Jeffery D Amato
+41 61 280 8434
jeffery.amato@bis.org

Jacob Gyntelberg
+41 61 280 9156
jacob.gyntelberg@bis.org

CDS index tranches and the pricing of credit risk correlations[1]

Standardised loss tranches based on credit default swap (CDS) indices have increased liquidity in the market for credit risk correlations. Although progress is being made, quantitative modelling of these correlations is complex and not yet fully developed.

JEL classification: G12, G13, G14.

One of the most significant developments in financial markets in recent years has been the creation of liquid instruments that allow for the trading of credit risk correlations. Prime among these instruments are CDS index tranches. Broadly put, index tranches give investors, ie sellers of credit protection, the opportunity to take on exposures to specific segments of the CDS index default loss distribution. Each tranche has a different sensitivity to credit risk correlations among entities in the index. One of the main benefits of index tranches is higher liquidity. This has been achieved mainly through standardisation, yet it is also due to the liquidity in the single-name CDS and CDS index markets. In contrast, possibly owing to the limited liquidity in the corporate bond market, securities referencing corporate bond indices have not been actively traded.

The standardisation of index tranches may prove to be a significant further step towards more complete markets. Credit risk correlations have always been key risk components in portfolios of credit-risky securities. However, up until now, standardised products for the trading of credit risk correlations have not been available. The emergence of index tranches therefore fills a gap in the ability of the markets to transfer certain types of credit risks across individuals and institutions.

We examine CDS index tranches in this article. In the first section we introduce these securities, focusing on the mechanics of CDS-based contracts and market liquidity. In the second section we discuss the pricing of CDS index

[1] We thank JPMorgan Chase for providing us with data; Rishad Ahluwalia, Jakob Due and Mike Harris of JPMorgan Chase for useful discussions; Henrik Baun, Claudio Borio, Ingo Fender, Frank Packer and Eli Remolona for helpful comments; and Marian Micu for research assistance. The views expressed in this article are those of the authors and do not necessarily reflect those of the BIS.

tranches, with an emphasis on how these instruments allow for the trading of credit risk correlations.

CDS-based contracts: characteristics and liquidity

To understand the advantages offered by CDS index tranches for the trading of credit risk correlations, it is first necessary to understand their composition, namely, the structure of CDS indices and the underlying single-name CDS contracts.

CDS contracts

Single-name CDSs are the building blocks

A single-name CDS contract is an insurance contract covering the risk that a specified credit defaults. Following a defined credit event, the protection buyer receives a payment from the protection seller to compensate for credit losses. In return, the protection buyer pays a premium to the protection seller over the life of the contract.[2]

There are two main reasons why CDS contracts are more liquid than most corporate bonds. First, they are more standardised. For instance, the credit events that trigger payment to the protection buyer are now clearly defined in the ISDA credit derivatives definitions (ISDA (2003)).[3] This is also the case for the settlement method.[4] Second, CDS contracts allow market participants to go long credit risk without a cash payment, as well as go short credit risk with less difficulty and at lower cost than with corporate bonds.

CDS indices

A CDS index contract is an insurance contract covering default risk on the pool of names in the index. Index contracts differ slightly from single-name securities. The main difference is that a buyer of protection on the index is implicitly obligated to pay the same premium, called the fixed rate, on all the names in the index. In addition, index contracts restrict the eligible types of credit events to bankruptcy or failure to pay.[5] In the case of a credit event, the entity is removed from the index and the contract continues (with a reduced notional amount) until maturity.

Liquidity of CDS indices is enhanced by ...

The market liquidity of CDS index contracts is enhanced by: (1) the emergence of widely accepted benchmark indices, which comprise the most

---

2 Several sources contain descriptions of CDS contracts and their features (eg Anson et al (2003) and O'Kane, Naldi et al (2003)).

3 Credit events include bankruptcy, failure to pay, repudiation and material restructuring of debt (including acceleration).

4 Payoffs can be settled either by cash (with the protection buyer receiving par minus the default price of the reference asset) or in physical form (where the protection buyer delivers the defaulted security to the protection seller in return for a cash payment of par).

5 This corresponds to the no-restructuring (XR) documentation clause in single-name CDS contracts, ie excluding debt restructuring as a triggering event (see ISDA (2003) for a description of documentation clauses). See O'Kane, Pedersen and Turnbull (2003) for a discussion of common market practices, as well as Packer and Zhu (this issue of the *Quarterly Review*).

liquid single-name CDS contracts in the market and have a group of global dealers committed to market-making; (2) a clear geographical focus, relatively stable sector-rating composition and standardised maturities for each index; and (3) the availability of two different contract formats. We consider each element in turn.

... the creation of benchmark indices ...

First, the main traded CDS indices have now been consolidated into a single family under the names DJ CDX (for North America and emerging markets) and DJ iTraxx (for Europe and Asia); see Table 1.[6] The composition of the new indices is chosen by participating dealers based on the liquidity of individual contracts, ie the most actively traded names are included. Once formed, an index remains static over its lifetime, except for entities that default, which are eliminated from the index. However, every six months a new rebalanced index is launched and associated "on-the-run" securities are issued.

... across regions and sectors ...

Second, indices have been created for the main currencies, investment grade and non-investment grade credits and the main sectors. At the

CDS indices[1]

By region

	North America	Europe	Japan	Asia excl Japan	Australia	Emerging markets
Master	CDX.NA.IG (125) CDX.NA.HY (100)	iTraxx Europe (125) iTraxx Corporate (52)[4] iTraxx Crossover (30)[5]	iTraxx CJ (50)[2]	iTraxx Asia (30)	iTraxx Australia (25)	CDX.EM (14)[3]
Sub-indices	Financials (24) Consumer (34) Energy (15) Industrials (30) TMT (22) HiVol (30) B (44) BB (43) HB (30)	Financials (15) Autos (10) Consumer cyclicals (15) Consumer non-cyclicals (15) Energy (20) Industrials (20) TMT (20) HiVol (30)	Financials (10) Capital goods (10) Tech (10) HiVol (10)	Korea (8) Greater China (9)[6] Rest of Asia (13)[7]	None	None

[1] Earlier generations of DJ Trac-x and iBoxx indices are still traded. This table summarises the composition of the most recently issued series, DJ CDX and DJ iTraxx, which are a by-product of the merger between the DJ Trac-x and iBoxx families. The number of reference entities in each index is given in parentheses. [2] Maximum of 10 names in a given sector. [3] Includes only sovereigns: Brazil, Bulgaria, Colombia, Korea, Malaysia, Mexico, Panama, Peru, the Philippines, Romania, Russia, South Africa, Turkey and Venezuela. [4] Includes the largest, most liquid non-financial names from the iBoxx EUR Corporate bond index. [5] Most liquid non-financial names rated BBB/Baa3 or lower and on negative outlook. [6] Includes China, Hong Kong SAR and Taiwan (China), with at least two names from each. [7] Includes India, Malaysia, the Philippines, Singapore and Thailand.

Table 1

[6] Two competing families of indices (Trac-x and iBoxx), supported by different dealers, were initially launched in 2003. Last year these indices were merged to form the new indices, which are administered by Dow Jones.

Sectors and ratings distributions

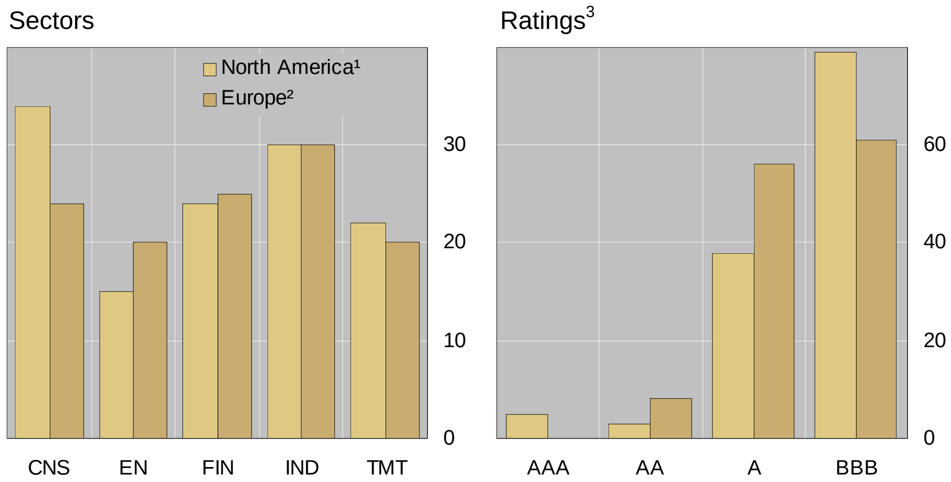

Note: CNS = consumer, EN = energy, FIN = financials, IND = industrials, TMT = technology, media and telecommunications. For Europe, CNS includes consumer cyclicals and consumer non-cyclicals, and IND includes industrials and autos.

[1] DJ CDX.NA.IG.3. [2] DJ iTraxx Europe Series 2. [3] Average rating from Moody's, Standard & Poor's and Fitch (where available).

Sources: Bloomberg; BIS calculations. Graph 1

investment grade level, the broad indices in North America (CDX.NA.IG) and Europe (iTraxx Europe), which are the most actively traded, are each composed of 125 reference entities, with an equal weighting given to each. There are also indices for selected sectors; an index based on names with high systematic exposures (ie high market betas); indices composed of speculative grade firms; and indices for regions other than North America and Europe, such as Japan, Asia (excluding Japan), Australia and a selection of emerging market countries. Graph 1 shows the distribution across sectors and ratings in the most recently issued versions of CDX.NA.IG and iTraxx Europe. Securities on the main indices are available at five- and 10-year maturities.

… and availability of different contract types

Third, two types of index contracts, unfunded and funded, are traded to better tailor the securities to investors' preferences with respect to funding format and counterparty risk exposure. An unfunded contract is simply a multi-name CDS; the funded version is a bond, where, at origination, the buyer of protection receives a pool of collateral securities from the protection seller and pays an upfront notional, in addition to paying a quarterly premium. In an unfunded contract, the protection buyer is exposed to counterparty risk, whereas in a funded transaction the protection buyer is exposed to the risk of credit deterioration in the collateral pool (but not to counterparty risk).[7]

Bid-offer spreads have been tight

The relatively liquid nature of these instruments, compared to other credit products, has been reflected in fairly tight bid-offer spreads, at least on the most actively traded contracts. For instance, bid-offer spreads on five-year unfunded contracts on the CDX.NA.IG index have typically been in the range

[7] In the event of defaults in the index, the protection buyer sells the collateral to recover losses on the CDS index.

0.5–4 basis points. To put the size of this bid-offer differential in context, spreads on the broad investment grade indices have averaged about 62 basis points in North America and 45 basis points in Europe since January 2004 (Graph 2, left-hand panel).[8]

CDS index tranches

Compared to other CDOs, index tranches are standardised and more liquid

CDS index tranches are synthetic collateralised debt obligations (CDOs) based on a CDS index, where each tranche references a different segment of the loss distribution of the underlying CDS index.[9] The main advantage of index tranches relative to other CDOs is that they are standardised. Standardisation applies to both the composition of the reference pool and the structure ("width") of the tranches.

Standardisation helps to foster greater liquidity in the secondary market. The development of a liquid secondary market for the trading of other CDO tranches has thus far been elusive largely because the structure of most CDOs has been highly customised.[10]

CDS index spreads[1]

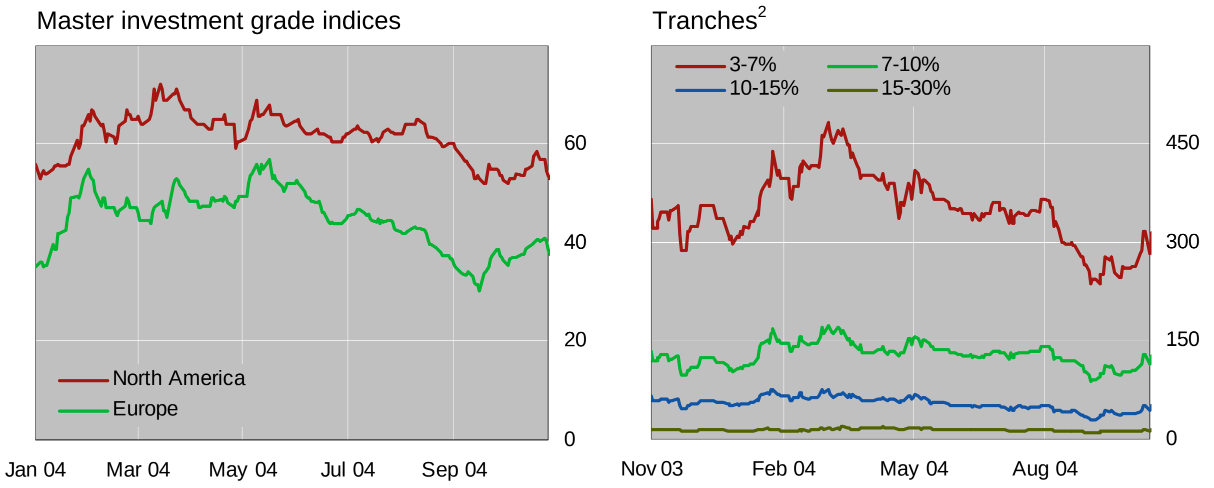

[1] On-the-run five-year swap spreads, in basis points. [2] North America master investment grade.

Source: JPMorgan Chase. Graph 2

---

8 At origination, the fixed spread for the index swap is set to be roughly equal to the average CDS spread for the names in the index. As time progresses, the index swap will have a positive value to the protection buyer when average spreads on individual names are high compared with the fixed rate. In this case, new buyers of protection would make a payment to the protection seller equal to this difference (and vice versa when average spreads are lower than the fixed rate).

9 In general, a CDO is a structured finance product in which the credit risk on a pool of assets is sold to investors. The claims issued against the assets in a CDO are prioritised (structured) in order of seniority, ie there are different levels or "tranches" of debt securities. This typically includes one or more investment grade classes and an equity (first loss) tranche. See CGFS (2005) for more detail on CDOs and their economics, and Gibson (2004) for a discussion of the risks in synthetic CDOs.

10 One of the main growth areas in the CDO market over the past couple of years has been so-called bespoke single-tranche CDOs. These are designed in accordance with a specific investor's wishes. It could be argued that market forces are pushing towards two extremes:

Index tranches target segments of the default loss distribution

Tranches have been issued on several indices, though most trading to date has been concentrated in the CDX.NA.IG index.[11] There are five tranches based on this index. The lowest tranche, known as the equity tranche, absorbs the first 3% of losses on the index due to defaults. If defaults occur over the lifetime of the tranche contract, the investor in an equity tranche is obliged to pay its counterparty an amount equal to the losses from default (the difference between par and the recovery price of the defaulted asset) up to a maximum of 3% of the total index. The next tranche (mezzanine) absorbs losses of 3–7% and is therefore fully insulated, by the equity tranche, from losses up to 3%. Further losses are absorbed by higher-ranking tranches. The 7–10% and 10–15% tranches are known as the senior tranches, while the super-senior tranche covers losses of 15–30%.[12]

In return for bearing the risk of losses, investors receive a quarterly payment from buyers of protection equal to a premium times the effective outstanding notional amount of a given tranche.[13] The premiums on the mezzanine and senior tranches are a running spread with no upfront payment. By contrast, buyers of protection on an equity tranche make an upfront payment that is a percentage of the original notional of the contract, in addition to paying a running spread premium of 500 basis points.[14] The presence of a (relatively large) upfront payment changes the prospective timing of cash flows to the investor in an equity tranche compared to the case of receiving a running spread only, and therefore the equity investor's exposure to the timing of defaults is different. Market quotes of the premiums on the mezzanine and higher tranches are shown in Graph 2 (right-hand panel).[15]

Trading credit risk correlations: pricing the tranches

Credit risk correlations affect the riskiness of index tranches

Credit risk correlations among the names in the index have a large impact on the riskiness of CDS index tranches. The high degree of sensitivity to credit risk correlations is clearly reflected in the pricing of the tranches. This implies

standardised index tranches (which can be used in active trading) and bespoke tranches (which are designed for buy-and-hold purposes).

11 Creditflux reports transactions volume of $10.2 billion in the second quarter of 2004, with 82% of this total referencing iBoxx CDX.NA.IG Series 2 and Trac-x NA combined.

12 Contracts for insuring against losses greater than 30% of the index currently do not exist.

13 The effective notional is the original notional less any losses incurred due to defaults that have impacted on the tranche (with a floor at zero).

14 A contract with an upfront payment can be converted into a contract with a running spread and no upfront payment. This is done by dividing the upfront payment by the (risky) duration of the tranche and adding any running spread. Thus, an equity tranche with an upfront payment of 37.5%, a running spread of 500 basis points and risky duration of 3.75 is equivalent to a contract with a running spread of (37.5*100/3.75) + 500 basis points = 1,500 basis points. See O'Kane and Sen (2003) for an analysis of upfront versus running spread quoting conventions.

15 Bid-ask spreads have been 1–2 basis points for the most senior tranche and 5–10 basis points for the mezzanine tranches, while they have been 15–70 basis points for the equity tranche.

that, in conjunction with the greater liquidity of these instruments relative to other multi-name credit products, these securities offer a relatively efficient way of trading this form of risk.

To illustrate the importance of credit risk correlations on the value of the tranches, consider tranches with a five-year maturity on a CDS index consisting of 125 names whose characteristics are similar to the average credit in CDX.NA.IG Series 3.[16] The left-hand panel of Graph 3 shows the five-year loss rate distribution, as a percentage of tranche size, from the equity to junior mezzanine tranche. The right-hand panel reports the expected loss as a percentage of the total index, on each tranche. This clearly illustrates that loss, both relative and absolute, is declining in tranche seniority. Indeed, the expected loss on the equity tranche is about 40–50% of notional in the cases shown in the graph.

The pricing of index tranches has focused on default time correlations

This example indicates that the market value of a given CDS index tranche will depend upon the joint default loss probability distribution for the reference entities in the index. In general, the joint default loss distribution incorporates both the correlations between individual default probability levels and the correlations between individual default times. In addition, the true loss distribution also incorporates correlations between losses-given-default and default probability levels (eg losses tend to be larger when the overall risk of default is higher, such as in recessions) and correlations between losses-given-default and default times (eg losses may be larger when defaults are clustered, such as when there are multiple defaults in an industry over a short period of time).

Index loss rate distributions and expected losses on index tranches[1]

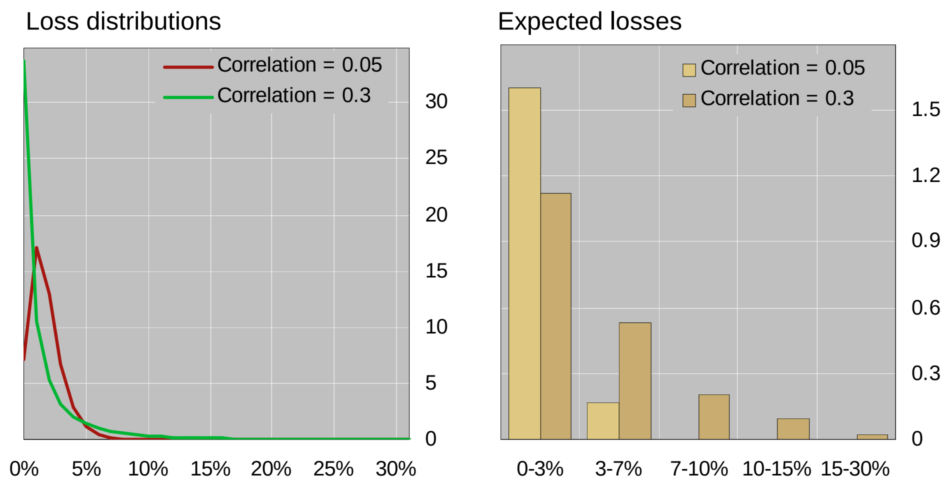

Note: the x-axis is percentage loss (left-hand panel) and tranche (right-hand panel).

[1] In per cent. See text for details on computations.

Source: BIS calculations. Graph 3

[16] To calculate the loss distribution, we use a one-factor Gaussian copula model (see below) and assume identical five-year default probabilities (2.97%), constant recovery rates (40%) and constant identical pairwise default time correlations (0.05 or 0.3). The default rate is estimated using Moody's data for US Baa-rated corporate issuers over the period 1983–2003. The recovery rate is the average for defaulted senior unsecured US corporate bonds. The chosen values of default time correlations are roughly in the range used by the rating agencies.

Pricing index tranches

The premium on an index tranche is the spread paid by the protection buyer that equates the expected present value of default costs to be borne by the protection seller ("protection leg") to the expected present value of investing in the tranche ("premium leg"). The value of the premium leg is the present value of the spread payments the protection seller receives from the protection buyer. Index contracts specify M quarterly payment dates, $t = t_1, t_2, \ldots, t_M$, on which the buyer of protection makes payments to the seller. Note that payments are only made as long as the (uncertain) effective notional of the tranche at time t_i, denoted by $N(t_i)$, is positive. Assume also that investors discount expected future income streams using the (uncertain) discount factors $D(0,t_i)$. Given the tranche premium S, the expected present value of the premium leg is:[①]

$$\mathrm{V_{prem}} = S \cdot E\left[\sum_{i=1}^{M} D(0,t_i) \cdot N(t_i)\right]$$

The expected tranche sizes depend on the number and timing of any future defaults and the expected costs of these future defaults (ie recovery rates).[②] The present value of the premium leg is lower if: the premium is low; the recovery rate is low; and default losses are incurred early. The expected present value of the protection leg is:[③]

$$\mathrm{V_{prot}} = E\left[\sum_{i=1}^{M} D(0,t_i) \cdot \left(N(t_i) - N(t_{i-1})\right)\right]$$

The present value of the protection leg is lower if: the tranche size does not change; the recovery rate is high; and defaults occur late during the contract period. The tranche premium is found by solving $V_{prem} = V_{prot}$ for S:

$$\mathrm{S} = \frac{E\left[\sum_{i=1}^{M} D(0,t_i) \cdot (N(t_i) - N(t_{i-1}))\right]}{E\left[\sum_{i=1}^{M} D(0,t_i) \cdot N(t_i)\right]}$$

Implementation

As can be seen from the equations above, two key factors are required to determine S: future effective tranche sizes and discount factors. Discount factors can be found via methods also used for other financial instruments (see Rebonato (2002)). To evaluate future tranche sizes, however, several inputs are needed: (1) the losses-given-default; (2) the number of defaults; and (3) the timing of defaults. All of these quantities are uncertain, and therefore expectations of them must be formed.

For the loss-given-default (or one minus the recovery rate), a simple approach is to assume that recovery rates are constant and equal to the average historical recovery rate on senior unsecured bonds for US corporations (typically around 40%). Recovery rates can also be estimated from CDS spreads.

Individual default probabilities for the names in the index can be estimated directly from single-name CDS spreads. Alternatively, they can be inferred indirectly from equity prices (eg Moody's KMV's expected default frequencies). Note that a recovery rate assumption is needed to extract default probabilities from CDS spreads.

The timing of defaults for the N entities over the lifetime of the contract can be calculated from a joint default time probability distribution. As this is unknown, a common approach is to assume that default times follow an N-dimensional multivariate normal distribution, ie the so-called Gaussian copula (see Nelsen (1999), Li (2000) and Cherubini et al (2004)).

[①] In practice, when defaults occur between payment dates, sellers of protection receive an accrual payment at the next payment date based on the previous effective tranche size. Note that any upfront payment on the equity tranche can be included in the present value of the premium leg by adding a constant. [②] Expectations are taken under a risk neutral measure, ie risk-adjusted expectations. [③] Assuming protection buyers receive compensation at the next scheduled payment date after a default has occurred.

In a *one-factor* Gaussian copula model, the correlations in default times are assumed to be equal and constant across entities. This is equivalent to assuming that there is a direct mapping from a latent random variable X_i to default times, where the evolution of X_i is given by:

$$X_i = \sqrt{\rho} \cdot M + \sqrt{1-\rho} \cdot Z_i$$

where M is a normally distributed random variable, the Z_i's are mutually uncorrelated and normally distributed random variables and $-1 < \rho < 1$ is the constant pairwise correlation between default times (see Hull and White (2004) for further details). One interpretation of the one-factor Gaussian copula approach is that X_i is the value of assets held by entity *i*, and entity *i* defaults if its assets fall below some threshold. This is similar in spirit to a Merton-type model, where the option to not repay debt is exercised when asset value reaches a given threshold. With this interpretation, M can be seen as the single common risk factor, while the Z_i's are N idiosyncratic risk factors, driving the values of firms' assets, and thus default times. The correlation parameter ρ can be estimated from correlations of equity returns, which are typically in the range 0–30%.

Up until now, the pricing of index tranches has focused on capturing the implications of default time correlations (see box). For this purpose, the so-called one-factor Gaussian copula model has become the market standard for gauging the prices on index tranches, similar to the Black-Scholes model for trading options. The term "copula" is meant to emphasise that this type of model "couples" individual-name default probability distributions together to form a joint default probability distribution (see Nelsen (1999)). The one-factor Gaussian copula assumes identical constant pairwise default time correlations across all firms, normally distributed default times and a normal joint default probability distribution. These simplifying assumptions make the one-factor Gaussian copula relatively easy to use to calculate valuations, which is one of the main reasons for its popularity.

Default time correlations and tranche pricing

Higher correlations imply more default clustering

The importance of default time correlation for the riskiness of the different index tranches is apparent in Graph 3. It is shown in the left-hand panel that, depending on the tranche, the probabilities of having either very small or very large loss rates are higher when default time correlation is higher. This can easily be seen by comparing two extreme, albeit unrealistic, cases.

First, if correlation is zero, the probability of zero names (out of 125) defaulting within a five-year period is $(100 - 2.97)^{125}$ = 2.31%, where 2.97% is the average historical five-year-ahead default rate of Baa-rated firms. By contrast, if correlation is equal to one (ie if the portfolio can be viewed as a single credit), the probability of zero names defaulting is 97.03%. Yet the index could lose one minus the recovery rate (= 1 – 0.4) with probability 2.97%, making the expected loss equal to 1.78%.[17] The right-hand panel of Graph 3 shows that the expected loss on the equity tranche is higher with low

[17] Increasing default time correlation is equivalent to making the default probability random but with the same mean default probability. Note that a mean-preserving distribution of this type implies a higher average joint survival rate due to convexity of the joint survival probability distribution. See Lando (2004) for further discussion.

Price sensitivity of CDX tranches to default time correlation[1]

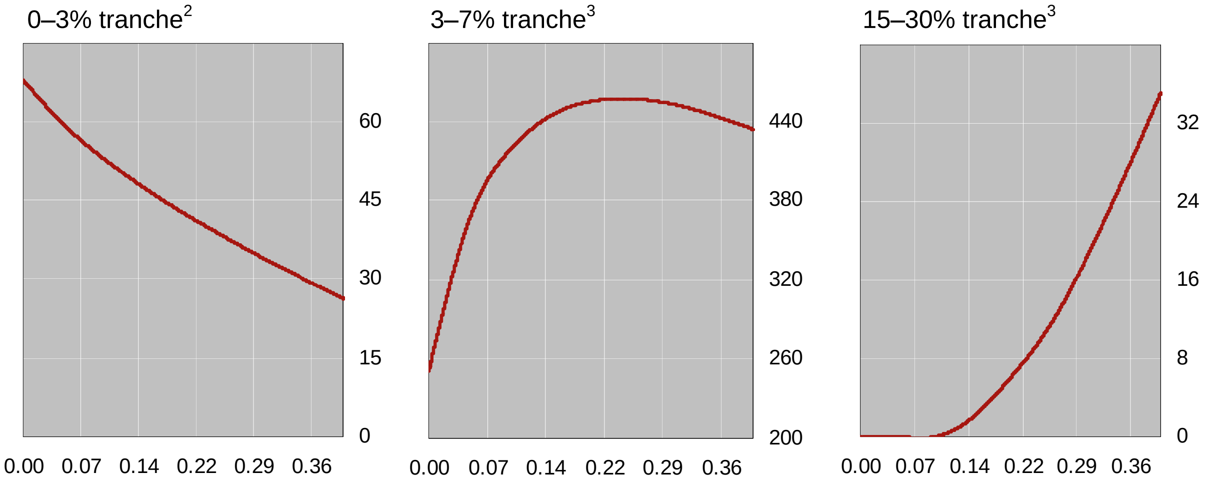

[1] Based on correlation sensitivities reported in Hull and White (2004). [2] Upfront payment, in per cent. [3] Spread, in basis points.

Sources: Hull and White (2004); BIS calculations. Graph 4

correlation. This is not the case for the mezzanine and senior tranches. Indeed, expected losses are higher on the senior tranches when correlation is higher.

As correlation increases, the spread on the equity tranche declines …

As the risk of different tranches varies with default time correlation, so does the pricing of the tranches. This is illustrated in Graph 4, which plots the model-implied upfront payment on the equity tranche and spreads on the mezzanine and super-senior tranches as a function of default time correlation.[18] Consider the equity tranche. More default clustering has little negative impact on the value of this tranche, as only few defaults are needed for this tranche to incur substantial losses. At the same time, a higher default time correlation increases the chance that no defaults will occur. Therefore, the upfront payment on the equity tranche declines as default time correlation increases. By contrast, the pricing of the senior tranche reflects its greater exposure to the risk of losses when defaults are more clustered. Unlike the equity and senior tranches, the price of the mezzanine tranche is generally not a monotonic function of default time correlation. With both high and low correlations, there is a high probability that this tranche will survive intact. However, for medium levels of default time correlation, there is a high risk that the mezzanine tranche will suffer substantial losses.

… whereas it increases on the senior tranche

Market prices and implied default time correlations

Default time correlations can be inferred from market prices

Evidence of the market's view on default time correlations can be inferred from market prices on CDS index tranches. This can be done by specifying a pricing model and all the necessary inputs for the model except the default time correlation. For instance, by specifying values for all of the inputs in the one-factor Gaussian copula model except for the constant pairwise default time correlation, it is possible to back out an implied correlation using market

[18] Tranche prices are based on Hull and White (2004).

quotes.[19] This is illustrated in the left-hand panel of Graph 5, which plots implied default time correlations for the index tranches over time.

Market-implied default time correlations have a smile …

The left-hand panel in Graph 5 illustrates one of the puzzles observed in market quotes: the so-called "correlation smile".[20] The correlation smile illustrates that, when using a one-factor Gaussian copula, market spreads on the mezzanine tranche (typically) imply a lower default time correlation than is implied by the spreads on equity and senior tranches. Thus, the degree of default clustering assumed by the market appears to be higher for the equity and senior tranches. If the one-factor Gaussian model is indeed the correct description of joint default dependence, then the same implied correlation value should be inferred for all tranches.

… and a skew

The right-hand panel in Graph 5 illustrates another implication of market quotes: the so-called "correlation skew". It plots the market-implied base correlation against the upper bound for each tranche. For example, in the case of the CDX.NA.IG index, the base correlation for the 0–10% interval would be defined as the correlation which equates the price of this synthetic first loss tranche to the combined observed market values of the 0–3%, 3–7% and 7–10% tranches. The base correlation can be interpreted, from the perspective of the protection buyer, as the correlation in an insurance contract which pays out up until a given level of losses is reached. The fact that the base correlation curve is upward-sloping, or "skewed", shows that market prices for index tranches imply that default time correlation is increasing with tranche seniority.

Implied and base default time correlations[1]

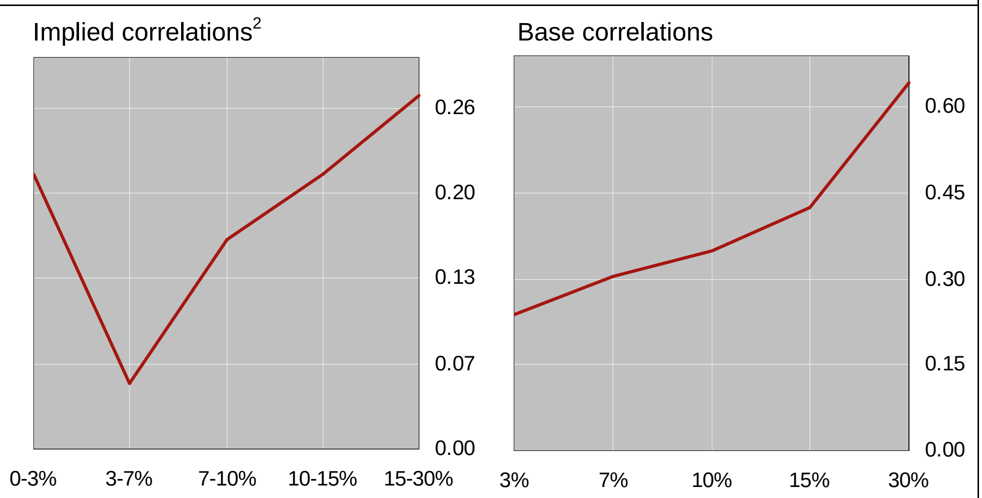

Note: the x-axis is tranche (left-hand panel) and tranche size (right-hand panel).

[1] Based on on-the-run five-year contracts. [2] Averages based on correlation sensitivities reported in Hull and White (2004) and market quotes from JPMorgan Chase spanning the period 13 November 2003 to 28 September 2004.

Sources: Hull and White (2004); JPMorgan Chase; BIS calculations. Graph 5

[19] Index tranches are sometimes quoted in terms of implied correlation instead of spread.

[20] The correlation smile is reminiscent of the volatility smile with respect to strike prices extracted from equity options using the Black-Scholes model.

This reflects the fact that spreads are high on the senior tranches, at least relative to the low level of expected losses on these tranches implied by the model. This is reminiscent of the positive relationship between risk premia and credit quality observed for corporate bonds.[21]

The correlation smile might reflect market segmentation …

There are several possible explanations for the correlation smile (and skew).[22] One is that there is segmentation among investors across tranches and that these different investor groups hold different views about correlations. For instance, the views of sellers of protection on equity tranches (eg hedge funds) may differ from sellers of protection on mezzanine tranches (eg banks and securities firms). However, there is no compelling reason why different investor groups would systematically hold different views about correlations.

… uncertainty about credit risk correlations …

A second possible explanation is that the smile reflects market participants' uncertainty about how best to model credit risk correlations. The implication is that the equity and senior tranches, which are more sensitive to correlations, contain a "model risk" premium embedded in their prices. While this explanation can account for the relatively large premium on the senior tranche, it is not consistent with the relatively low equity tranche premium.

… local demand conditions …

A third explanation is that, even though the index tranche market has grown significantly over the past year, prices might still be subject to local demand conditions. For example, the implied correlation on the mezzanine tranche may reflect strong interest by banks in selling protection on this segment of the index loss distribution. This could be due to the hedging demands of banks, which may be short credit risk of this type as a result of their role as originators of other, notably single-tranche, CDOs.

… or the use of different pricing models

A fourth explanation is that market participants may, in fact, use other models for pricing than the one-factor Gaussian copula. Possibilities include: (1) using fatter-tailed distributions (eg Student's-t); (2) relaxing the restriction of constant pairwise correlations; (3) allowing individual default probabilities to depend on macroeconomic risk factors; and (4) letting recovery rates vary over time and be correlated with default times and default probabilities.[23] For instance, the impact on pricing from using a fatter-tailed distribution, which implies more clustering of defaults, increases break-even spreads for senior tranches and lowers them for junior tranches. Alternatively, a positive correlation between losses-given-default and clustering of default times would lower the price on the most senior tranches for a given level of default time correlation. In this case, the implied correlation inferred from senior tranches (under a constant recovery rate assumption) would be upward biased. This could also explain the pricing of equity tranches, as higher recovery rates during times of little default clustering would imply that this tranche is more valuable.

---

21 For further discussion of this, see Amato and Remolona (2004).

22 See also Bernand et al (2004).

23 The importance of these elements for the modelling of credit risk have been discussed, respectively, by Hull and White (2004), Gregory and Laurent (2004), Duffie and Singleton (2003) and Altman et al (2004).

Looking forward

For CDS index-based markets to mature …

… diversification must increase …

Despite rapid growth, the market for CDS index tranches is still relatively small. Furthermore, even though they have improved diversification opportunities at a lower cost to investors, these instruments still contain significant idiosyncratic risk because they only reference 125 names in five different sectors.[24] However, as these markets continue to mature, the number of underlying names is likely to increase and improve diversification. Thus, in future, index tranches should provide further scope for more efficient trading of credit risk correlations.

… and credit risk modelling should improve …

… to capture more types of credit risk codependencies

To improve market efficiency and limit the risk that exposures are accumulated in ways that are not fully appreciated, it is important for credit risk modelling to develop further. The main challenge appears to be developing frameworks that realistically capture credit risk correlations (see Duffie (2004)). As noted above, the valuation of CDS index tranches has so far mainly focused on modelling the correlation of default times. By contrast, correlations among default probabilities and losses-given-default (ie credit spread correlations), have received less attention. No doubt, progress is being made in developing more general models to capture credit risk codependencies.[25] For instance, some models incorporate contagion effects, which allow them to capture the impact on credit risk from declines in overall market liquidity, the failure of large firms or adverse industry-level developments.[26] Examples of large defaults that have had a market-wide impact include Enron and WorldCom; a recent example of an adverse industry development is the investigation by the New York Attorney General's office into insurance industry practices in the United States. Looking ahead, practitioners, as well as policymakers monitoring these markets, will face the challenge of designing robust models that capture these types of systematic and systemic events.

References

Altman, E I, B Brady, A Resti and A Sironi (2004): "The link between default and recovery rates: theory, empirical evidence and implications", *Journal of Business*, forthcoming.

Amato, J and E Remolona (2004): *The pricing of unexpected credit losses*, Bank for International Settlements, mimeo.

Anson, M, F Fabozzi, M Choudhry and R-R Chen (2003): *Credit derivatives: instruments, applications and pricing*, Wiley Finance.

---

[24] For a discussion of the importance of idiosyncratic risk in credit portfolios, see Amato and Remolona (2004).

[25] These dependencies could also include correlations between discount factors and credit risk.

[26] See Davis and Lo (2001) and Collin-Dufresne et al (2003) for theoretical models of credit risk contagion. Schönbucher and Schubert (2001) show how certain types of more general copulas are able to capture these general credit risk codependencies.

Bernand, A, F Pourmokhtar, B Jacquard, D Baum, L Gibson, L Andersen and J Sidenius (2004): "The Bank of America guide to advanced correlation products", supplement, *Risk* magazine, May.

Cherubini, U, E Luciano and W Vecchiato (2004): *Copula methods in finance*, Wiley, New York.

Collin-Dufresne, P, R Goldstein and J Helwege (2003): *Is credit event risk priced? Modeling contagion via the updating of beliefs*, Carnegie Mellon University, mimeo.

Committee on the Global Financial System (2005): *The role of ratings in structured finance: issues and implications*, Bank for International Settlements, Basel.

Davis, M and V Lo (2001): "Infectious defaults", *Quantitative Finance*, 1, pp 382–87.

Duffie, D (2004): "Time to adapt copula methods for modelling credit risk correlation", *Risk* magazine, April, p 77.

Duffie, D and K J Singleton (2003): *Credit risk: pricing, measurement and management*, Princeton University Press.

Gibson, M (2004): "Understanding the risk of synthetic CDOs", *FEDS Discussion Papers*, no 2004-36, Board of Governors of the Federal Reserve System.

Gregory, J and J-P Laurent (2004): "In the core of correlation", *Risk* magazine, October, pp 87–91.

Hull, J and A White (2004): "Valuation of a CDO and an n-th-to-default CDS without Monte Carlo simulation", *Journal of Derivatives*, forthcoming.

International Swaps and Derivatives Association (2003): "ISDA Credit Derivatives Definitions"*, Supplements and Commentaries*.

Lando, D (2004): *Credit risk modeling: theory and applications*, Princeton University Press.

Li, D (2000): "On default correlation: a copula function approach", *Journal of Fixed Income*, March, pp 43–54.

Nelsen, R (1999): "An introduction to copulas", *Lecture Notes in Statistics*, Springer, Berlin.

O'Kane, D, M Naldi, S Ganapati, A Berd, C Pedersen, L Schloegl and R Mashal (2003): "The Lehman Brothers guide to exotic credit derivatives", supplement, *Risk* magazine, November.

O'Kane, D, C Pedersen and S Turnbull (2003): "The restructuring clause in credit default swap contracts", *Fixed income quantitative credit research*, Lehman Brothers, April.

O'Kane, D and S Sen (2003): "Up-front credit default swaps", *Quantitative Credit Research Quarterly*, Lehman Brothers, Third Quarter.

Packer, F and H Zhu (2005): "Contractual terms and CDS pricing", *BIS Quarterly Review*, March.

Rebonato, R (2002): *Modern pricing of interest-rate derivatives*, Princeton University Press.

Schönbucher, P and D Schubert (2001): *Copula-dependent default risk in intensity models*, Department of Statistics, Bonn University, mimeo.