CAMBRIDGE

Cambridge University Press & Assessment
978-0-521-89788-4 — Synthetic CDOs
C. C. Mounfield
Excerpt
More Information

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.

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE

Cambridge University Press & Assessment
978-0-521-89788-4 — Synthetic CDOs
C. C. Mounfield
Excerpt
More Information

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

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE

Cambridge University Press & Assessment
978-0-521-89788-4 — Synthetic CDOs
C. C. Mounfield
Excerpt
More Information

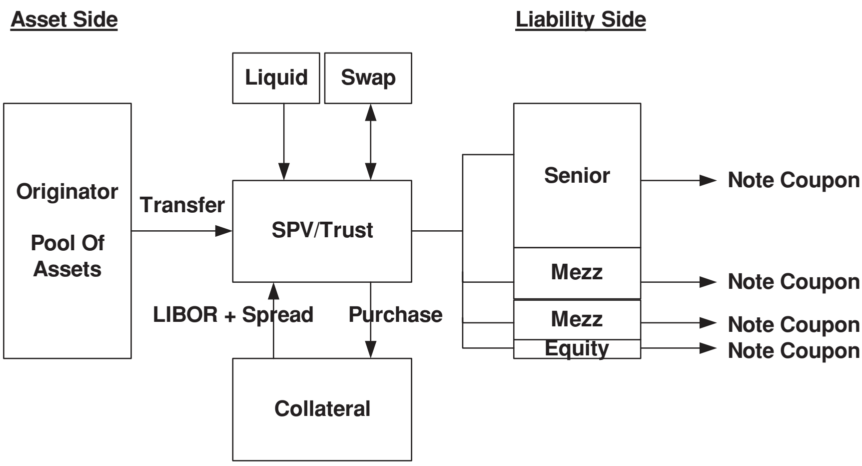

Figure 1.1 Securitisation of a pool of illiquid assets into tradable securities via the mechanism of an SPV. See the text for a full discussion.

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

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE

Cambridge University Press & Assessment
978-0-521-89788-4 — Synthetic CDOs
C. C. Mounfield
Excerpt
More Information

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,

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE
Cambridge University Press & Assessment
978-0-521-89788-4 — Synthetic CDOs
C. C. Mounfield
Excerpt
More Information

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 × (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

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE

Cambridge University Press & Assessment
978-0-521-89788-4 — Synthetic CDOs
C. C. Mounfield
Excerpt
More Information

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 subprime 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].

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE
Cambridge University Press & Assessment
978-0-521-89788-4 — Synthetic CDOs
C. C. Mounfield
Excerpt
More Information

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

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE

Cambridge University Press & Assessment
978-0-521-89788-4 — Synthetic CDOs
C. C. Mounfield
Excerpt
More Information

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

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE

Cambridge University Press & Assessment
978-0-521-89788-4 — Synthetic CDOs
C. C. Mounfield
Excerpt
More Information

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.

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE

Cambridge University Press & Assessment
978-0-521-89788-4 — Synthetic CDOs
C. C. Mounfield
Excerpt
More Information

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

© in this web service Cambridge University Press & Assessment www.cambridge.org