Mathematics, Finance and Risk

Synthetic CDOs

Modelling, Valuation and Risk Management

C. C. Mounfield

CAMBRIDGE

This page intentionally left blank

SYNTHETIC CDOs

Modelling, Valuation and Risk Management

Credit derivatives have enjoyed explosive growth in the last decade. One of the most important assets in this industry is synthetic Collateralised Debt Obligations (synthetic CDOs). This book describes the state-of-the-art in quantitative and computational modelling of these instruments.

Starting with a brief overview of the structured finance landscape, the book introduces the basic modelling concepts necessary to model and value simple vanilla credit derivatives. Building on this the book then describes in detail the modelling, valuation and risk management of synthetic CDOs. A clear and detailed picture of the behaviour of these complex instruments is built up. The final chapters introduce more advanced topics such as portfolio management of synthetic CDOs and hedging techniques, often not covered in other texts.

Mathematics, Finance and Risk

SYNTHETIC CDOs

Modelling, Valuation and Risk Management

CRAIG MOUNFIELD

CAMBRIDGE
UNIVERSITY PRESS

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

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

www.cambridge.org
Information on this title: www.cambridge.org/9780521897884

First published in print format 2008

ISBN-13 978-0-511-46084-5 eBook (Adobe Reader)

ISBN-13 978-0-521-89788-4 hardback

Dedicated to my parents, my wife and my daughter.

Contents

Preface

This is a book about the modelling, valuation and risk management of synthetic collateralised debt obligations (or synthetic CDOs or simply CDOs for short). Synthetic CDOs are an example of a structured credit product. This is a financial product that takes targeted risk for the purpose of achieving targeted returns. Structured credit products utilise two financial engineering technologies: credit derivatives and asset securitisation. Synthetic CDOs have played an increasingly important role in the expansion of the global credit derivatives market which has grown rapidly since the turn of the century. Indeed, it is estimated that by the end of 2006 the total credit derivative notional amount outstanding was over $20 trillion (from virtually zero only a decade earlier). Increased trading volumes naturally led to market participants becoming more sophisticated (in terms of their risk/return characteristics and the strategies they employ) as well as to a commensurate increase in the complexity and subtlety of the products available. This in turn drives the evolution of the mathematical and computational models used to value these products. The objective of this book is to collate, summarise and critically assess the current state-of-the-art in quantitative and computational modelling of synthetic CDOs. The key word here is *modelling*; the book is about mathematical models and their properties. This book is not intended to provide detailed descriptions of the business and economic rationales for trading credit derivatives; there are better resources available that describe this and due reference will be given to these sources. It is meant to provide a detailed quantitative description of the modelling techniques currently employed in the marketplace for characterising synthetic CDOs.

It will be assumed that the technical level and experience of the reader is relatively high. Basic financial concepts will not be described in detail (except insofar as when such detail is necessary). Instead reference will be made to the appropriate resources. The use of financial and technical jargon will hopefully be kept to a minimum, although in a specialised, technical text such as this some jargon is inevitable. The rationale for this approach is to ensure the volume is concise and to the point. It is intended to describe just enough of the mathematical and computational modelling to enable the reader to understand the relevant issues (along with a discussion of the practical implementation considerations) and help the reader to form their own opinion as to the merits, or otherwise, of the models presented. I will consider the book to be a success if it enables readers to understand the behaviour of models and to build better versions of them. This lean approach will hopefully make the volume attractive to practitioners (who do not always have the time to study a subject in detail) who wish to understand more about the properties of the credit derivative models commonly used in the marketplace. In particular it is envisaged that the volume will be of interest to a range of different types of practitioner.

- Quantitative analysts (quants) and quant developers wanting to understand more about credit modelling and credit derivatives. The book is written with a strong emphasis on models, implementation and understanding of the model behaviour. It is therefore well suited to quants in model validation teams, for example.
- Quantitative risk managers wanting to understand better the models used for valuation, to interpret synthetic CDO risk sensitivities (e.g. spread and correlation sensitivities) and risk manage complex credit portfolios.
- Traders and product controllers seeking a better understanding of the mechanics going on in the black boxes when ‘F9’ is pressed (and to understand the relative strengths and weaknesses of different models).
- Structurers wanting to understand better the properties of the instruments they are using to construct strategies with specific risk/return characteristics.
- Researchers in academia looking to understand some of the practical issues surrounding the common models used in the marketplace.

The downside to this lean approach is that for less experienced readers the material may at times not give as much explanation as would be liked, or some (basic) concepts are not described fully. However, for the motivated and intelligent reader this should present not a problem but a challenge and (as the author knows from experience) the rewards in terms of deeper understanding are worth the effort.

At the beginning of a project such as writing a book one has a vision as to what the finished product will look like. The vision for this book was that it would be very much model focused, with a strong emphasis on the practical, pragmatic implementation details that are of crucial importance in a live banking environment. This means there is less focus on the ‘business’ topics of the economics, mechanics and structures of credit derivatives than can be found in other texts. To include this information would have detracted from the core message of models and their properties. Also, when writing a book it is necessary to make compromises and be pragmatic in terms of content. At the beginning of the project one’s vision of what will be achieved is vast and expansive. By the end of the project one is simply happy to stumble across the finish line. There are occasions throughout the book when more detailed analysis of a particular model or scenario would have been very useful indeed to illustrate a particular point further, but due to time constraints was not included. On these occasions it is suggested that the reader build the models and do the analysis themselves as an exercise.

This leads into the next important point about the approach taken in the text. In the modern world of quantitative finance it is almost impossible to develop models of complex derivative trades that are wholly tractable analytically. It is therefore difficult to separate a model's mathematical description from its actual implementation. When it comes to building models suitable for use within a live investment banking environment the devil really is in the details. Full understanding of a model only comes from implementing it, analysing its properties and understanding its weaknesses. An important objective of this volume, therefore, is to provide not only the mathematical descriptions of the models, but also details of the practical implementation issues. To achieve this objective, liberal use is made of pseudo code to illustrate the implementation of an algorithm. The purpose of this code is to allow the reader to convert quickly a description of a model into the programming environment of their choice (although the author is most familiar with C++, and there may appear to be a bias towards the syntax of this language on occasion).

The volume is structured into three distinct sections. Broadly speaking Chapters 1–3 motivate the main topic, synthetic CDOs, and introduce some of the basic modelling tools necessary to describe them. Chapters 4–10 analyse the mathematical and computational modelling techniques applied to synthetic CDOs. Chapters 11–14 look at more advanced topics in the analysis of synthetic CDOs. Each of the chapters can in principle be read in isolation and each is relatively self-contained. However, there is a clear path from chapter to chapter (which reflects the author's own train of thought), particularly in Chapters 4–10. Reading each chapter sequentially will build a clearer and more coherent picture of the subject matter as a whole, but it is by no means a prerequisite.

In the first part of the book we motivate the study of synthetic CDOs by understanding their importance and usage within the broader credit derivatives marketplace. Chapter 1 provides a brief overview of the credit derivatives market in terms of instruments and introduces the crucial concepts of securitisation and tranching which are the basis of CDO technology. In this first section we also provide some of the basic mathematical building blocks necessary for later chapters. Chapter 2 describes the current market standard modelling methodologies for capturing the arrival of default risk of an obligor. This chapter also introduces the concepts and methods used for the modelling of default correlation, which as we will see is one of the most fundamental concepts in the characterisation of synthetic CDOs (and indeed any multi-name credit derivative). The first section of the book ends with a discussion, in Chapter 3, of the valuation models for the simplest and most vanilla of credit derivatives – credit default swaps or CDSs. The market for single-name default protection CDSs is extremely liquid and a good understanding of the valuation methods for these basic building blocks is a necessary prerequisite for understanding the more complex multi-name products.[1] For a reader already conversant with single-name credit derivatives, the material in Chapters 1–3 will be familiar. Indeed these chapters are only included in order to provide a reference guide to the concepts underpinning the rest of the book.

The second part of the volume, Chapters 4–10, which is its mathematical and computational core, focuses specifically on the valuation and risk analysis of multi-name credit derivatives and synthetic CDOs in particular. Chapter 4 introduces the credit indices that have emerged and evolved over the course of the last few years. The introduction and subsequent trading of these indices has provided enormous impetus to the growth of the credit derivatives market. Chapter 5 then introduces default baskets. In terms of materiality, default baskets are a very small fraction of the overall structured credit marketplace. However, they are the simplest form of multi-name credit derivative and an understanding of their valuation and risk sensitivities can provide substantial insight into the behaviour of more complex synthetic CDOs.

Chapters 6 through 8 develop and analyse the core mathematical models for valuing synthetic CDOs. Chapter 6 describes a number of different methodologies for valuation and, in particular, introduces the current market standard valuation model, the so-called normal copula model. Chapter 7 investigates the fundamental behaviour of the model as certain key parameters are varied systematically. As will be seen in this chapter, the phenomenology of the model is relatively complex and subtle. Chapter 8 analyses the risk sensitivities of the standard market model to variations of input parameters. More importantly this chapter discusses the different risk sensitivity measures such as credit spread 01 (CS01) and value-on-default (VoD) that are necessary to capture and characterise the risk inherent in synthetic CDOs.

The next chapters look at the implications for the standard market model that standardised tranches and the development of a liquid market have had. Initially the market for synthetic CDOs was relatively illiquid and deals were done on a bespoke basis. The introduction of standardised credit indices and the subsequent development of a market for trading tranched exposures to slices of the index provided enormous impetus to the liquidity and volume of trades in single-tranche synthetic CDOs (STCDOs). Eventually the market became sufficiently liquid to allow transparent price discovery for the prices of these standardised index tranches. At this

[1] The main focus of the book is synthetic CDOs. Therefore we will not spend a great deal of time talking about CDSs and other credit derivatives – there are better texts available that describe these products in great detail.

point the role of the standard model changed; it became a mechanism whereby market participants could express and trade their views on default correlation. Chapter 9 introduces the concepts of implied and base correlations that have been developed to capture implied pricing information from market observed prices. As the prices of instruments become transparent in the open market it is crucially important for the standard model to be able to reproduce these prices accurately. Chapter 10 describes some of the different methodologies that have been developed to allow calibration of models of synthetic CDOs to market observed prices (the so-called ‘correlation skew’).

The final part of the volume, Chapters 11–14, looks at more advanced topics in the characterisation and analysis of synthetic CDOs. Chapter 11 introduces a number of exotic CDOs. Examples include CDOs with asset backed securities as the underlying pool of obligors as well as CDOs with CDOs as the assets in the underlying pool (so called CDO squareds). Correlation trading is the term used to refer to trading strategies designed to exploit the risk/return characteristics of portfolios of CDO tranches. Chapter 12 analyses the risk/return characteristics of a number of popular CDO trading strategies. Chapter 13 considers extending the models developed thus far for a single-tranche position to a portfolio of tranches and assesses how the risk in the tranche portfolio can be quantified and controlled.

Finally, a natural extension of analysing the static (in time) performance of CDO trading and hedging strategies is to look at the through life performance of the trading strategy. In the pricing of simpler derivatives, the value of the derivative is equal to the cost of the dynamic hedging strategy. If a hedging strategy is good at capturing all the risks a position is exposed to then the overall P/L generated from the process of selling the derivative instrument and rebalancing the hedging portfolio as the market risk factors evolve should be small. If the hedging strategy is not adequate there will be significant P/L leakage. Chapter 14 sets up and analyses a simple hedging simulation of synthetic CDO tranches. This chapter is more speculative in nature than previous chapters as it represents the cutting edge of technology applied to the analysis of complex derivative securities.

Acknowledgements

A book is never written in isolation, and it is a pleasure to acknowledge the contribution that a number of individuals have made to the current text. I would like to thank all the people I have worked with in the Model Validation and Risk Management teams of Credit Suisse and Barclays Capital as well as my co-workers at Cheyne Capital Management. A lot of the experience that is encapsulated in this text is a direct result of day-to-day interactions with my colleagues at these institutions. In particular, I would like to thank Dr Niclas Sandstrom of Barclays Capital and Dr Andrea Petrelli of Credit Suisse for their detailed reading of the original manuscript, and for making numerous suggestions as to how it could be improved.

I would also like to thank my editor at CUP, David Tranah (and all the other staff who have contributed to the bringing to fruition of this project), for providing me with an opportunity to write this book. Finally I would like to acknowledge the contribution of my Ph.D. supervisor Professor Sir S. F. Edwards of the Cavendish Laboratory, Cambridge. The scientific training I received under his tutelage has proven to be of enduring value throughout my career. I hope this text reflects some of what I learnt from him.

1

A primer on collateralised debt obligations

Credit – Derived from the Latin verb credo meaning 'I trust' or 'I believe'.

1.1 Introduction

In this book we will introduce and describe in detail synthetic collateralised debt obligations (or synthetic CDOs for short). Synthetic CDOs are a sophisticated example of a more general asset class known as credit derivatives. In their simplest form credit derivatives facilitate the transfer of credit risk (the risk that a counterparty may fail to honour their outstanding debt obligations such as paying coupons or repaying principal on bonds they issued) between different counterparties to a trade. The rationale for trading credit derivatives is to allow this risk to be transferred efficiently between counterparties, from those who are unwilling or unable to hold it, to those who want it. This chapter will introduce some of the important credit derivative products that will be analysed in detail later in the book. The chapter will also introduce the financial engineering concepts that underlie synthetic CDOs.

Section 1.2 introduces the concepts of securitisation and tranching. These are the key financial innovations that underpin CDOs and indeed much of structured finance technology. Section 1.3 then provides an overview of some of the most common credit derivative instruments. These include credit default swaps, credit indices and most importantly synthetic CDOs. The key features of the different instruments will be described and some discussion given of the motivations for trading them (although the level of detail of this final point is by no means exhaustive since there are other resources available which already extensively cover this material [Das 2005, Kothari 2006, Rajan *et al*. 2007]). Finally in Section 1.4 we briefly summarise the key points introduced in the chapter and set the scene for the remainder of the book.

1.2 Securitisation and tranching

In this section we provide an overview of the concepts of securitisation and tranching (a very detailed reference on this topic is Kothari [2006]). These are the fundamental financial engineering techniques that underpin CDOs and indeed most of structured finance. We motivate the discussion of securitisation by considering a simplified model of a bank's business.

The business of banks is to invest money and speculate with the expectation of making a positive return on their investments. They will, for example, provide loans or extend lines of credit to corporate entities for them to invest in expanding their business. In return for these loans the corporate pays interest to the bank and at the maturity of the loan repays the initial principal back (or alternatively the principal is paid back gradually over time). The risk the bank runs is that, for one reason or another, they will not get back the income due to them from the periodic coupons or their original investment (return of principal). For example, if the corporate were to go into bankruptcy or administration due to poor management or a global recession, it is unlikely the bank would receive all of their investment back.

The key component in the whole of the global financial system is liquidity (as was painfully apparent during the latter half of 2007 – a good history of financial crises past can be found in Kindleberger and Aliber [2005] and Bookstaber [2007]). Banks need cash in order to survive day-to-day. If all of the loans that a bank has made were to go bad simultaneously, the income the bank receives from this business would evaporate, forcing them to raise their costs of borrowing in other areas to recoup some of the lost income (in turn putting pressure on other parts of the economy such as consumer spending). Or worse, the bank could go out of business. In order to mitigate against the risk of loans going bad, banks are required by their regulatory bodies to hold capital against their investments. For example, if it was assumed that loans on average default at a rate of 5% per year the bank may be required to hold in readily available assets (not illiquid securities such as retail mortgages) a total of 8% of the value of their book. To a bank seeking the maximum possible return on their capital to keep shareholders happy this regulatory capital is dead money. Any means for reducing this amount is most welcome.

Unfortunately investments such as loans to corporate entities, mortgages to individuals, automobile loans, credit card receivables, home equity loans etc. are very illiquid assets. There is no secondary market for actively trading individual loans in the same way that there is for trading, for example, shares in IBM. It is difficult therefore for the bank to do anything with these assets. This is where the concept of securitisation enters. The basic concept of securitisation is to bundle up large numbers of the illiquid securities (for example pooling many thousands of

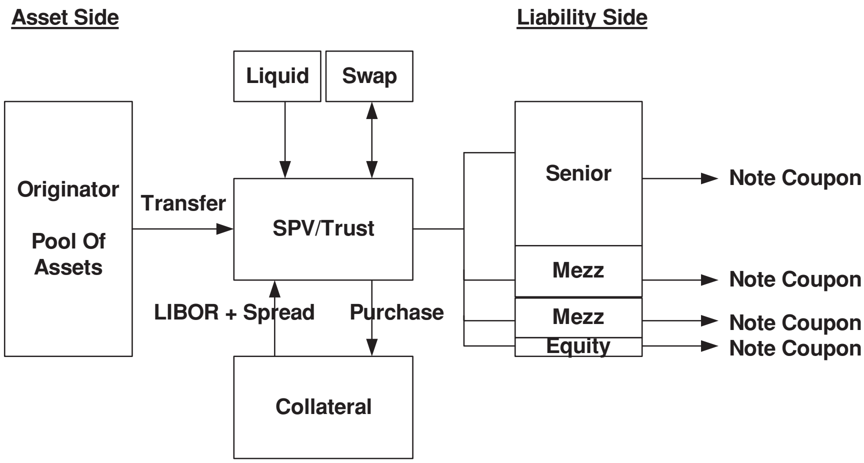

Figure 1.1 Securitisation of a pool of illiquid assets into tradable securities via the mechanism of an SPV. See the text for a full discussion.

mortgage commitments to individual domestic customers) into a new 'super' security. Figure 1.1 shows this schematically.

In this figure we have on the left-hand side the originator of the transaction (for example the bank). Let us assume the originator has a pool of illiquid assets which they own and wish to securitise. For example this might be a large number of corporate loans which are currently sitting on the balance sheet eating up regulatory capital. To securitise these assets the originator will physically transfer the ownership of these assets to a bankruptcy remote special purpose vehicle (or SPV, sometimes also referred to as the Trust). The SPV in essence purchases the assets from the originator. The funds for this are provided by the note investors, as described below, because the SPV has no funds of its own. From the originator's point of view the future (and potentially uncertain) cashflows from the assets have been transformed into an immediate cash payment, which can be beneficial to the originator's liquidity. The value of this cash payment is presumably the fair value of the expected future cashflows. The fundamental problem in mathematical finance is to develop realistic models for estimating the value of these future cashflows.

The SPV is a separate entity and most importantly is bankruptcy remote from the originator. This means that if some of the assets in the pool default, it will have no impact upon the originator (since these assets no longer sit on their balance sheet). Conversely, if the originator itself defaults it has no impact upon the SPV (and the notes that the SPV issues). Because the assets have been physically transferred the originator no longer has to hold regulatory capital against them, thereby freeing up the aforementioned '8%' for further investment in new business opportunities. Regulatory capital relief was one of the initial motivations behind securitisation.

The effect of this transfer of assets upon the underlying collateral (the corporate loans or individual mortgages) is minimal; the loans still have to be serviced, meaning that the SPV receives coupon payments (typically LIBOR plus a spread) and principal from the loans. However, it is the SPV (not the original owner) that will now be sensitive to any interruption to these cashflows due, for example, to defaults in the underlying pool. To facilitate all this, the role of the servicer (often the originator) in Figure 1.1 is to manage the collection and distribution of payments from the underlying pool (distributed to where we will now describe).

So far the discussion has focused on the 'asset' side of the structure. We now discuss the 'liability' side and introduce the concept of tranched exposures. The assets in the pool pay their owner income. The assets in turn can be used to fund further debt obligations, i.e. bonds or notes. The next step in the securitisation process is to sell the rights to the cashflows that the SPV is receiving (using these asset cashflows as security for the new debt to be issued). However, rather than selling the rights to individual cashflows or loans, the SPV sells exposure to a particular slice, or *tranche*, of the aggregate cashflows from the entire pool. For example, if the collateral is composed of 100 loans each of \$10 m then the total notional amount of loans issued is equal to \$1 bn. Each individual loan will pay a coupon of LIBOR plus a certain spread. The originator slices up this capital into a series of notes of sizes (notional amounts) \$800 m, \$100 m, \$70 m and \$30 m (for example). Each of these notes pays a coupon of LIBOR plus a spread based on the (aggregated) notional of that note. For example, the note with a notional of \$800 m may pay an annual coupon of 30 bps over LIBOR quarterly. Hence each coupon payment is (roughly) equal to $\$800\text{ m} \times (\text{LIBOR} + 30\text{ bps}) \times 1/4$. The investors in the notes pay the principal upfront, which is used to fund the purchase of the assets in the collateral pool, in return for receiving the periodic coupons and principal redemption at maturity. The risk, of course, to the investors is that the assets on the asset side do not deliver the expected returns (due to default, prepayment etc.).

The tranches are named according to their place in the capital structure and the legal seniority that the notes associated with the tranches have in terms of distribution of payments from the SPV. The most senior tranches have the first legal claim to the aggregate cashflows from the collateral pool and are referred to as the 'senior' tranches. The next most senior tranche has the next claim (typically the tranches in the middle of the capital structure are referred to as 'mezzanine' or mezz), all the way down to the most junior note at the bottom of the capital structure which is referred to as the equity tranche (or residual or first-loss piece). In the example shown in Figure 1.1 the capital structure has a senior tranche, two mezz tranches (typically referred to as junior and senior mezz) and an equity tranche. The (notional) sizes of the tranches are arranged so that the senior tranches have the largest notional and the equity tranche has the smallest amount (\$800 m and \$30 m respectively in the example given).

In general the income from the collateral pool is allocated down the capital structure starting with the most senior notes and working their way down to the most junior. Losses on the other hand are allocated from the bottom up. For example, if one of the assets in the pool defaults and 40% of the notional amount is recovered (leading to a loss of \$10 m $\times$ (100%–40%) $=$ \$6 m) it is the equity tranche that is impacted first. This results in a reduction of the notional amount of the equity tranche from \$30 m to \$24 m, reducing the payments that the equity note holder receives. In addition to this, going forward the asset pool now has less collateral and will therefore make fewer coupon payments. This leads to less cash being fed into the top of the capital structure, meaning less for the junior note investors once all the senior liabilities have been met.

The tranches are also rated by an external rating agency such as Moodys, S&P or Fitch. One of the upfront costs of securitising a pool of assets is the fees paid to the rating agency to provide a rating for the issued liabilities. The rating of a note is determined by the level of losses that can be sustained by the collateral on the asset side before the note cashflows on the liability side are impacted. Obviously the equity tranche is immediately impacted by losses and is therefore the riskiest tranche. For this reason it is typically unrated, and is often held by the originator of the deal (as a sign of confidence to investors that the assets in the underlying pool do not represent a moral hazard). To compensate the equity tranche holder for the enhanced risk they are taking on, the spread on this note is typically much larger than that on more senior tranches.

More senior tranches have a greater layer of protection (subordination) and so warrant higher ratings. It is important to note that a pool of assets that individually have poor ratings can, when securitised (with a priority of payments from senior to junior liability), result in new notes which have substantially better credit quality. This immediately broadens the appeal of the notes issued by the SPV to a whole range of new investors. For example, pension funds may be prohibited from investing in assets that are rated BBB due to their default risk (but which have a substantially enhanced yield compared to say AAA rated assets making them attractive to investors who are prepared to take on the risk). But a pool of BBB assets that are securitised and reissued as a series of notes including an AAA rated one is a different matter (the AAA rating being awarded based on the level of subordination that this note has relative to more junior notes). If the original BBB rated assets perform well then the pension fund benefits from this; on the other hand if the BBB rated assets do not perform well and default, the subordination provided by the equity and mezz tranches insulates the AAA notes from this. Everyone's a winner. That is, of course, unless large fractions of the underlying collateral start to default. For example, if all the underlying collateral were composed of US sub-prime mortgages which suddenly reset from a low teaser rate to 10%, this might have an impact on the returns of the notes.

One practical consideration of importance is the actual process of building up the collateral on the asset side. It is unlikely that an SPV will simply be able to go out and buy all of the collateral at a single instant in time. It is much more likely that the collateral pool will be assembled over an extended period as and when suitable assets that the manager of the structure deems fit to include in the pool become available. This is known as the ramp-up period and can last for several months. This represents a potential risk to the manager as they have to purchase and warehouse all of these assets until the structure is ready to sell on to investors. During the ramp-up period market conditions can change adversely, leading to the manager holding collateral which is not as attractive as initially anticipated. A solution to this ramp-up problem is provided by the use of credit derivative technology to construct the exposures to the assets synthetically, without actual physical ownership (more on this later). Another practical difficulty with the process described so far is that there is unlikely to be much standardisation amongst the type of collateral in the underlying pool. This means that for the types of structure described there is unlikely to be a highly liquid secondary market.

Finally there are two other components of the securitisation structure that need explanation. The role of the swap counterparty in Figure 1.1 is to provide a macro hedge against interest rate and FX rate fluctuations. There is also a liquidity provider. One of the less obvious risks of the structure described is mismatches in the timing of cashflows. For example, all of the assets on the asset side may pay coupons semi-annually, but the notes issued by the SPV may be quarterly. This would lead to short-term liquidity problems for the SPV in meeting its liabilities. To provide protection against this the liquidity provider (which may for example be the originating bank) will give the SPV lines of credit that it can draw down on, on an as-and-when needed basis.

1.3 Credit derivative products

In the previous section we described in quite general terms securitisation and tranching. In this section we discuss the application of these concepts to cashflow and synthetic CDOs. We also briefly describe some of the other important credit derivative products in the marketplace. More detailed business and economic descriptions of many of the products described in this section can be found in, for example, Gregory [2003], Das [2005], Chaplin [2005] and Chacko *et al*. [2006].

1.3.1 Credit default swaps (CDSs)

CDSs are the simplest example of a single-name credit derivative [Gregory 2003, Das 2005, Rajan *et al.* 2007]. The principal motivation of a credit derivative is to transfer credit risk (risk of default on outstanding obligations of a specified reference entity) between investors. A credit derivative will therefore usually have three counterparties to the trade: the counterparty wishing to purchase protection, the counterparty willing to sell protection and the reference entity to whom the bought and sold protection refers. For example counterparty ABC may own bonds issued by a separate reference entity C. ABC might be concerned about C defaulting (meaning ABC would receive no further coupons or its principal back if C did default) and may want to purchase protection against this risk. This protection is purchased by entering into a bilateral trade with counterparty XYZ who is willing to provide protection in return for a fee. A CDS provides the legal and financial mechanisms to achieve this transfer of risk.

Reference counterparties in the CDS market can include corporate entities as well as sovereign states (allowing protection to be purchased against a sovereign defaulting on its debts – this sort of protection is particularly popular for sovereigns in emerging markets where geopolitical risk can be a significant factor). The type of reference obligor asset that protection is bought or sold on has also evolved over time. Originally CDSs referenced the plain bonds of the reference asset. This has grown to include leveraged loans (LCDS) as well as asset backed securities (ABSCDS) as the underlying assets. CDSs are usually quoted on a spread basis, which is the coupon rate that is applied to the periodic protection payments. The par CDS spread is the spread (given the prevailing market conditions) which gives a fair value of the CDS at contract inception of zero. Protection is purchased for a specified period of time. During this period the protection purchaser makes periodic fee payments to the protection seller. These payments continue until the reference entity defaults or the protection period expires. If the reference entity defaults, subsequent coupon payments cease and the protection seller makes a contingent payment to the protection purchaser to compensate them for any loss. The contingent payment is a fraction of the notional amount of protection purchased. The fraction is termed the recovery rate and is determined in the market (by a dealer poll) at the time of the default.

As the credit derivative market has grown the uses of CDSs have evolved. They are now used as much for speculation and relative value trading (playing the default risk of one obligor off against another) as for providing long-term protection against the risk of a particular obligor defaulting. One of the important developments has been the growth of the market for trading protection over different time horizons. Initially, protection was purchased for a period of, typically, five years. As the market grew, investor demand for different time horizons led to the emergence of contracts specifying protection for maturities ranging from a few months up to ten and more years. As with the bond market, this introduced an additional degree of freedom that investors can express a view on: the likelihood of default over a certain time horizon. For example, a corporate that is subject to a private equity buy-out might be viewed by the market as having a higher long-term default risk than short-term. This is because the buy-out may typically be financed by the corporate taking on long-term debt (two-thirds of the buy-out cost is normal). Its liabilities in the short term are therefore less onerous than in the long term. Conversely, a whole sector may be perceived as having significant short-term default risk. For example, banks experiencing short-term liquidity problems might be viewed as a short-term risk, but not long term (if they survive the short term, they will go from strength to strength).

Having a term structure of CDSs also allows for investors to implement trading strategies based on the relative dynamics of different CDS maturities. This is analogous to what is observed in the interest rate market where interest rates are set for borrowing over a specific time horizon. Examples include so-called curve steepeners and flatteners [Rajan *et al*. 2007] where opposite trades are placed at different ends of the term structure of par CDS spreads.

Variations on the basic CDS trade have also appeared over time. Some of these variations are now briefly described.

1.3.1.1 Forward starting CDSs

A forward starting CDS is a CDS where the protection (purchased or sold) is specified to begin at a future point in time.

1.3.1.2 Credit default swaptions

Options on CDSs, or CD swaptions, are an important class of credit derivative because they allow investors to speculate on the volatility of CDS spreads. A CD swaption gives the holder of the option the right to enter into a CDS at a future date if the prevailing par spread at that time is such that the option is in the money. CD swaptions can in principle be of European, American or Bermudan exercise variety [Hull 1999, Wilmott 2000]. More details about the mechanics and strategies for trading CD swaptions may be found elsewhere [Rajan *et al*. 2007].

1.3.1.3 Recovery rate plays

For a plain, vanilla CDS the protection purchaser receives a payment upon default of the recovered amount of notional (assuming cash settlement for the moment – the different settlement mechanisms will be discussed in Chapter 3). The amount of notional recovered is a function of the prevailing market conditions at the time the payment is due. It is usually determined by a dealer poll (taking an average of the quotes received, having stripped out the highest and lowest quotes). Although this process is (relatively) transparent it does introduce an uncertainty into the amount that the protection purchaser will actually receive since the quotes provided by dealers will depend on a lot of different factors. Recovery rate strategies can be used to express outright views on recovery rates or to fix recovery rates at a desired level.

As the name suggests, the *recovery rate lock* is a contract that enables investors to 'lock-in' a specified recovery rate. The recovery rate lock was released in May 2006 as a specific contract by ISDA and is a product for trading views on recovery rates. Prior to the recovery rate lock, two separate CDS contracts (a standard CDS and a digital CDS – where the actual recovery rate paid in the event of a default is fixed at contract inception) were needed to express a similar view on recovery rates. The economics of the recovery rate lock and dual CDS position are the same, but the recovery rate lock is a single contract. In a recovery rate lock there is an agreed fixed recovery rate set at contract initiation. There are no upfront payments or periodic coupon payments. The only cashflow in the contract is the payment in the event of a credit event. The contract is physically settled. The lock buyer is the protection seller; they want the recovery rates to increase. The lock seller is the protection buyer and they want recovery rates to decrease.

1.3.1.4 Constant maturity CDS (CMCDS)

In a vanilla interest rate swap product the two counterparties swap a periodic stream of cashflows. One stream of cashflows may be based on a fixed coupon rate, and the other on a coupon rate which is a function of the prevailing LIBOR rate. A variation of this is a constant maturity swap (CMS). In a CMS the coupon rate for the floating leg payment is based on the prevailing par swap rate of a swap with a constant maturity at each coupon date [Hull 1999]. For example, a CMS may have a coupon based on the 10 year swap rate. At the first coupon date, say after 3 months, the par swap rate of a swap with a maturity of 10 years (from the 3 month point) will be determined and used as the coupon rate for the next floating leg. At the next coupon date at 6 months, the par swap rate for a swap of maturity 10 years (from the 6 month point) is determined and is used as the coupon rate for the next floating leg payment, and so on.

A CMCDS is a similar concept [Pedersen and Sen 2004, Das 2005, Rajan *et al.* 2007]. However, instead of a par swap rate, the rate that is determined is the par CDS spread for a CDS of a specified (constant) maturity. The coupon spread that the protection purchaser pays therefore changes (resets) at each coupon payment date. The CMCDS references an obligor's debt just like a normal CDS (popular CMCDS trades reference sovereign debt) or can reference a credit derivative index.

If there is a credit event during the lifetime of the contract, the contract terminates like a CDS with either cash or physical settlement.

The buyer of protection pays periodic coupons based on the current par spread (determined at each reset date). At the contract inception a participation rate (less than unity) is agreed. The participation rate is the multiple of the prevailing par spread which is actually paid at each coupon date. The buyer of protection is taking the view that the par CDS spread on the credit will increase by less than the spread implied by existing forward rates. If the spread remains low, then a low rate will continue to be paid at each reset date (and the protection buyer receives protection for a cheaper rate than would be currently available in the market). If the spread is higher than the expected forward value, the buyer of protection will have to pay a larger coupon for the protection. The initial participation rate reflects the steepness of the credit curve at contract inception. If a CMCDS position is combined with an offsetting CDS referencing the same obligor, then the overall position is default neutral. However, the spread payable allows investors to take curve and directional spread exposures to the obligor having isolated and removed the default risk.

1.3.2 Default baskets

A natural extension of a vanilla CDS is an instrument which provides protection against not one but a basket of obligors. Default baskets provide this [Chaplin 2005]. For example, an automotive manufacturer relies on different suppliers to provide the raw materials from which cars are constructed. The automotive manufacturer may want to purchase protection against any of the suppliers defaulting. This could be achieved by purchasing individual CDSs referencing each of the suppliers. Alternatively a default basket could be constructed composed of all the suppliers, which paid out on the first default of any of the individual constituents. Default baskets are leveraged positions with respect to a single-name CDS (they generate a higher yield than an equivalent – in notional terms – position in a single-name CDS). This is because the investors are exposed to a specific fraction (e.g. second default) of the default risk of the pool of obligors. Basket sizes are usually small (of the order of 3–20 obligors) with clip sizes typically \$10–50 m. They are usually unfunded products entered into by both counterparties at par (meaning that no initial exchange of cashflows takes place). Funded, or note variants (where the investor pays a principal amount upfront which they receive back at maturity or upon termination of the contract) are also traded, but these are not as common.

Higher-order extensions to default baskets, such as second-to-default, are also possible variations on the basic theme. Some baskets can also be extendable/cancellable and with digital payoffs. However, by far the most common form of contract is the first-to-default basket. Default baskets also tend to be bespoke instruments. That is, each default basket provides protection on a unique pool of obligors. This is because baskets are often used to provide protection against particularly worrisome credits in a larger synthetic CDO. For example, a hedge fund may sell protection on an equity tranche of a bespoke, synthetic CDO. If they have particular concerns about a small group of obligors within the pool (for example, obligors in the banking sector), they may consider buying protection on a small basket (a typical basket size is five obligors) referencing these obligors. This provides protection against default of these names, but at the expense of reducing the positive net carry for the hedge fund (since the basket protection requires some funding).

1.3.3 Credit indices

Credit indices are another natural extension of the CDS concept [Rajan *et al*. 2007]. Like default baskets they are a product that provides exposure to a portfolio of assets, in this case the CDSs of a reference portfolio of obligors. The crucial feature of credit indices is that they are constructed, traded and managed according to transparent and standardised rules. This degree of standardisation has proven absolutely pivotal in the development of the credit derivative market. The first-generation credit products were highly bespoke, requiring not only traders and structurers to make the deal, but also legal and documentation experts to ensure that there was a clear understanding between the counterparties as to what constituted a default event, what suitable *pari passu* deliverable obligations were etc. This meant that it was not possible to transact large volumes of deals. Standardisation of trades obviously is a key facilitator of liquidity.

The credit indices in their current incarnation were introduced in 2004. There are two broad families of indices: the CDX family based on North American investment grade corporates and the iTraxx family based on European and Asian investment grade corporates. The indices are managed and administered by an independent calculation agent, MarkIt [MarkIt] who provide the transparent construction and trading mechanics of index products.

Unlike a single-name CDS which terminates on the occurrence of a single default, an index CDS does not terminate. Instead the protection seller compensates the protection purchaser (as in a normal CDS), but the contract then continues with the defaulted obligor removed from the pool. The protection purchaser continues to make coupon payments with the same coupon spread, but based on a reduced notional amount (the reduction corresponds to the notional amount of the defaulted obligor). This continues to the overall maturity of the contract.

Index products are very good for rapidly expressing views, both long and short, as to the macro state of the credit markets. They are also relatively effective for hedging other credit derivative exposures. Index products are followed not just by those actively engaged in the credit markets, but also by market participants in other sectors. Specifically, the common indices that are used to gauge the mood of the structured finance and credit markets include the following:

- the iTraxx and CDX indices (main and crossover indices) to gauge corporate credit quality;
- the ABX indices (see Chapter 11) to gauge ABS referencing home equity collateral quality;
- the LCDX (an index of 100 LCDS contracts) to gauge the corporate loan market.

1.3.4 Collateralised debt obligations (CDOs)

CDOs and synthetic CDOs in particular are the most prominent and important example of multi-name credit derivatives. CDOs were first developed in the 1980s as a financing vehicle that allowed an owner of a pool of assets to finance the purchase of a pool of assets on a non-recourse basis at tighter spreads than were paid by the underlying assets [Gregory 2003, Das 2005, Kothari 2006].

From the original bespoke deals that were arranged individually between originator and investor (see the above references for examples of early deals and their economic rationale) they have evolved into a sophisticated and liquid asset class in their own right, particularly the single-tranche CDO (STCDO) market. The introduction of standardised credit derivative indices and subsequent trading of tranched exposures to portions of the portfolios' losses has resulted in the appearance of correlation as a new asset class to trade (although correlation is only the mechanism traders use to communicate prices and assess the relative value of different instruments, the fundamental tradable is still price).

A CDO is a specific example of the general securitisation structure shown in Figure 1.1. To characterise a CDO we need to specify the following (in addition to standard contract conventions such as holiday calendars, day-count conventions, calculation agents etc.):

- the composition of the underlying asset pool (type of asset, maturities, notional amounts, coupons etc.);
- the capital structure of the issued liabilities (the tranche structure);
- the maturity of the CDO;
- the priority of payments to the investors (the cashflow waterfall);
- the allocation of losses on the asset side to the investors on the liability side;
- whether the deal is managed or static (can names be substituted in and out of the underlying pool).

By an appropriate choice of these parameters the issuer and investor can tailor the benefits of the structure to fit their risk/return requirements. A big motivation for structuring and issuing CDOs is regulatory capital relief and ratings arbitrage. Typical regulatory regimes require a CDO investor to hold 100% capital against equity tranche exposure and of the order of 8% against exposures to other tranches. This is to be compared with the requirement to hold 8% against the entire pool of assets which has been securitised. The one-for-one holding of capital against equity exposures reflects the enhanced risk of the equity tranche which is a highly leveraged exposure to the overall pool of assets (it is leveraged because it is a thin slice of the overall portfolio losses, meaning the holder of the tranche is highly exposed to the first losses in the underlying pool). This leverage enables the originator to enhance their yield on the underlying assets by securitising the assets and purchasing back the equity tranche. This facilitates ratings arbitrage.

In addition to this form of arbitrage, it is also possible that the securitisation and tranching process can lead to a beneficial spread differential between the asset and liability sides. For example, the (notional) weighted average spread of the liabilities may turn out to be less than the spread received from the assets. The equity tranche holder can therefore earn a positive carry simply by arranging the transaction. In order to achieve this, however, the assets have to be chosen carefully (in terms of their ratings etc.) as does the structure of the liabilities. Why might this be possible? One possibility is that CDO liquidity is less than that of the corporate bond market, meaning that CDO investors will require an enhanced spread to compensate for the liquidity risk. CDOs are also leveraged instruments meaning that mark-to-market volatility is higher compared to straight corporate bonds, again leading to investors demanding an increased spread to compensate for this.

1.3.4.1 Cashflow CDOs

A cashflow CDO has the following general characteristics [De Servigny and Jobst 2007, Rajan *et al.* 2007]:

- an asset pool drawn from a potentially diverse range of collateral (see below);
- a tailored, highly bespoke capital structure;
- a long legal maturity;
- a detailed cashflow waterfall that sets out how the interest, principal and recovery amounts generated by the asset pool are propagated through the issued liabilities from the top of the capital structure down to the bottom (including appropriate collateral and interest coverage tests to protect senior note holders from losses in the asset pool);
- a detailed specification setting out how losses on the asset side are propagated up from the bottom of the capital structure;
- typically cashflow CDOs are actively managed deals where the portfolio/asset manager can choose to substitute collateral dynamically in and out of the asset side. For the benefit of their investment expertise and market insight the manager will charge a fee. The indenture of the deal will set out the terms and conditions for substitution and reinvestment.

The securitisation example discussed in the previous section refers primarily to a cashflow CDO. Cashflow CDOs are a natural extension of asset backed securitisation (ABS) technology. The most common cashflow CDOs are constructed from underlying collateral which is the following:

- corporate loans (a collateralised loan obligation or CLO);
- corporate bonds (a collateralised bond obligation or CBO) although the popularity of these products has diminished somewhat;
- senior or mezzanine tranches of existing asset backed securities (ABS), known as ABSCDOs. Suitable ABS collateral can include residential or commercial mortgages, home equity loans, automobile loans, student loans, credit card receivables etc. (even future gate receipts from a football stadium can be securitised in order to raise cash to finance other projects – the risk to the investors being that the team is not successful);
- combinations of all the above collateral types.

The complexity of cashflow CDOs arises principally from the diversity of collateral that can be incorporated in the underlying pool (as described above). An additional complexity arises if a note issued by the master CDO is backed by collateral which is itself securitised providing tranched exposures to other assets, which in turn might be securitised and tranched etc. That is to say, there may exist multiple layers of securitisation and tranching within the overall CDO structure. CDO managers and investors need to understand very clearly what the underlying collateral they are purchasing exposure to is. It is very easy to purchase assets which have deeply embedded exposures to all sorts of unwanted asset classes. This also raises the possibility that the same underlying assets might appear multiple times in different pools all referenced by the master CDO.

Therefore, on the one hand CDOs, being portfolio products, provide investors with diversification (compared to a position in, for example, a single corporate bond). On the other hand, if the same assets appear in multiple securitisations the degree of diversification might be far less than expected or desired. Indeed if there are a significant number of assets which appear in multiple pools all referenced by the same master CDO, 'cliff' risk becomes important. This refers to the possibility that a master CDO might have, on paper, a large degree of subordination and thus appear to be well protected against default. This may indeed be the case. On the other hand if the same assets appear in different pools, it may be that it only takes one or two particularly highly connected obligors to default before all the subordination is rapidly eaten away. As a consequence of this, a very important aspect of risk managing cashflow CDOs (including ABSCDOs) is constant monitoring of the composition of the underlying collateral to identify hot-spots of asset-class concentration (drilling right down as far as reasonably possible to the atomic level of individual exposures).

Another complex issue in the management of cashflow CDOs is the timing of cashflows. If, on the asset side, there is a rich mixture of different types of collateral, it is highly likely that the different assets will have coupons which pay with different tenors leading to a cashflow profile that is relatively stable over time. In an extreme example, however, all of the collateral on the asset side might pay coupons semi-annually. On the other hand, the issued liabilities may require quarterly payments. This means that there is a funding mismatch between the cash the SPV has to pay to meet its liabilities, and the cash it receives from the asset side. This is an issue that has to be monitored on a daily basis to ensure that the SPV has enough liquidity from its long-term investments to meet its short-term liabilities. To mitigate against this risk, cashflow CDOs usually have access to short-term revolving credit facilities (from banks) to provide emergency liquidity.

Other types of risk that cashflow CDOs are sensitive to stem principally from interruptions to the stream of cashflows on the asset side of the structure. Interruptions can occur due to defaults amongst the collateral or amortisation of the asset notional due, for example, to prepayments of loans (residential or commercial mortgages, automobile loans, home equity loans etc.). When interruptions to the asset cashflows occur this impacts the rated liabilities because there is now less cash being fed into the top of the waterfall.

1.3.4.2 Synthetic CDOs

Cashflow securitisation physically removes assets from the originator's balance sheet. In some cases, however, selling assets can be detrimental to a relationship with a client. Selling a corporate loan to an SPV is hardly a sign of confidence or respect for that corporate, or the value the originator places on doing business with them. Synthetic securitisation can alleviate this problem by generating the exposure to the obligors synthetically, without the physical transfer of ownership of assets. Synthetic CDOs utilise credit derivative technology to achieve this [Joannas and Choudhry 2003, De Servigny and Jobst 2007, Rajan *et al.* 2007].

Synthetic CDO issuance was, like cashflow CDO issuance, initially driven by regulatory capital relief motivations (capital relief is achieved synthetically via the use of credit derivative technology). The next evolution of the market was full capital structure synthetic CDO deals (which in turn helped the development of the CDS market). These positions required relatively simple post-trade risk management. Single-tranche CDOs were the next step in the evolution of the product. These products allowed dealers to retain and hedge individual slices of the capital structure that could not be placed in the marketplace. These products were also attractive to new types of investors, in particular hedgies wanting leveraged exposure to a particular slice of the tranche.

A synthetic CDO (also sometimes referred to as a collateralised swap obligation or CSO) has the following general characteristics.

- An asset pool composed solely of vanilla CDSs each referencing an individual obligor (the pool can be bespoke or standardised).
- In some cases a tailored, highly bespoke capital structure; in other cases (STCDOs) the capital structure is standardised.
- A relatively short legal maturity and matching of the maturities of the assets (CDSs) and liabilities (notes).
- A simple cashflow waterfall structure. Compared to cashflow CDOs, synthetic CDOs have very simple cashflow mechanics, only requiring the management of default losses (and not having complex waterfall structures). Each note pays a coupon based on the remaining notional of the note's tranche.
- Losses on the asset side are propagated up from the bottom of the capital structure (and also sometimes the recovered amounts are amortised from the most senior tranche downwards). Each loss reduces the tranche notional until all of the tranche's notional is exhausted. After this point subsequent losses have no impact upon the tranche.
- Cashflow mechanisms can be on an all-upfront basis (the only coupon cashflow is at contract inception), an all running coupon basis (there is no upfront payment, but there are periodic coupon payments) or a mixture of the two (combining an upfront payment with periodic coupons).
- The composition of the asset pool can be either static or dynamic, depending principally upon who the investor is (hedge funds for example are big users of managed tranches since it is their business to invest funds and dynamically manage the investment to achieve the best return, whereas investment banks tend to be users of static tranches to hedge bespoke positions).

Synthetic CDOs are created by applying credit derivative technology to portfolios of assets. They are portfolios of CDSs whose cashflows are tranched and these tranches are sold to investors. They facilitate the transfer of risk synthetically, without physically transferring the assets off the originator's balance sheet (thereby removing the potential problem with client management). In terms of regulatory capital relief, synthetic CDOs do act as a hedge for some of the credit risk. This allows some capital to be reduced, but the reduction is not as great as for cashflow CDOs where the assets are physically taken off the balance sheet. Because the assets are not physically purchased the ramp-up, availability of collateral and warehousing issues experienced with cashflow CDOs are greatly alleviated, although not removed completely.

The simplest rationale for investing in synthetic CDOs is to express a view on the default risk of the obligors in the underlying portfolio (purchasing or selling protection accordingly). A counterparty who is concerned about default risk can purchase protection from another counterparty willing to sell them the protection.

Conversely a yield hungry hedge fund (whose portfolio managers get paid to take risks and outperform the market) might be willing to sell protection on a portfolio of obligors if their view of the default risk is sufficiently sanguine. If the deal is also managed it gives the manager additional freedom to dynamically replace poorly performing assets in the pool. Other uses for synthetic CDOs include relative value trading (the CDO may have a greater yield than an equivalently rated bond), expressing leveraged views (both spread leverage and leveraged exposures to the slices of tranche losses) and hedging of other exposures.

Initially synthetic CDO issuance was small compared to that of cashflow CDOs. However, as the credit derivative market (for CDSs) has grown, synthetic CDO issuance has outgrown that of cashflow CDOs. Part of the growth of the CDS market has in fact been driven by the synthetic CDO market which needs large volumes of CDSs for hedging purposes. The synthetic CDO market is very dependent on the CDS market for its operation.

1.3.4.3 Single-tranche CDOs (STCDOs)

One of the disadvantages with bespoke CDOs (either cashflow or synthetic) is that it takes time to ramp up the underlying collateral pool. This is even true for synthetic CDOs despite the fact that the exposure to the underlying obligors is generated synthetically via credit derivatives. The problem is that the protection must still be purchased or sold, meaning there needs to be a counterparty on the other side of the trade willing to do the deal at an acceptable market level, and that all of the tranches must be placed into the market near simultaneously. The protection seller then has to hedge the relevant exposures separately.

An elegant solution to this problem is provided by the standardised indices which have clear and transparent rules for construction, modification and trading (standardisation of course boosts liquidity by making trading between counterparties easier). The composition of the indices is also static through time (except at the roll dates when the portfolio is updated). Tranched exposures on these indices have standardised

- pool composition (the index),
- attachment/detachment points,
- maturities.

An index tranche provides investors with exposure to a slice of the loss distribution of the credit index portfolio. Greater standardisation of product has led to greater liquidity since investors are familiar with the terms and conditions of the deal leading to less difficulty in trading and encouraging more trading etc. An STCDO trades only a single tranche of the full capital structure. STCDOs therefore have

Table 1.1 *Standardised tranches defined on the iTraxx main index*

Tranche	Attachment point (%)	Detachment point (%)	Attachment point ($ m)	Detachment point ($ m)	Tranche notional ($ m)
Equity	0	3	0	37.5	37.5
Junior mezz	3	6	37.5	75	37.5
Senior mezz	6	9	75	112.5	37.5
Senior	9	12	112.5	150	37.5
Super senior	12	22	150	275	125

Table 1.2 *Standardised tranches defined on the CDX main index*

Tranche	Attachment point (%)	Detachment point (%)	Attachment point ($ m)	Detachment point ($ m)	Tranche notional ($ m)
Equity	0	3	0	37.5	37.5
Junior mezz	3	7	37.5	87.5	50
Senior mezz	7	10	87.5	125	37.5
Senior	10	15	125	187.5	62.5
Super senior	15	30	187.5	375	187.5

the advantage that they allow a very targeted exposure to the portion of the capital structure that the investor wishes.

Standardised tranches are defined on both the iTraxx and CDX main indices. The attachment/detachment points are shown in Tables 1.1 and 1.2. Each obligor in the index has a notional of $10 m and there are 125 obligors in each index. To hedge against ‘small’ spread movements, index tranches trade with an opposite index position (of size delta $\times$ tranche notional). The price that is quoted is known as the delta exchange price (see Bluhm and Overbeck [2007] for an example of the mechanics of delta exchange).

STCDOs based on the standardised indices are almost always static, where the underlying asset pool composition does not change over time. Investment banks are big users of STCDOs for hedging bespoke exposures. The bespoke exposures are typically with respect to investors such as hedge funds who will want to construct bespoke STCDOs to express a view on the behaviour of the credit markets. Bespoke synthetic CDOs enable customisation in terms of

- structure of protection payment coupons (fixed, floating or zero-coupon),
- funded (note)/unfunded (swap) form,
- static or managed portfolio and the composition of that portfolio,
- custom yield/leverage,

- tranche size and subordination,
- principal protected structures (to ensure that investors receive their original investment back).

If the investment bank is on the other side of the trade they will want to hedge their exposure. Slicing up the indices and trading only specific slices of the risk solves the problem of ramping up the underlying portfolio. However it introduces new complexities:

- correlation sensitivity (later chapters will describe in detail how correlation impacts STCDOs, but for the moment we simply observe that STCDOs facilitated the introduction of correlation as a new financial variable to trade);
- how to hedge the exposure to the bespoke tranche.

This latter point is one of the most important issues in the STCDO market. One of the major (if not the major) modelling challenges in the synthetic CDO world is determining how to value bespoke tranches consistently with standardised tranches (the bespoke may only differ from the standard tranche by a few obligors in the underlying pool, non-standard attachment/detachment point etc.). If the bespoke tranche had a set of properties that were equivalent to an index tranche, then it would be easy to hedge the position. However, it is unlikely this will be the case. Therefore to hedge the bespoke exposure the investment bank must try and ‘map’ the bespoke to an appropriate index position. Prices of bespokes are mapped to index products because these are the only readily observable prices in the marketplace.

In terms of risks, synthetic CDOs do not have the prepayment type of risks that are common to cashflow CDOs since the only early termination mechanism for a CDS is default. The main risks to synthetic CDOs arise from the spread movements of the underlying CDSs as well as the possibility of outright obligor default. In addition to this STCDOs introduce another very important risk sensitivity: correlation (because a STCDO exposes the investor only to a certain fraction of a pool’s loss exposure and the pool’s loss distribution is heavily dependent on the correlation between the constituents).

Cashflow CDOs are typically buy and hold investments, where the exposure is usually long only (i.e. the investor buys the notes issued by the SPV). Synthetic CDOs on the other hand, being highly liquid, allow for investors to take long and short positions (selling or buying protection respectively). They also allow for unleveraged (index) positions or leveraged (tranche) positions to be taken. The synthetic CDO market principally consists of a flow business based on the trading of standardised index tranches. These trades are used for relative value trading as well as for hedging certain exposures. More importantly there is trading in bespoke synthetic CDO tranches. These bespoke tranches may have non-standard attachment/detachment points, maturities or underlying pool composition (or all three).

1.3.4.4 Options on index tranches

Recently we have seen the introduction of derivatives which have synthetic CDOs as their underlying security. Examples include forward starting CDOs and options (European, American and Bermudan) on synthetic CDOs. Options on tranches are options on the spread of the tranche and allow investors to trade the volatility of the tranche spreads. A put is defined as the right to buy protection (selling the risk to another counterparty), and a call is defined as the right to sell protection (buying the risk).

1.3.4.5 Tranchelets

Tranchelets are very thin tranches on a given portfolio. Because the tranche width is smaller than usual, tranchelets provide an even more granular and leveraged exposure to the portfolio losses.

1.3.4.6 Zero-coupon equity tranches

In a zero-coupon equity tranche there are only two cashflows: one at trade inception and one at maturity (it is therefore an all-upfront cashflow mechanism). This is similar to a zero-coupon bond. The payment at maturity depends only on the notional losses incurred by the tranche during the lifetime of the tranche (but does not depend upon the timing of the losses). The structure can provide a high internal rate of return with a payout whose timing is known (as opposed to a normal synthetic CDO structure where the timing of contingent payments is unknown). The high rate of return stems from the fact that the initial upfront payment prices in the full expected losses of the portfolio (upfront). If this amount is higher than the actual realised loss then the investor benefits at the maturity of the tranche.

1.3.4.7 Structured investment vehicles (SIVs)

A SIV is not technically a credit derivative, it is more correctly described as a structured finance vehicle. However, the technology and investment rationale for SIVs is quite similar to that described for CDOs so it is worth digressing for a moment to mention this type of vehicle. It is also true that SIVs have achieved a certain amount of notoriety due to their role in the credit crunch of 2007. A very good description of SIVs can be found in De Servigny and Jobst [2007].

Like an SPV, a SIV is a structured finance vehicle that is bankruptcy remote from the originating institution. Also in common with an SPV, a SIV issues rated liabilities (notes) and uses the proceeds from investors to purchase assets (collateral). The core investment rationale of a SIV is that it will purchase high-yielding, long-term assets and fund these purchases with short-term, lower yielding paper. The SIV manager will therefore be able to make a profit on the difference between the monies received from the assets and the monies paid to investors to service the debts. In essence the SIV is taking long-term exposure (risk) to high-grade assets and funding it with short-term debt. SIVs are typically buy-and-hold investment vehicles. Active trading of assets by the portfolio manager is only undertaken when assets mature and the proceeds need to be reinvested (and of course trading is undertaken when investors initially invest in the issued liabilities). SIVs usually do not have a well-defined lifetime, instead simply continuing to operate indefinitely (although so-called SIV-lites have a well defined legal maturity beyond which the SPV ceases to operate).

The collateral purchased will typically have high credit ratings (usually AAA to A with little BBB rated exposure). A broad range of collateral can be purchased, similar to the types of collateral that appear in cashflow CDOs. Some SIVs also have synthetic credit derivative exposure. Assets are typically well diversified according to type, rating, sector and geography. Strict concentration limits (for example, limiting the geographical concentration in any one region) are specified as part of the SIVs operating manual. The capital model (see the following paragraph) will impose penalties for breaches of these concentration limits. For example, if a rating concentration is breached the SIV will have to liquidate assets in order to re-establish the appropriate capital buffer.

Because the assets are highly rated, the default risk to the collateral on the asset side is expected to be low. When an asset matures the income stream from that asset disappears. Therefore it is necessary to replace the assets that are rolling off with new assets in order to ensure there is enough income coming in to service the liabilities. Similarly, short-term liabilities must be rolled in order to fund new asset purchases. SIVs dynamically manage the debt that they issue. In particular, existing debt can be rolled and new debt issued at a very high frequency (weeks). This provides a constant stream of paper. Typically the principal amounts raised from the liabilities will be used to fund the purchase of (longer maturity) assets on the asset side. Usually in a SIV there are two tranches of liabilities. Senior liabilities are rated AAA and are issued in several classes. Capital notes are the mezz piece and were originally one tranche. But the capital notes can also be tranched into a rated and unrated (first-loss) piece. SIVs therefore incorporate the securitisation and tranching technologies that are the basis of the CDO world.

An important role of the SIV manager (in addition to choosing the collateral to purchase and the liabilities to issue) is to manage the mismatch between assets and liabilities and the consequences of a temporary shortfall in liquidity. Because of this mismatch, temporary liquidity shortfalls must be managed with internal/external liquidity in the form of bank lines, breakable deposits, committed