDS spreads. Examination of CDS spread time series also reveals that the log-normal spread assumption often is inappropriate. It is not uncommon for CDS spreads to make relatively large jumps as reaction to firm specific news. However calibrating a mixed jump and Brownian process to spread dynamics is not easy.

To price default swaptions with Bermudan exercise we could construct a lattice for the forward spread, but criticism analogous to that for using a yield diffusion model to price bond options applies. Instead, we recommend using a stochastic hazard rate model.

To illustrate the use of the Black formula, consider a payer swaption on Ford Motor Credit with option maturity 20/12/2003 and

swap maturity 20/9/2008. The valuation date is 28/8/2003. The strike is 260, which is the five year CDS spread on the valuation date. Our trading desk quoted the swaption at 2.29%/2.52%, which means that \$10m notional can be bought for \$252,000 and sold for \$229,000. These are prices of non-knockout options. The fair value of protection until swaption maturity given the credit curve on the valuation date is 0.211%, the $PV01$ used in the Black formula is 4.027 and the forward spread is 274.7bp. These numbers imply that the bid/offer volatilities quoted by our desk are approximately 75% and 85%.

Interest rates and credit risk

One way to model credit spread dynamics is via a hazard rate process. This provides a consistent framework for modelling spreads of many different maturities. The differences between directly modelling the credit spread and modelling the hazard rate are similar to the differences between modelling the yield of a bond and using a term structure model as discussed above.

A stochastic hazard rate model is naturally combined with a term structure model to produce a unified model that can, at least in theory, be used to price both bond options and credit spread options. However, a number of practical complications arise in getting such an approach to work. Since bond and the CDS markets have their own dynamics, usually not implying the same issuer curves, it may not be appropriate to naively price all options in the same calibrated model without properly adjusting for the basis between CDS and bonds.

Also, calibration of the volatility parameters of the hazard rate process, when possible, is less than straightforward, especially when calibrating to bond options struck on

price or yield, or to bonds with embedded options. Interest rate volatility parameters should be calibrated separately using liquid prices of interest rate swaptions. These parameters are then taken as given when determining the hazard rate volatility parameters by fitting to time series estimates or calibrating to market prices.

Finally, it is necessary to determine the correlation between interest rates and hazard rates. When estimating correlations it is important not to use yield spread or OAS (option adjusted spread) as a proxy for the hazard rate, because such spread measures do not properly take into account the risk of default. Using OAS is especially a problem for long maturity bonds with high default probabilities and high recovery rates. For such bonds there can be a significant decrease in OAS when interest rates increase even if hazard rates remain unchanged. Correlations should be estimated directly using bond prices or CDS spreads.

Modelling hybrids

The behaviour of hybrid credit derivatives is driven by the joint evolution of credit spreads and other market variables such as interest and exchange rates, which are commonly modelled as diffusion processes. As a consequence, the reduced form approach is the natural framework for pricing and hedging these products.

We illustrate the main ideas via the example of default protection on the MTM of an interest rate swap. Suppose an investor has entered into a receiver swap with fixed rate k with a credit risky counterparty. If the MTM of the receiver swap, RS_t , is positive to the investor at the time of default, this is paid by the protection seller. If RS_t is negative, the investor receives nothing, so that the payoff at default is $\max(RS_t, 0)$. This is

an option to enter into a receiver swap with fixed rate k for the remaining life of the original trade at default. For simplicity, we assume that default can only take place at times t_i . If B denotes the price process of the savings account, then computing the expected discounted cash flows gives the value V_0 for the price of the default protection, where

$$V_0 = \sum_{i=1}^n E \left[\frac{\max(RS_{t_i}, 0)}{B(t_i)} \middle| \tau = t_i \right] P[\tau = t_i]$$

The interpretation of this equation is that the value of default protection is a probability weighted strip of receiver swaptions, where each swaption is priced conditional on default happening at t_i .

Representing the default protection via a strip of swaptions is a very useful framework for developing intuition around the pricing and helps us understand the importance of the shape of the interest rate curve. In terms of volatility exposure, the protection seller has sold swaptions and is therefore short interest rate volatility.

If the rate and credit process are correlated, then there will also be a spread volatility dependence. However, under the assumption of independence, the volatility of the credit spread does not enter into the valuation, only the default probabilities.

A strip decomposition, as we have shown above, is the basic building block for most hybrid credit derivatives. A cross-currency swap for example could be dealt with in exactly the same way, with the exchange rate as an additional state variable. For this instrument, the exchange rate exposure is of prime importance.

In terms of tractability, it is important to be

able to calculate the conditional expectations appearing in the strip decomposition. A further consequence of the hazard rate method is that we do not actually have to condition on the realisation of the default time itself, conditioning instead on the realisation of the hazard rate process. This is particularly advantageous for Monte Carlo simulation, as one does not have to explicitly simulate default times and is a useful variance reduction technique.

The parameters needed for pricing hybrids are essentially volatility and dependence parameters. Calibrating volatilities for interest rates and FX is relatively straightforward. Credit spread volatilities are somewhat more involved. Until recently, we have had to rely on estimates of historical volatilities; now the growing options markets are starting to make the calibration of implied spread volatilities feasible.

Determining the correct dependence

structure between credit spreads and the other market variables is the main challenge in modelling hybrids. The simplest approach is to work within a diffusion setting where spreads and interest rates/FX are correlated. Such a model will not necessarily generate levels of dependence representative of periods of market stress where investment grade defaults are likely. As a starting point, however, it appears to be reasonable.

Even within this framework, the effect of correlation on the valuation of a hybrid instrument can be marked. At this stage of the market it is safe to say that this correlation is a 'realised' parameter as opposed to an implied one, ie pricing and hedging must proceed on the basis of a view on correlation founded on historical estimates. Going forward, it will be interesting to see to what extent a market in implied correlations will develop via standardised hybrid credit derivatives.

References

Altman E, Resti A and Sirone A, 2001

Analyzing and explaining default recovery rates

Report submitted to ISDA, Stern School of Business, New York University, December

Berd A and Naldi M, 2002

New estimation options for the Lehman Brothers risk model

Quantitative Credit Research Quarterly, Lehman Brothers, September

Berd A and Rangelova E, 2003

The co-movement of interest rates and spreads: implications for credit investors
High Grade Research, Lehman Brothers,

Embrechts P, Lindskog F and

McNeal A, 2003

Modelling dependence with copulas and applications to risk management

In Handbook of Heavy Tailed Distributions in Finance, Elsevier, New York, pages 329–384

Finger C, 1999

Conditional approaches for CreditMetrics portfolio distributions

RiskMetrics Group, New York

Finger C, 2000

A comparison of stochastic default rate models

RiskMetrics Group, New York

Fitch Ratings, 2003

Default correlation and its effect on portfolios of credit risk

Structured Finance, Credit Products Special Report, February

Frye J, 2003

A false sense of security

Risk August, pages 95–99

Ganapati S, Ha P and O’Kane D, 2001

Synthetic CDOs – how are they structured? how do they work?

CDO Monthly Update, Lehman Brothers, November

Ganapati S and Ha P, 2002

Evolution of synthetic arbitrage CDOs

Structured Credit Strategies, Lehman Brothers, November

Ganapati S, Berd A, Ha P and Rangelova E, 2003

The synthetic CDO bid and basis convergence – which sector is next?

Structured Credit Strategies, Lehman Brothers, March

Gregory J and Laurent J-P, 2003

I will survive

Risk June, pages 103–107

Harrison J and Kreps D, 1979

Martingales and arbitrage in multiperiod securities markets

Journal of Economic Theory 20, pages 381–408

Hull J, 2000

Options, futures and other derivatives

Fourth edition, Prentice Hall

Hull J and White A, 2001

Valuing credit default swaps II: modelling default correlations

Journal of Derivatives 8, pages 12–22

Hull J and White A, 2003

The valuation of credit default swap options

Working paper, University of Toronto

Isla L, 2003

Hedged CDO equity strategies

Fixed Income Research, Lehman Brothers, September

Jarrow R and Turnbull S, 1995

Pricing derivatives on financial securities subject to credit risk

Journal of Finance 50, pages 53–85

Li D, 2000

On default correlation: a copula function approach

Journal of Fixed Income 9, pages 43–54

Lindskog F, 2000

Linear correlation estimation

Working paper, RiskLab, ETH Zurich

Mashal R and Naldi M, 2002a

Extreme events and default baskets

Risk June, pages 119–122

Mashal R and Naldi M, 2002b

Beyond CADR: searching for value in the CDO market

Quantitative Credit Research Quarterly, Lehman Brothers, September 2002

Mashal R, Naldi M and Pedersen C, 2003

Leverage and correlation risk of synthetic loss tranches

Quantitative Credit Research Quarterly, Lehman Brothers, April

Mashal R, Naldi M and Zeevi A, 2002

The dependence structure of asset returns

Quantitative Credit Research Quarterly, Lehman Brothers, December

Mashal R, Naldi M and Zeevi A, 2003

Extreme events and multi-name credit derivatives

In Credit Derivatives The Definitive Guide, Risk, London

Mashal R and Zeevi A, 2002

Beyond correlation: extreme co-movements between financial assets

Working paper, Columbia University, New York

Munves D, Berd A, O'Kane D and Desclee R, 2003

The Lehman Brothers CDS indices

Fixed Income Research, Lehman Brothers, September

Naldi M, Chu K and Wang G, 2002

The new Lehman Brothers credit risk model
Quantitative Credit Research Quarterly, Lehman Brothers, May

O'Kane D and McAdie R, 2001

Explaining the basis: cash versus default swaps

Lehman Brothers, March

O'Kane D, Pedersen C and Turnbull S, 2003

Valuing the restructuring clause in CDS
Lehman Brothers, June

O'Kane D and Schloegl L, 2001

Modelling credit: theory and practice
Lehman Brothers, February

O'Kane D and Schloegl L, 2003

Leveraging the spread premium with correlation products

Quantitative Credit Research Quarterly, Lehman Brothers, July

O'Kane D and Sen S, 2003

Upfront credit default swaps

Quantitative Credit Research Quarterly, Lehman Brothers, July

O'Kane D and Turnbull S, 2003

Valuation of credit default swaps
Quantitative Credit Research Quarterly, Lehman Brothers, June

Schonbucher P, 2003

A note on survival measures and the pricing of options on credit default swaps
Working paper, ETH Zurich

Vasicek O, 1987

Probability of loss on loan portfolio
Working paper, KMV Corporation

Approved by Lehman Brothers International (Europe), which is authorised and regulated by the Financial Services Authority. Rules made for the protection of investors under the UK Financial Services Act may not apply to investment business conducted at or from an office outside the UK.

Published by Incisive Risk Waters Group,
Haymarket House, 28–29 Haymarket,
London SW1Y 4RX
© 2003 Lehman Brothers and Incisive RWG Ltd


All rights reserved. No part of this publication may be reproduced, stored in or introduced into any retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the copyright owners.

Risk Waters Group and Lehman Brothers and the author have made every effort to ensure the accuracy of the text; however, neither they nor any company associated with the publication can accept legal or financial responsibility for consequences, which may arise from errors, omissions, or any opinions given.

Notes

An Established Tradition of Research Excellence

Weekly commentary across asset classes



Award-winning website www.lehmanlive.com



In-depth product analysis



The Lehman Brothers Global Fixed Income Research Group has been consistently recognised as a top-ranked research department.

- For the fourth year in a row, the Group was ranked #1 in the Institutional Investor 2003 All-America Fixed Income Research Team Survey
- Institutional Investor ranked Lehman Brothers #1 for Fixed Income Research in 10 out of the last 14 years
- Lehman Brothers website, LehmanLive - winner of 8 "Best Site" categories in the Euromoney.com Awards
- LehmanLive - #1 dealer website on the Bond.Hub portal
- Dominant global marketshare in Fixed Income Indices and Index Analytics - available on our POINT™ platform

We have an established tradition of providing the highest quality research services to our clients. With over 375 dedicated professionals worldwide covering the whole spectrum of Fixed Income products, we provide strategic advice and analysis to our institutional clients.

All of our research and analysis is accessible through our website, LehmanLive.

For further information please contact the Global Fixed Income Research Group or visit www.lehmanlive.com

New York: **Ravi Mattu**
+1 212 526 0206 rmattu@lehman.com

Europe: **Robert Campbell**
+44-(0)20-7260-2639 rcampbel@lehman.com

Prafulla Nabar
+44-(0)20-7256-4480 pnabar@lehman.com

Asia: **Hua He**
+81-3-5571-7215 hhe@lehman.com

LEHMAN BROTHERS
Where vision gets built.®

All Rights Reserved. Member SIPC. Lehman Brothers International (Europe) is regulated by the Financial Services Authority. ©2003 Lehman Brothers Inc.

