BANK OF ENGLAND

Working Paper No. 407

Extracting information from structured credit markets

Joseph Noss

December 2010

BANK OF ENGLAND

Working Paper No. 407
Extracting information from structured credit markets

Joseph Noss[(1)]

Abstract

Structured credit instruments offer an insight into markets' perceptions of the extent of future credit defaults. Claims of different seniorities incur losses only if defaults reach different magnitudes, so their relative value offers an insight into the likelihood of losses being of different severities. This paper matches the traded values of structured credit products by modelling the defaults of the underlying credits and their interdependence. It offers an improvement on the industry-standard 'Gaussian copula' model in its ability to capture the 'tail event' of multiple firms defaulting together. This allows policymakers to draw better inference as to the likely scale of defaults implied by structured credit prices. It offers an indication of the extent to which defaults are driven by systemic shocks to firms' balance sheets. It may also be of use to those who trade structured credit products and may offer an improvement in risk management.

Key words: Structured credit instruments, systemic risk, asset prices.

JEL classification: G01, G12.

---

(1) Bank of England. Email: joseph.noss@bankofengland.co.uk

I am grateful to Nicola Anderson, Jens Larsen, Nicholas Vause and Chris Yeates for their help and advice. Any remaining errors are of course my own.

This paper was finalised on 2 November 2010.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/publications/workingpapers/index.htm

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email mapublications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

Contents

Summary

Assessing the stability of an economy frequently involves assessing the risk of bad states of the world materialising. It is often necessary to judge how many firms are likely to default on their debt obligations over a certain time horizon. The likelihood of a large number of firms defaulting is of particular interest to policymakers, particularly if this is caused by some 'systemic shock' that presents a particular threat to financial stability.

Structured credit instruments are created by collecting defaultable assets, such as mortgages or corporate bonds, into portfolios and issuing claims of different seniority against these portfolios. Claims' seniorities determine the order in which they receive cash flows from the underlying assets, with more senior claims being paid first. Their prices therefore reflect market perceptions about the chance of these cash flows materialising, or equivalently, the likely extent of defaults of the underlying credit instruments. While the values of standard credit instruments, such as corporate bonds, offer an insight into the market-perceived probability of a given firm defaulting, the values of structured credit instruments provide a richer view of the likely extent of corporate defaults away from this central case. Claims of different seniorities incur loss only if defaults reach different magnitudes; their relative value therefore affords an insight into the likelihood of losses being of different severities.

Information can be recovered from the prices of structured credit by modelling the default of the different underlying credit instruments and then fitting the resulting modelled prices to those observed in the market. Correctly modelling the distribution of defaults, and in particular their codependence, is crucial in order to find a model whose tranche premia fit those traded in the market. For example, only if a large number of firms default together will senior claims incur loss. Previous attempts to model this interdependence have used a 'Gaussian copula model', based on the Gaussian or normal distribution, to capture the correlation between firms' defaults. However, this gives insufficient weight to the 'tail event' of multiple firms defaulting together.

The framework presented here instead uses a gamma distribution that is more able to capture the possibility of extreme dependence between defaults. It is therefore more successful in matching the traded prices of structured credit products. The model is also extended to include 'catastrophe' and 'becalmed' states that represent the possibility of very high degrees of systemic risk in credit markets, and its reduction perhaps due to government intervention; it therefore offers an intuitive explanation for the large fluctuations in codependence witnessed during the recent credit crisis.

This work offers three key outputs. First, it allows the market-implied probability distribution of firms' defaults to be inferred from the traded value of structured credit instruments. These distributions may be of use to policymakers, particularly because they offer an insight into the risk of 'tail outcomes' involving the default of large numbers of firms. This is likely to be of particular interest to policymakers seeking to measure and mitigate systemic risk. Second, the model offers an insight into the nature and magnitude of the risks firms face. It allows the average probability of a firm defaulting to be decomposed into components relating to default events of different severities. For example, it can estimate how the probability of a particular firm defaulting depends on the likelihood of a very severe event such as widespread financial crisis. Finally, in common with other models of structured credit that go beyond the Gaussian copula, this work is of potential use to those who trade structured credit products. It gives rise to a set of parameters that determine the structure of the codependence of default between credits, which could form the basis of an investor's 'hedging strategy' that allows positions in different tranches to be hedged against each other. This has the potential to protect them from changes in the nature of default codependence that reduce the value of their portfolio.

Introduction

There is a strong tradition of central banks and other policymakers extracting information from the prices of financial securities. Under the efficient market hypothesis,[1] traded prices reflect all publicly available information relating to their future pay-offs. Policymakers can therefore use these prices to back out information about these expected future cash flows. For example, information can be extracted from equity prices on the perceived value of their future dividends.[2] Bond and swap market prices can communicate investors' views on the future path of interest and inflation rates.[3]

Expectations derived from financial asset prices have a number of advantages over other measures of private sector expectations such as surveys. They are available almost continuously, whereas surveys are produced intermittently and take some time to compile. The prices from which they are derived combine the information of a large number of investors. The fact that the expectations are derived from prices in financial markets where actual investments are made means that they are likely to reflect more careful consideration than survey responses. There remains the problem that the price of a financial asset is driven by more than just its expected future pay-offs: some prices are said to include an 'illiquidity premium' that compensates its holder for any difficulty they may encounter in finding a buyer for the asset should they wish to sell it at a future date. Expectations derived in this way will also be those of a 'risk-neutral' investor; in the likely case that investors are averse to risk, the market-implied probability of bad states of the world materialising will be less than under risk-neutral valuation.

Financial derivatives can provide information on the future path of their underlying asset's price that goes beyond its central expectation. Because their pay-offs are contingent on the asset's price, derivatives provide a guide to the relative likelihood of different future price movements, as perceived by investors. Bahra (1996 and 1997) fit a mixture of two log-normal distributions to asset prices and fit the resulting distribution to that implied by options prices. Breeden and Litzenberger (1978) observed that the risk-neutral probability density of an asset's future price is proportional to the second derivative of the price of options written upon it with respect to their strike prices. This has now become a standard technique adopted by central banks for extracting information on the future course of asset prices from options contracts.[4] These methods also

[1] See Fama (1970). The efficient market hypothesis has, however, been criticised in recent years: there is some empirical evidence suggesting that it may not hold true in practice, and debate as to whether some agents' predictions could in fact be systematically wrong.
[2] See Panigirtzoglou and Scammel (2002).
[3] See Anderson and Sleath (2001)
[4] Clews *et al* (2000) gives the implementation used by the Bank of England.

highlight two different approaches to extracting information from derivatives prices. In contrast to that of Bahra, the framework of Breeden and Litzenberger is 'non-parametric' and makes no attempt to model the underlying asset's price. While this has the advantage of being able to fit a large constellation of derivative prices without the constraint of a single model of asset dynamics, it has the disadvantage of offering no information as to the causes of these prices. The presentation offered here is parametric and seeks to model the behaviour of the underlying assets in order to offer some insight into the economic drivers of their prices.

Structured credit instruments are created by collecting defaultable assets, such as mortgages or corporate bonds, into portfolios and issuing multiple 'tranches' that represent claims of different seniority against these portfolios. The relative seniority of a tranche determines the order in which it accrues losses incurred on the underlying credits. If any of the assets in the portfolio default during the life of the product, the resulting losses accrue first to junior tranches and only on to senior tranches if losses reach a sufficient magnitude. An understanding of what drives the relative value of exposures of differing level of seniority enables market practitioners to determine their relative value. This valuation process can also be reversed to infer information on the likely severity of defaults.

Just as equity derivatives can be used to extract information on the relative likelihood of stock price movements, structured credit instruments can be used to extract information on the relative likelihood of the defaults of their underlying credits reaching different magnitudes. Because tranches of different seniorities incur loss only if defaults reach different magnitudes, their relative value therefore affords an insight into the likelihood of losses being of different severities.

The approach taken here is to value structured credit products by modelling the defaults of the underlying credits. These modelled values are then fitted to those observed in the market. The implied values of the model's parameters are then used to make inferences about the likely scale of defaults and their codependence. Correctly modelling the codependence of defaults is crucial in order to find a model whose tranche premia fit those traded in the market. A standard means of modelling this interdependence has been to use a 'Gaussian copula model', based on the Gaussian, or normal, distribution to capture the codependence between firms' defaults. However, this gives insufficient weight to the possibility of multiple firms defaulting together.

The presentation given here builds on a strand of literature that seeks to better capture the dependence structure between defaults of individual firms and consequently matches traded tranche premia better. There has for some time been interest in models of financial asset prices that go beyond the limited scope of the Gaussian distribution in their ability to capture 'tail events' witnessed in markets. Madan and Seneta (1990) model the value of equity derivatives using a gamma distribution, rather than a Gaussian distribution, to capture the possibility of extreme price movements beyond those incorporated in the standard Black-Scholes[5] framework. Cariboni and Schoutens (2007) developed a similar model to price simple credit contracts. Both Albrecher *et al* (2007) and Baxter (2007) suggest using a gamma distribution to improve the pricing of structured credit products and better capture the 'tail risk' of extreme codependence between defaults. This work builds on this literature and offers a model based on a gamma distribution that accounts for the more complex dynamic for codependence of default witnessed during the recent credit crisis. In contrast to some recent literature however, the presentation here offers an intuitive explanation for the economic drivers of systemic risk. In particular, it includes 'catastrophe' and 'becalmed' states that represent the possibility of very high degrees of systemic risk in credit markets, and its reduction perhaps due to government intervention; it therefore offers a natural economic explanation for the large fluctuations in codependence witnessed during the recent crisis.

The model presented here offers an improved calibration to traded prices of structured credit compared to that based on the Gaussian copula. This allows superior inferences to be drawn as to the nature of credit default risk priced into these products and *offers two key outputs of use to policymakers*. First, it allows for the extraction of the market-implied probability distribution of the underlying firms' defaults. These distributions offer insight into the risk of 'tail outcomes' involving the default of large numbers of firms. This is likely to be of particular interest to policymakers seeking to measure and mitigate systemic risk. Second, the relative values of the parameters of the model give insight into the drivers of corporate defaults. An index of structured credit is decomposed into portions relating to different dynamics of the underlying model. This allows the average probability of a firm defaulting, reflected by the value of the index, to be decomposed into components relating to default events of different severities. The model can extract the proportion of default risk arising, for example, from systemic shocks to firms' asset values that cause the default of multiple firms.

[5] See Black and Scholes (1973).

Finally, this framework is of use to those who trade such securities. It gives rise to a set of parameters that give an intuitive explanation for the structure of the codependence of default between credits. These can form the basis of an investor's 'hedging strategy' that allows positions in different tranches to be hedged against each other. In common with that of Masol and Schoutons (2008), this model could offer an improvement in risk management and valuation compared to those based on a Gaussian distribution, because it is able to fit the prices of all traded tranches simultaneously. It also allows non-standard 'bespoke' products to be priced consistently.

This paper proceeds as follows. Section 1 introduces structured credit indices and, in particular, the CDX index to which the models here are applied. Section 2 presents three competing models of structured credit premia, and compares their ability to capture the default codependence implied by the traded prices of structured credit products. It begins with a benchmark model of structured credit along with a description of its shortcomings. This is extended to a framework based on a gamma distribution that is shown to better capture default codependence as implied by market prices. A further refinement deals with the high levels of default codependence observed during the recent financial turmoil, as well as the possibility of state intervention to remove this codependence. A third section presents the results of calibrating this model to the CDX index, along with various measures of systemic risk implied by their prices which may be of interest to policymakers. A fourth section shows how, in common to other models that price structured credit products, this work may be relevant to industry professionals who wish to price new credit exposures and hedge existing ones. A final section concludes.

Technical details are confined to an appendix. This gives details of how losses on each tranche are calculated under the various models, and how these then give rise to a premia on each tranche. An analytical approximation is used in the course of this calculation that radically reduces the computation time compared to alternative techniques. This is important as the subsequent fitting of observed premia in the market requires a number of successive computations of the modelled spreads.

1 Structured credit indices

This study focuses on the Dow Jones CDX five-year North American Investment Grade Index, which is based on a basket of 125 credit default swaps (CDS) that track the cost of insuring against the default of the underlying firms. Each CDS contract functions as an insurance arrangement under which a buyer of credit protection pays a fixed quarterly premium over some given time such as five years. If there is a default on the underlying bond during that period then the buyer of protection may give the defaulted bond to the protection seller and receive the full face value of the bond in return. The index tracks the average CDS premia on the underlying firms, and hence their average level of credit risk.

Figure 1: Tranche structure of the CDX index

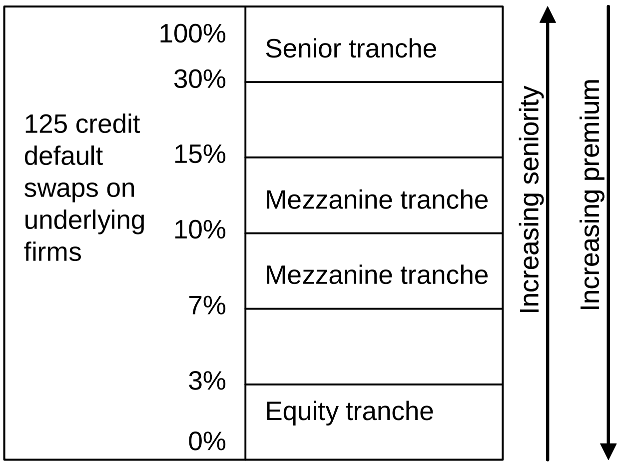

Chart 1: Premia on tranches of CDX index

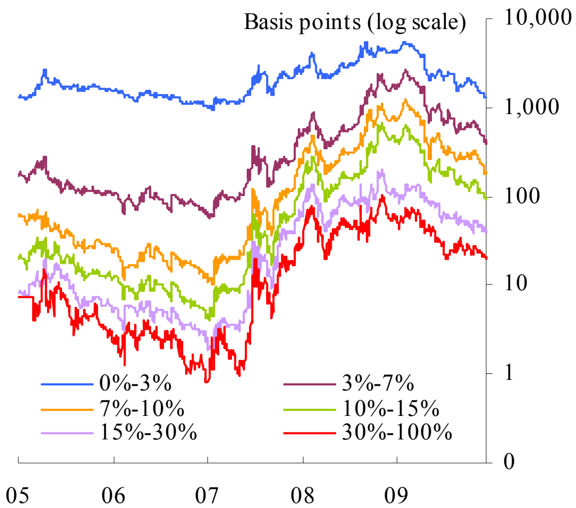

Indices such as the CDX also include a number of tranches, each of which has an 'attachment' point that the proportion of defaults on the underlying firms must exceed before the tranche itself suffers loss. Attachment points on 'junior' tranches are low, meaning that losses accrue to them first; those on the 'senior' tranches are higher, so that they only suffer loss once more junior tranches have been depleted. An investor in a given tranche therefore functions as a seller of credit protection on a 'slice' of the underlying securities, receiving a fixed premium payment, but suffering a loss should defaults exceed a certain amount. Those investors requiring a highrisk investment can invest in junior tranches, to which losses accrue first, but which pay a high premium. Those requiring a lower risk investment can invest in senior tranches, which only suffer loss if large numbers of firms default, and so offer lower premia. In this way, tranches redistribute among investors both the credit risk of the underlying firms and the returns to bearing that risk. Attachment/detachment points of the tranches of the CDX index are set at 0%-3%, 3%-7%, 7%-10%, 10%-15%, 15%-30% and 30%-100%. The tranche structure and its associated premia are shown in Figure 1 and Chart 1.[6] The data are daily mid-quotes between 1 February 2005 and 10 January 2010. The liquidity of the CDX index and its associated tranches has grown rapidly since its inception.

[6] For further details on the CDX index see Rogoff and Ursino (2007).

The compensation demanded by investors to hold different tranches of structured credit offers policymakers an insight into the market's perceptions of the nature of corporate default risk. These tranche premia reflect not only the probability of each underlying asset defaulting, but also market perceptions about the codependence of these defaults - the extent to which large numbers of firms are likely to default together.

While the value of the main index tracks the average probability of default of the underlying firms, the tranche premia reflect how the market expects these losses to be distributed between tranches. Moves in tranche premia not accompanied by a change in the perceived risk of the index as a whole should, in principle, reflect only a change in the market view of the nature of the codependence of default. A higher degree of codependence increases the likelihood of polar outcomes in which either a majority of underlying assets default, or very few default. This reduces the likelihood of intermediate outcomes in which a modest proportion of the underlying credits default. This benefits investors in the equity tranche, because the transfer of probability mass from intermediate outcomes to outcomes with very few defaults increases the chance of the equity tranche being preserved. At the same time, investors in senior tranches become worse off because there is more chance of clustered defaults that could erode some of their tranche. The increased degree of clustering in defaults could be due to defaults being driven by a 'common economic shock' that affects many firms simultaneously.

Table 1: Examples of varying levels of codependence				**Premium for holding:**	
State of the world	A	B	C	**…the equity tranche**	**...the senior tranche**
Probability	1/3	1/3	1/3		
Case 1–Low codependence	Asset X defaults	No defaults	Asset Y defaults	**Higher**	**Lower**
Case 2–High codependence	No defaults	Assets X and Y default	No defaults	**Lower**	**Higher**

This codependence is illustrated in Table 1. Each of three states of the world, A, B and C, occur with equal probability. In the first case, assets default in two of three states of the world - X in state A, Y in state C. In contrast, under case two assets only default in one state of the world, but now both default together. Crucially, *the expected number of defaults in each case are equal*; but in case 1 defaults occur idiosyncratically, whereas in case 2 they occur together.

2 Three competing models of portfolio credit risk

Three models of structured credit premia are examined here, each of which seeks to model the codependence between defaults in order to fit the premia observed on tranches of the CDX index. Each is based on the same 'copula'[7] structure that seeks to capture the dependence structure between the defaults of individual credits, but offer successive improvements in their ability to capture the levels of codependence implied by market prices. The tranche premia given by each model are fitted to those traded in the market on each day separately, and the resulting model parameter values allow us to draw inference on the nature of the codependence between defaults and how it changes over time.

The first model, based on the 'Gaussian copula', became the 'industry-standard' model for pricing structured credit products. However, it is shown to be unable to capture the degree of codependence between defaults implied by market prices. This motivates its comparison with a second model, of a type first proposed by Baxter (2007), which is based on a gamma distribution that enables it to better fit the traded values of tranches. A refinement of this, presented as a third model, offers an additional improvement that is able to fit the fluctuating degrees of codependence witnessed during the recent credit crisis.

2.1 The Gaussian copula

Early attempts to model the value of structured credit products have focused on the Gaussian copula.[8] This represents the value of each firm's assets by a state variable, X_i, which is the sum of two normally distributed random variables; W_g and W_i:

$$X_i = \sqrt{\rho} W_g + \sqrt{1-\rho} W_i . \qquad \textbf{(1)}$$

W_g is a 'global' shock that affects all firms simultaneously; W_i is an idiosyncratic shock affecting only the *i*th firm. The default of the *i*th firm occurs if $X_i \le \theta_i$, that is if the value of the firm's assets, over the life of the contract, breaches some barrier θ_i, typically representing the face value of the firm's liabilities. θ_i is calibrated so that each firm within the portfolio defaults with some probability that, for example, can be determined by its CDS premium, as this reflects its individual probability of default.[9] The value of the correlation parameter, $\rho \in (0,1)$,

[7] For a general introduction to copulas see Nelsen (1999).
[8] See Li (2000).
[9] We require that $\theta_i = \Phi^{-1}\left(P^i{}_{CDS}\right)$, where $P^i{}_{CDS}$ is the i*th* firms probability of default.

determines the extent to which defaults are all driven by the common shock rather than by idiosyncratic shocks, and thus the extent to which their defaults codepend.[10]

Calibrating this model to market prices determines the probability of, and correlation between, defaults on the underlying firms such that the model produces tranche premia that fit those prevailing in the market. This technique became the industry standard for structured credit valuation; see Belsham *et al* (2005) for further details.

The Gaussian copula does however have a significant shortcoming in that it does not accurately capture the nature of the codependence between defaults implied by market prices. No single value of the correlation parameter produces tranche premia that are close to those observed in the market. The Gaussian distribution fails to capture the possibility of very high codependence between defaults implied by the premia at which tranches are traded. This is analogous to how the univariate normal distribution underestimates the probability of extreme movements in the price of a single financial asset; thus no single volatility parameter allows the Black-Scholes model[11] to fit the market prices of options of all strikes.

Chart 2: Base correlations for tranches of the CDX index

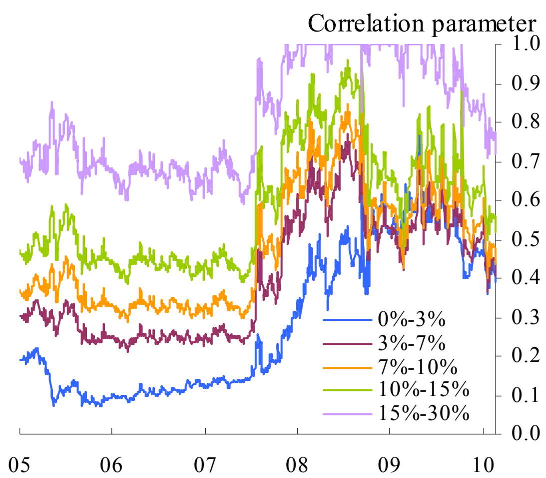

Chart 3: Correlation 'smiles' for the tranches of the CDX index on particular dates

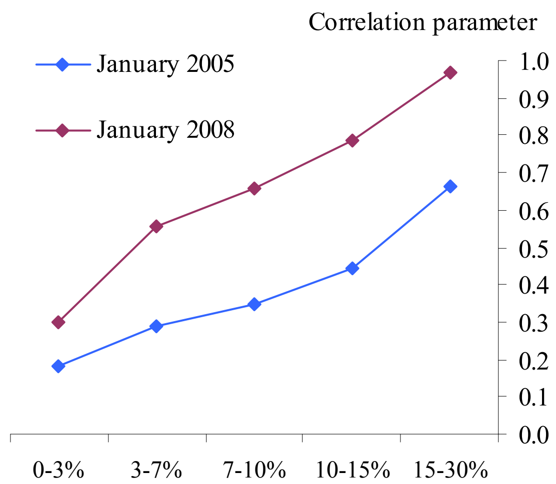

In practice, practitioners have found a way to 'force' the Gaussian copula to fit values observed in the market, through a technique now commonly known as 'base correlation'. This applies different correlations to each tranche, modelling the loss on each tranche as though it existed

[10] The Gaussian copula is therefore a multidimensional extension of a Merton model (see Merton (1974)), which models the default of a single firm.
[11] Black and Scholes (1973).

independently of the other tranches.[12] Base correlations for the tranches of the CDX are shown in Chart 2. Even with correlation set at 100%, the correlation structure provided by the Gaussian model is insufficient to fit premia on senior tranches, as shown by how the top-most line on Chart 2, relating to the 15%-30% tranche, appears 'capped' at 100%.

The way in which base correlation allows the Gaussian copula to be fitted to the tranche market is analogous to how the Black-Scholes model can be made to fit options with different strike prices. With Black-Scholes, different 'implied volatilities' are fitted to options that pay-off if the underlying asset's price moves by different amounts; here, different 'implied base correlations' are fitted to tranches of different seniorities – that is tranches which payout if losses accrue to different amounts. On a given day, 'implied volatilities' arising from the Black-Scholes model can be plotted for options of different strikes, resulting in 'correlation smile' plots. In the same way, Chart 3 shows a plot of implied correlations relating to the Gaussian copula for CDX tranches of a different seniority.

These 'correlation smiles', while betraying the Gaussian model's inadequacies, do however serve as a 'prism' through which to visualise levels of correlation implied by tranche premia and how they compare to those pertaining to the Gaussian model. The way in which the correlation smile moved upwards between January 2005 and 2008 shows how correlation levels implied by the CDX rose. However, because each tranche is a subordinated claim on the same pool of assets, a coherent pricing model ought to price all tranches using the same set of parameters. What is required is a model that correctly captures the dynamics of defaults on the underlying credits and allows all tranches to be priced with a single parameterisation.

2.2 An improvement: the gamma copula model

The Gaussian copula, and the normal distribution on which it is based, fails to capture the degree of codependence between defaults implied by the tranche premia on the CDX index. For each firm's assets, the normally distributed process X_i is far too 'smooth', and fails to capture how downward movements in firms assets can be sudden and jump-like, giving rise to 'tail' outcomes for assets that cause the firm to default.[13] It is this that causes the Gaussian copula to fail to match the premia observed on senior tranches.

[12] The problem is complicated by there being multiple correlations that will fit a given X%-Y% tranche under the Gaussian model. The approach of base correlation is therefore to price a 0%-X% tranche and a 0%-Y% tranche, allowing the 'implied correlation' for the X%-Y% tranche to be extrapolated. For further details see McGinty *et al* (2004).

[13] Jumps in the value of firms' assets could occur, for example, following the unexpected release of price-relevant information.

A better approach, suggested by Baxter (2007), is to model firms' assets through a gamma distribution, arising due to the firm's assets experiencing a number of downward 'jumps'. Firms' assets now face a series of shocks that follow a 'gamma process' governed by two parameters, γ and λ. These determine the intensity and the (inverse) size of jumps respectively. The firm's assets can either experience a series of small shocks with high frequency, or only occasional shocks of a greater size. Chart 4 shows a sample plot of a realisation of this process for two extreme parameterisations. The red line in this chart shows how these shocks might evolve when it has frequent (high γ) small (high λ) jumps, while the blue line shows the opposite case. As before, default occurs if these assets fall to a value below the face value of liabilities. The distribution of asset values arising from the realisation of a gamma process follows a gamma distribution, notated $\Gamma(\gamma,\lambda)$. The probability density function for this is plotted in Chart 5 for different values of γ. As gamma increases this approaches the normal distribution and the case of high intensity, low impact 'normal' shocks akin to those obtained from the normal distribution.

Chart 4: Gamma processes: two extremes

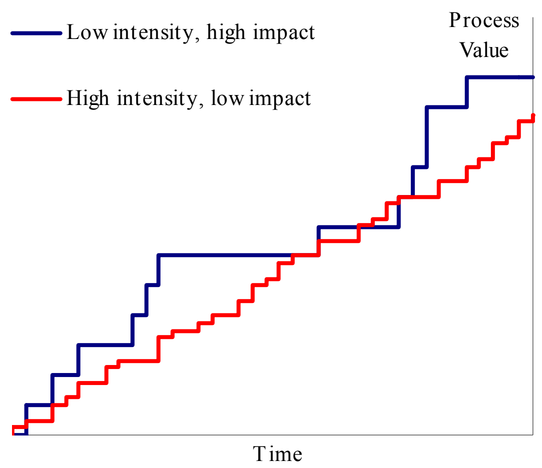

Chart 5: Gamma distributions for different values of gamma

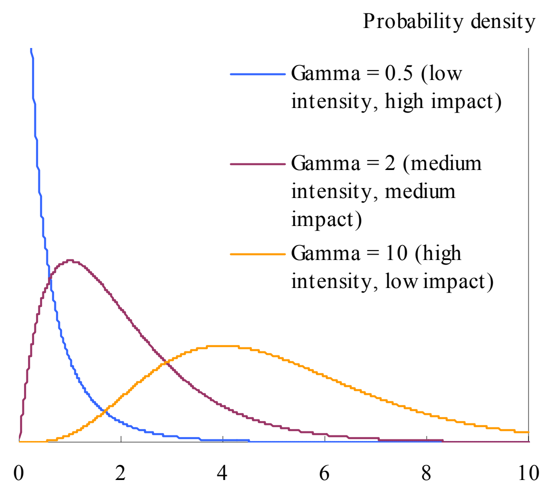

The values of the gamma process's parameters therefore give rise to a spectrum of credit risk models, which capture the nature of a firm's defaults on a scale between two extremes:

- The firm's assets can either evolve smoothly, enduring a series of small but frequent shocks. The cost of credit protection increases slowly as the firm's assets decrease, and default becomes steadily more likely;
- The firm's assets witness a series of infrequent but very severe shocks. In this case default occurs unexpectedly, with no prior increase in the cost of credit protection.

The gamma model therefore has the attractive feature of nesting the Gaussian model as a limiting case as γ tends to infinity, but has added flexibility to fit the constellation of market prices.

Under the gamma framework, the new expression for X_i is:

$$X_i = -\Gamma_g(\phi\gamma, \lambda) - \Gamma_i((1-\phi)\gamma, \lambda). \qquad \textbf{(2)}$$

This is analogous in spirit to equation **(1)**. We again assume that each firm defaults if X_i is below a threshold θ_i, which can be calibrated to each firm's default probability. The shape parameter, γ, is now split between the global and idiosyncratic shocks by parameter ϕ. The properties of the gamma distribution mean that X_i is gamma distributed, and has shape parameter, γ - the sum of those of its two constituents. ϕ is therefore analogous to the correlation parameter in the Gaussian model, and controls the degree of codependence between defaults. The λ parameter is made redundant due to time-scaling.[14]

[14] It is a property of the gamma distribution that if $G \sim \Gamma(\gamma, \lambda)$ then $tG \sim \Gamma(\gamma, \lambda / t)$. Therefore, $P(X_i(t) \le \theta_i) = \Gamma(t\theta_i \mid \gamma, \lambda) = \Gamma(\theta_i \mid \gamma, t\lambda)$, hence lambda merely rescales our choice of default threshold. For further details see Applebaum (2009).

Figure 2: Illustration of the gamma copula

Copula 1

Low correlation, of default (low phi).

Diffusive asset values (high gamma).

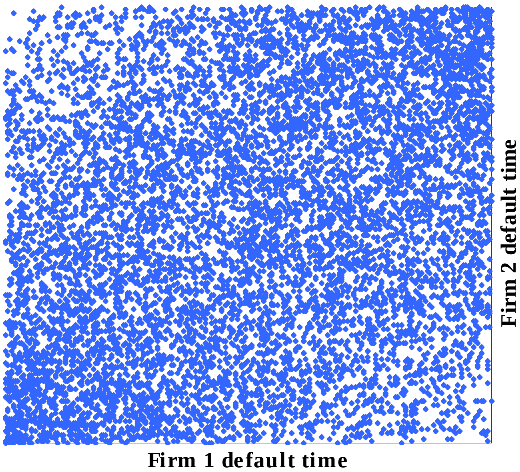

Phi = 0.3; Gamma = 10

Copula 2

High Correlation, of default (high phi).

Diffusive asset values (high gamma).

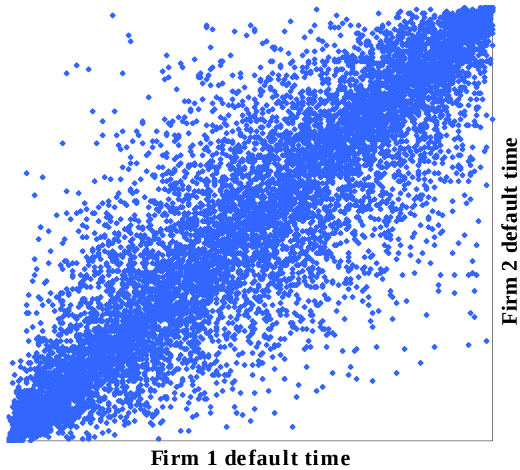

Phi = 0.9; Gamma = 10

Copula 3

Low correlation of default (low phi).

Sudden default (low gamma).

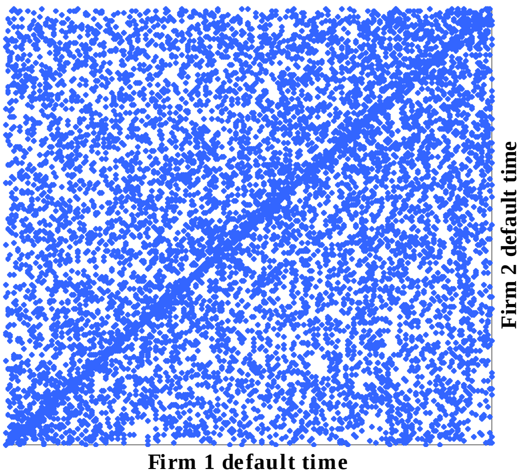

Phi = 0.3; Gamma = 0.1

Copula 4

High correlation of default (high phi).

Sudden default (low gamma).

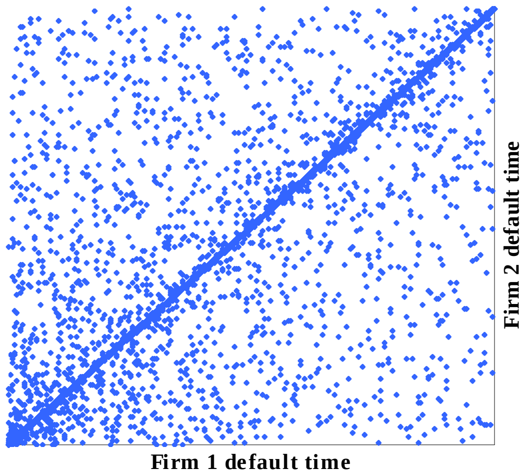

Phi = 0.9; Gamma = 0.1

The gamma copula model therefore has two parameters: ϕ and γ:

- ϕ controls the codependence of default;
- γ controls the propensity of firms to 'jump to default'. Higher values of γ mean the model is more like the Gaussian copula model, where firms' asset values gradually approach a default level. Lower values of gamma mean that firms jump to default more often. It thus determines the 'normalness' of the correlation structure.

The effects of the two parameters are shown in the copula scatter diagrams in Figure 2, which show realisations of the value of individual firms' assets, X_i. Higher values of ϕ (copulas 2 and 4) increase the codependence of firm's asset values, drawing them together in the diagonal of the chart. It is this that increase the correlation of defaults. Altering γ changes the structure of this codependence: high values of γ (copulas 1 and 2) mean the gamma copula approaches that of the Gaussian. Lower values of γ (copulas 3 and 4) make this correlation more extreme: some firms jump to default together, making their defaults highly correlated. Other firms' defaults are less correlated than under the 'high gamma' parameterisation. It is this that causes the tight band of simultaneous defaults down the diagonal of the chart.

2.3 An extension: 'becalmed' and 'catastrophe' regimes

The gamma copula fitted the market well prior to August 2007. But as the credit crisis got under way, senior tranche premia widened to new-found highs that the gamma model was unable to match, as shown in the times series of fit scores in Chart 9. Later, around September 2008, the tranche market moved in strange ways. Premia increased, especially that of senior tranches, which moved to unprecedented highs (Chart 6). Later, equity tranche premia decreased a little. This caused their base correlations, as given by the Gaussian copula, to increase (shifting risk up the tranche structure) and meant that the base correlation curve in January 2009 (Chart 7) adopted a 'smile' shape.

The movements can be explained in terms of risks faced by the underlying firms. The rise in the senior tranche premium reflected perceptions that a catastrophic credit event could occur, which would cause a large number of underlying firms to default. This would increase the chance of losses accruing to the senior tranche. The tightening of equity tranche premium could be attributed to government intervention creating the possibility of a 'becalmed' state where intervention mitigated system-wide risks to firms, and prevented them defaulting for systemic reasons. The analogy to the Black-Scholes 'volatility smile' is again helpful: it is as though the market began to see high levels of codependence as an 'out-of-the-money option' that was less likely to materialise. This is analogous to how options that pay off in extremely 'bad' states of the world have high implied volatilities.

Chart 6: Senior (30%-100%) tranche premia of the CDX index

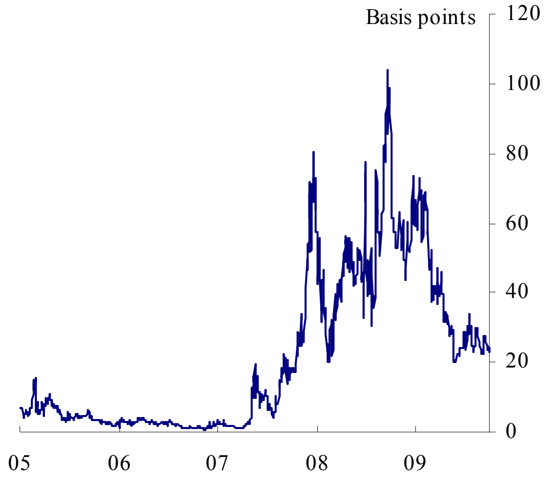

Chart 7: Correlation 'smiles' for the CDX index on particular dates

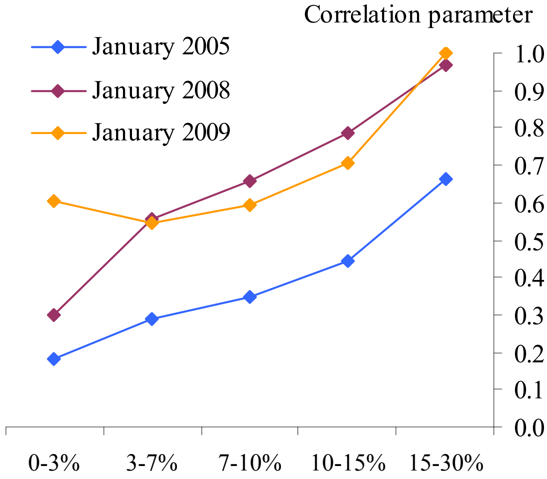

The onset of crisis motivates a variant of the model to cope with such 'catastrophe' and 'becalmed' regimes. This includes provision for extreme comovement of defaults that go beyond those provided by the gamma model. We superimpose a 'switching regime' over the process X_i that controls the default of each firm, which now moves between three states in the course of the life of each CDS contract:

1. **Becalmed state:** At the contract's inception, systemic risk is mitigated perhaps by successful government intervention. Any defaults that occur are entirely idiosyncratic; all codependence between defaults is removed.

2. **Gamma state:** At some point in the life of the contract, government support is withdrawn. The model enters the standard gamma model described in Section 2.2, and firms default for both systemic and idiosyncratic reasons.

3. **Catastrophe state:** At some later date, there is a major systemic shock that causes the default of every firm. In this state, either all or no firms default. Codependence is total.

It is important to appreciate that the occurrence of the becalmed and catastrophe states does not alter the probability of any single firm defaulting. As in the earlier example in Table 1, this remains constant (and determined exogenously) across all models. The additional states merely serve to minimise/maximise the codependence between defaults if they occur.

Our equation for X_i, which controls the default of each firm, now has additional terms compared to that in **(2)**:

$$X_i = \underbrace{-\left(-kI_{\{benignState\}} + kI_{\{catastropheState\}} + \Gamma_g(\phi\gamma, \lambda)I_{\{gammaState\}}\right)}_{globalFactor} - \Gamma_i((1-\phi)\gamma, \lambda). \qquad \textbf{(3)}$$

As both the catastrophe and the becalmed states affect all firms simultaneously, the new terms, which quantify the effects of the becalmed and catastrophe states, are added to the global factor. The becalmed/catastrophe states subtract/add a (arbitrary) large constant k on to the value of the global factor, reflecting the way in which they remove/maximise codependence of the defaults. This adds extra mass in the tails of the distribution of the global factor, in addition to that provided by the gamma distribution over the Gaussian. This density function is illustrated in Chart 8.

Chart 8: The probability density function of the global factor, with becalmed and catastrophe regimes

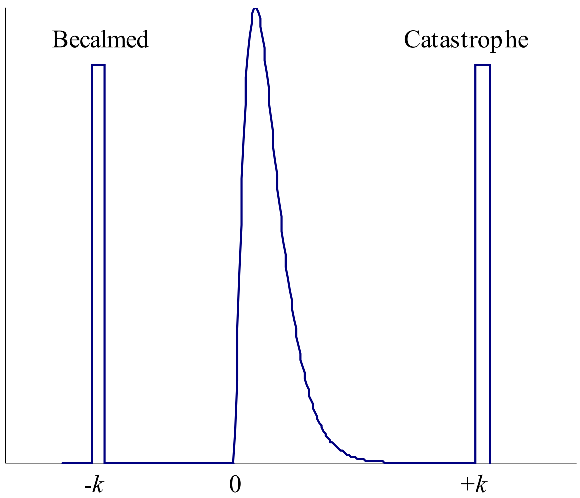

The occurrence of the becalmed and catastrophe states is controlled by two 'Poisson processes'.[15] These consist of jumps at unpredictable occurrence times, t_i, $i = 1,2,\ldots$. The jump times are assumed to be independent of one another, and each jump is assumed to be of the same size. During some small time interval, the probability of observing more than one jump is infinitesimal. The total number of jumps observed up to time t is a Poisson counting process, denoted N_t. The probability that during a finite interval of time Δ there will be n jumps is given by:

$$P(\Delta N_t = n) = \frac{e^{-\lambda\Delta}(\lambda\Delta)^n}{n!}, \qquad \textbf{(4)}$$

[15] For further details and rigorous definitions see Cox and Miller (1965).

where λ is the intensity parameter that controls the expected number of jumps per unit time.

The model here adopts two independent processes C_t and B_t which control the occurrence of the catastrophe and becalmed states respectively. t is the time in years from the CDS contracts' inception. At time t = 0, both processes take value zero. The becalmed state occurs for as long as B_t = 0, and hence, given **(4)**, with probability $P_b(t) = e^{-bt}$. The catastrophe state occurs when $C_t > 0$, which occurs with probability $P_c(t) = 1 - e^{-ct}$, and $B_t > 0$, which occurs with probability $1 - P_b(t) = 1 - e^{-bt}$. The occurrence of the 'becalmed' state and the extension of state support rules out the occurrence of a catastrophic credit event; the occurrence of a catastrophic event is conditional on the exit from the becalmed state. These state probabilities, and their effect on the value of the global factor, are summarised in Table 2.

Table 2: Catastrophe and becalmed states

	Becalmed state	**No becalmed state**
Catastrophe	[Ruled out] **P(this state occurring) = 0**	Catastrophe state **P(this state occurring) = (1-$P_b(t)$)$P_c(t)$** Global factor = +k
No catastrophe	Becalmed state **P(this state occurring) = $P_b(t)$** Global factor = -k	Standard gamma model **P(this state occurring)= (1-$P_b(t)$)(1-$P_c(t)$)** Global factor = $\Gamma_g(t;\phi\gamma,\lambda)$

The becalmed-catastrophe extension to the model has the effect of introducing two new parameters, b and c, which are the intensities of the Poisson processes that control the exit from the becalmed state, and entry into the catastrophe state. Higher values of b increase the propensity of firms to exit from the becalmed state earlier, reducing the probability of government support applying at a given time t. Higher values of c increase the propensity of firms to enter the catastrophe state, increasing the probability of the catastrophe state applying.

3 Calibration

In this section, the Gaussian and gamma models, along with the becalmed-catastrophe extension, are fitted to the traded premia of the tranches of the CDX index on different days. The models' parameters are optimised to achieve the best fit over the observed values. Two objective functions are used: root mean squared error, $\sqrt{\frac{1}{n}\sum_{tr\neq equity}\left(Mkt_{tr} - Model_{tr}\right)^2}$, and mean proportional error, $\frac{1}{n}\sum_{tr\neq equity}\frac{|Mkt_{tr} - Model_{tr}|}{Mkt_{tr}}$, where both sums are taken over the n tranches of the capital structure. Mkt_{tr} are the tranche premia observed in the market on tranche tr. The probability of default of the underlying firms is computed from the value of the main CDX index, as this is the average of their individual CDS premia – the cost of insuring their debt.[16]

Calibrations based on different objective functions serve different purposes. Root mean squared error prioritises minimising the absolute error on the junior tranches, as it is these that have larger premia. They are useful to gauge the models' success at modelling tail risk inherent in senior tranches, as insufficient tail risk immediately shows up in the model giving too small a premium on the senior tranches. Proportionate errors apportion the fitting error more equally across tranches, and thus give equal weight to the information provided by each tranche. They are, therefore, used to extract information from the entirety of the tranche structure.

The modelled tranche premia, $Model_{tr}$, can be calculated using a variety of methods, the most simple of which is Monte Carlo simulation. This involves the repeated simulation of underlying defaults allowing the loss incurred by each tranche to be sampled. While this is straightforward to implement, it is computationally expensive to run – a particular problem here as we intend to iterate over different sets of parameters to find that which produces tranche premia that best fit those observed in the market.

We therefore prefer the 'conditional normal' approach of Li and Liang (2005), who, through an application of the central limit theorem, generates an analytical approximation for the loss on each tranche conditional on the value of the 'global' factor. Intuitively, the central limit theorem gives that the sum of independent random variables, which in this case determine the binary default/non-default of each firm, approaches a normal distribution as the number of summands approaches infinity. However, an important property of both the Gaussian and the gamma copula models is that defaults are *not* independent of each other: indeed, their mutual dependence is ensured by the presence of the global factor which drives the default of multiple firms. However defaults are independent *conditional upon the value of this global factor*, as the remaining source of variation lies in the remaining 'idiosyncratic factor' that is independent between firms. This conditional independence allows for the application of the central limit theorem and yields an analytical solution for the loss on each tranche, conditional on the value of the global factor. All that remains is to integrate this over the distribution of values of the global factor. This last step is undertaken numerically, but the method is far more computationally efficient than that of Monte Carlo. Further details are confined to the appendix.

[16] Extracting firms' probability of default from their CDS premium is fairly straightforward. See, for example, Hull (2007).

The remainder of this section present the results in three different forms. First, the tranche premia given by the different models on a given day are compared, allowing the models' relative success at capturing traded premia to be evaluated. Second, time series of the models' parameters offer an insight into how the nature of default codependence, as implied by the CDX index, changes over time. Finally, the models allow probability density functions of the likely extent of future defaults on firms underlying the index to be extracted from the traded tranche premia on a given day. This offers a means of assessing the likely severity of future defaults on the underlying firms.

3.1 Model comparison

The results of calibrating the gamma and the Gaussian copula models on a single day, using the root mean square error objective function, are shown in Table 3. This offers a view of how well the two models fit premia on the mezzanine and senior tranches, when the 3%-7% tranches are fitted well. The equity tranche is excluded from the calibration, but the modelled values for this tranche are included. This facilitates comparison of the becalmed state, which improves the fit on this tranche.

Table 3: Modelled and actual (market) premia for the CDX index on 2 January 2009 (fitted via root mean squared errors)

Tranche (%)	**Market**	Gaussian model	Gamma model	Gamma model, with catastrophe extension	Gamma model, with becalmed-catastrophe extension
0-3	4108	8521	1302	2170	3948
3-7	1515	1473	1436	1443	1529
7-10	718	650	724	721	384
10-15	439	100	210	477	397.6
15-30	110.5	8.6	22	88	78.6
30-100	57.5	0.01	4.99	47	52.5
Fit score		**4428.27**	**2818.32**	**1939.87**	**374.31**

The improved ability of the gamma model to fit the correlation structure of the tranches is clear. The Gaussian model (without using separate base correlations) does not come close to fitting tranches more senior than the 3%-7% and 7%-10%, such is its inability to capture the 'tail correlations' embodied in these. In contrast, the gamma model offers a considerable improvement in terms of goodness-of-fit over the Gaussian model, with a radically lower fit score.[17] The fit to the mezzanine tranches is vastly improved, though the modelled super-senior premia remains small compared to that observed in the market. This indicates that although the

[17] The 'fit score' – a measure of the goodness of fit obtained on a given date, is the value of the objective function for the optimal parameter values.

gamma model and its ability to capture 'jump-to-default' risk goes some way towards fitting prices better, it cannot account for the exceedingly high premia on senior tranches. The catastrophe extension goes far further in enabling the model to fit the premia on the senior tranche.

Time series of average percentage errors across all tranches for the four models are shown in Chart 9. These confirm that the successive improvements to the Gaussian model are progressively more successful in capturing the nature of the codependence of defaults observed in the market. The gamma model's inability to capture the movements in tranche premia from mid-2007 can be clearly seen by the tick up in its fit scores. Similarly, the catastrophe extension alone is unable to capture the movements in the tranche premia from late 2008.

3.2 *Time series of parameters*

Time series of the models' parameters are presented in two forms: both as their raw values, and as a decomposition of the CDX index, which relates to the total credit risk of its underlying firms, into components relating to the global and idiosyncratic drivers of the model.

Chart 9: Time series of fit scores for the successive models calibrated to the CDX index

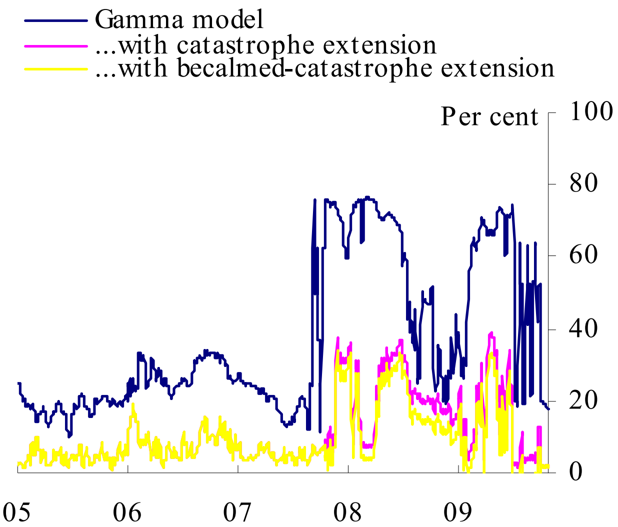

<u>The parameters</u>

The parameters of the becalmed-catastrophe model are plotted in Chart 10. Before the credit turmoil of 2007, ϕ had slowly been decreasing and γ increasing, indicating a decrease in both correlation and jump-to-default risk. Conditions in credit markets were perceived to be increasingly becalmed: defaults were largely driven by idiosyncratic shocks and firms' assets were evolving relatively smoothly without sharp jumps. The implied probability of the catastrophe state occurring was both small and decreasing. That of the becalmed state was negligible.

The initial phases of the credit crisis in late 2007 saw a rapid increase in ϕ and a sharp decrease in γ. The supply of credit to firms contracted and market worries increased regarding both the codependence of default and the risk of 'nasty surprises' in the form of jumps to default.

Early in 2008, market worries about systemic risk increased.[18] The catastrophe probability rose sharply, and correlation continues to increase. In mid-2008 there is a partial reversion as the US Federal Reserve supported Bear Stearns, and both decrease a little, consistent with a perceived reduction in systemic risk. At this point the probability of the becalmed state ticks up, consistent perhaps with an increased perceived likelihood of state support.

In the autumn of 2008 the second phase of the crisis began, with the failure of two major US banks. The probability of catastrophe reached new-found highs, as did the perceived level of government support, reflected in the becalmed probability. γ remained low, reflecting the chance of sudden bouts of bad news causing firms to jump to default.

<u>A decomposition of traded index</u>

All of these models allow the average risk of firms' default, reflected in the value of the main CDX index, to be decomposed into systemic and idiosyncratic components. The proportion of systemic risk is proxied by the value of ϕ, the correlation parameter, which controls the relative importance of the global shock. The remainder of the index premium is attributed to the risk of idiosyncratic default. This is perhaps a more intuitive presentation than presenting the parameter values in isolation, as it focuses on the *relative* importance of different drivers of default risk, as the overall default risk varies over time. With the inclusion of the catastrophe extension, the probability of the catastrophe state occurring can be plotted as a proportion of the index, giving an indication of the proportion of the risk of the index accounted for by the risk of a catastrophic credit event. The proportion of the global shock attributable to jump-to-default risk, through the γ parameter, is also given.[19]

Chart 11 shows this decomposition, both in terms of the absolute values, and as proportions. The latter is given on a logarithmic scale, so that the (small) proportion relating to the catastrophe state is visible. The effect of the becalmed state occurring can be represented by computing what the size of the systemic component would be in the counterfactual world where government support did not exist, and so the probability of the becalmed state occurring were zero. This is shown by the blue dotted line.[20]

[18] See, for example, *Bank of England Financial Stability Report*, May 2008.
[19] Recall that a large value of γ corresponds to low jump-to-default risk, and *vice versa*. The proportion of the global factor attributable to jump-to-default risk is calculated by scaling the gamma parameter to $e^{-\gamma}$.
[20] In the case where government support is extended, which occurs with probability *Pb*, the proportion of spread driven by the global factor would be zero, as all systemic risk is removed. In the case where government support is not extended, which occurs with

Chart 10: Parameter values for the CDX index

Phi (correlation):

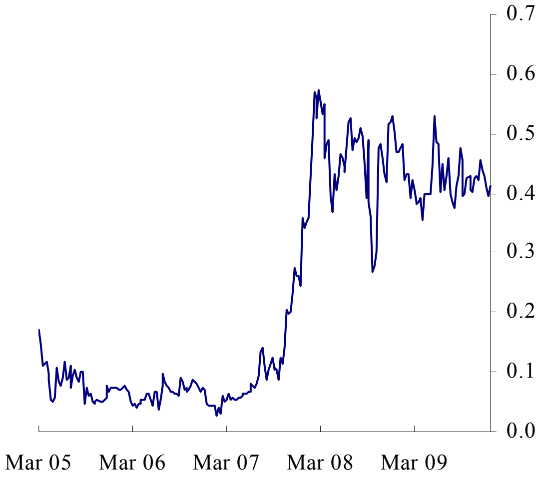

Gamma (jump-to-default risk):
(lower values of gamma gives greater jump-to-default risk)

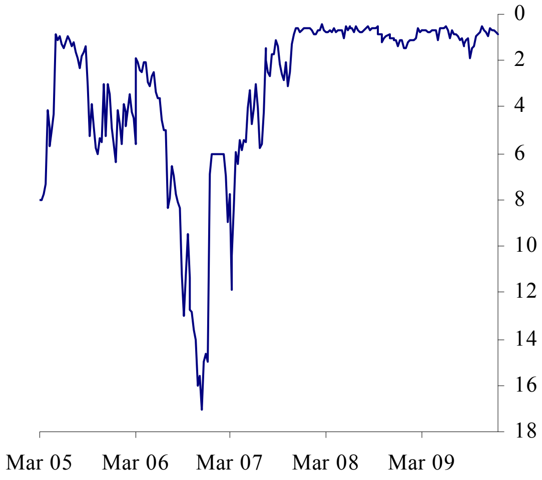

Probability of catastrophe state:

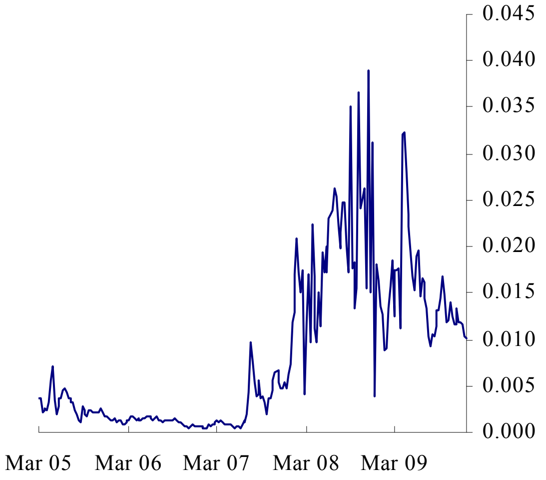

Probability of becalmed state:

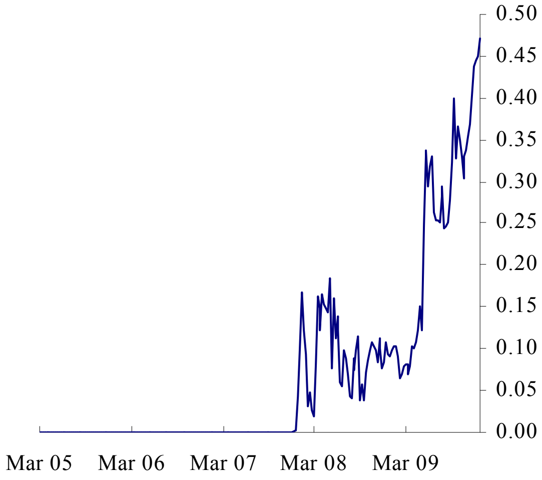

---

probability (1-*Pb*), the proportion driven by the global factor takes some unobserved 'counterfactual' value, which we seek to recover. A little algebra reveals that this unobserved proportion of the global factor takes value of the 'observed' global factor, equal to (ϕ * index spread) / (1 – *Pb*).

Chart 11: Decomposition of the CDX index: absolute values

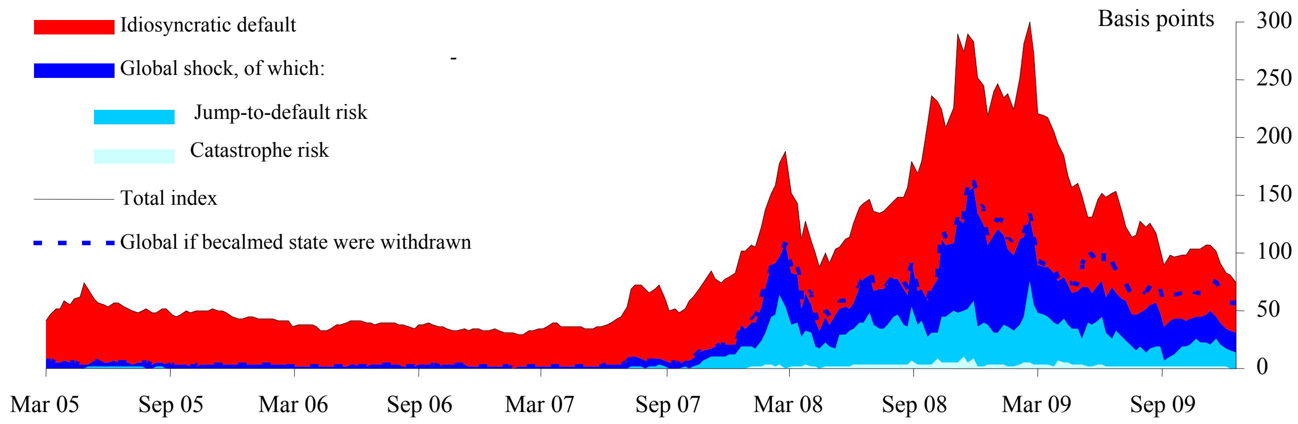

Decomposition of the CDX index: proportionally

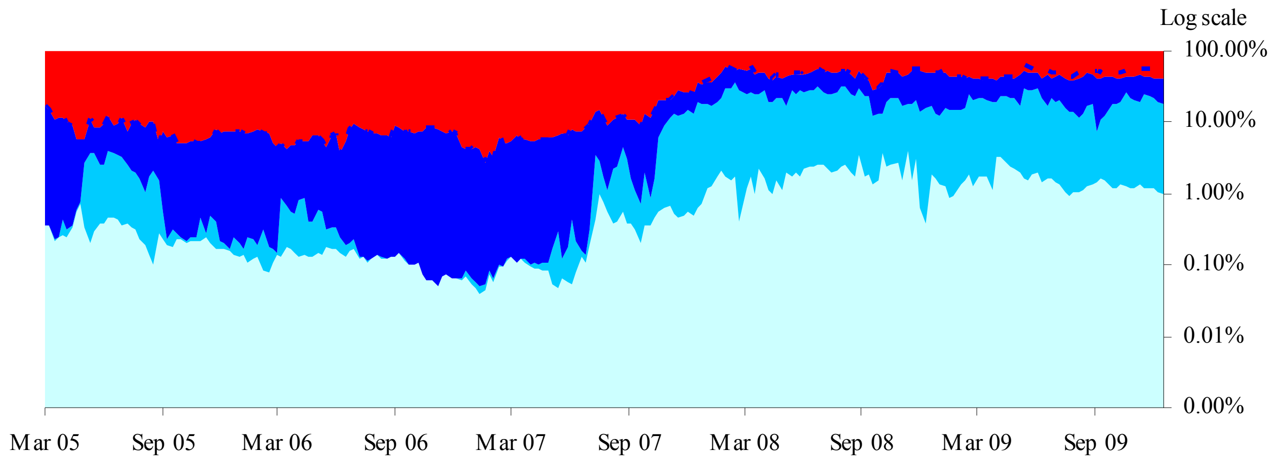

Several interesting conclusions arise from this presentation of the drivers of default as a proportion of the index. The increase in the systemic component, both at the onset of the crisis in 2007, and its intensification in late 2008, is clear (in absolute values). The proportion of jump-to-default risk, though this increased substantially during the crisis, has remained broadly constant throughout it. The impact of the becalmed state on the level of systemic risk increases substantially in 2009, as a range of government policy measures were announced to support financial markets.

3.3 Probability density functions

The model can be used to extract probability distributions for losses on the portfolio of firms underlying the CDX index. These provide a useful means both for assessing the severity of corporate defaults implied by the model and how this changes over time, and also for gaining insight into the performance of each model.[21] Chart 12 shows probability densities for the CDX index extracted under both the Gaussian and the gamma copulas – one for a pre-crisis day and one for during the crisis. For both models, the bulk of the loss density moves to the right after the onset of the crisis, indicating an increase in loss severity.

The probability density functions also offer insight into the performance of the two models. Those produced by the gamma model, indicate a more polar outcome for defaults than those of from the Gaussian, with the mode of the distribution being centred on a lower number of defaults, and a 'spike' in the tail indicating the possibility of a high number of defaults. It is this that improves the fit of senior tranche premia. In contrast the Gaussian model indicates a more intermediate outcome, with a slightly higher mode, but a thinner tail. The gamma distribution's extra degree of freedom allows superior inferences from tranche premia to be drawn more accurately.

Chart 12: Probability densities of defaults on firms underlying the CDX index:[(a)]

Pre-crisis (1 June 2007) **Mid-crisis (1 April 2009)**

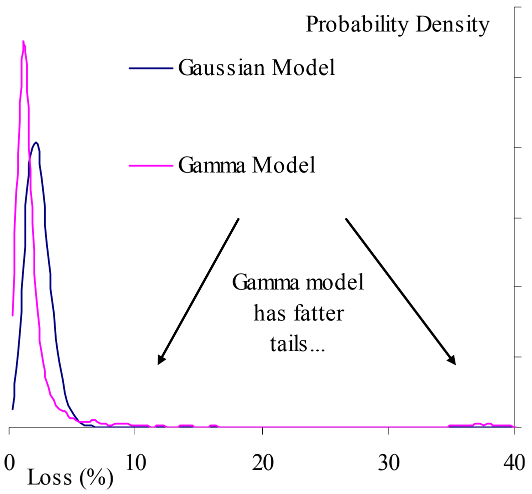

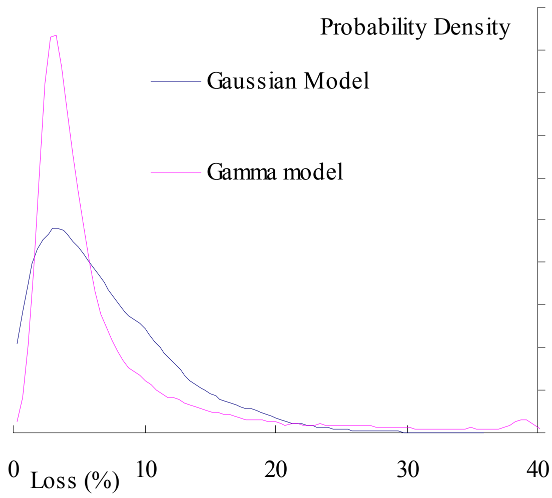

(a) These are 'risk-neutral' probability densities. In the likely case that investors are averse to risk, the perceived probability of high loss rates would be lower than under the risk-neutral measure.

Were densities to be extracted under the model with the addition of the 'becalmed/catastrophe' states, 'lumps' of mass would appear at zero and total losses, representing the realisation of the becalmed and catastrophe states. While these enable the model to accurately price the tranches, they expose a weakness of the framework. If government support consisted of assistance only to selected firms that pose a material risk to the stability of the system as a whole, some firms may still default in the becalmed state. Similarly, catastrophe might therefore not result in the default

[21] For an example of the use of probability densities of defaults in a policy context see *Bank of England Financial Stability Report*, December 2009, page 9.

of all names. The market values of only a limited number of tranches do not hold enough information to allow these subtleties to be deciphered. The resulting simplification (of zero or total default) allows prices to be fitted accurately, but may present problematic density functions.

4 Valuation and risk management

Having a coherent model of default codependence that fits tranches coherently has a number of potential benefits for those that trade structured credit products. It could facilitate better risk management of portfolios of structured credit. Dealers of structured credit run so-called 'correlation books' comprising positions in single-name CDS contracts and credit indices such as the CDX. In managing the risk of their aggregate portfolios, dealers attempt to control their overall exposure to possible market movements, including changes in credit spreads, and movements in the correlation structure.

The development of structured credit indices has given rise to 'correlation trading', where investors take views on the future direction of the codependence of credit defaults. By calculating the 'delta' of each tranche – that is how its premium changes as the value of the main index changes – investors can take a position in the index that offsets the credit risk embodied in the tranche. This allows trades to be constructed which are, in principle, unaffected by movements in the overall credit risk of the underlying firms and exposed only to changes in codependence. These 'deltas' can be determined under the Gaussian copula, given an assumed base correlation for each tranche.

However, base correlation does not allow dealers to risk-manage their exposure to correlation of default – because a different correlation is assumed for each tranche. To hedge against the risk of shifts in the nature of correlation it is necessary to have a coherent model for correlation across all tranches, such as the gamma model.

By modelling the correlation of the underlying names and prices all tranches simultaneously, the gamma model and its variants may allow correlation risk to be better managed. Monitoring parameter values that fit all tranche prices simultaneously and tracking how they evolve over time could provide some useful risk management information. Prior to the auto crisis of mid-2005,[22] many investors sold protection on the equity tranche and bought that on the mezzanine tranche. By delta hedging under the Gaussian copula, investors aimed to be

[22] See *Bank of England Quarterly Bulletin*, Autumn 2005, pages 313-16.

unaffected by changes in the credit risk of the underlying firms, but aimed to profit from increasing levels of correlation that would cause the premia on the equity tranche to decrease and that on the mezzanine tranche to widen. However, this position was vulnerable to a decrease in gamma, which would produce the reverse effect. As several firms 'jumped to default' together, the market witnessed a swift reduction in gamma that widened equity tranche premia and narrowed mezzanine. The gamma model may have informed investors *a priori* that their position was exposed to a reduction in gamma, and so might have provided some useful risk management information that could go beyond that made possible under the Gaussian copula.

Similarly, prior to mid-2007, many investors reportedly thought that senior tranche premia offered good value, in that they seemed to be trading at levels of implied codependence that, at the time, were unthinkably high, and were impossible to calibrate using the Gaussian copula. The catastrophe gamma model would have warned that these premia were consistent with a high-impact systemic shock, whose probability of occurrence had declined since 2005, but which was nonetheless non-zero.

The ability of the gamma copula model to fit tranche premia coherently allows investors to carry out parameter-based hedging to protect themselves against such movements in the correlation structure. This is impossible under the Gaussian copula, which offers no coherent 'view' of correlation to fit the entire tranche structure. Table 5 shows the sensitivity of each tranche to a basis point movement in each parameter. These allow different parts of the tranche structure to be hedged against each other, by allowing a holder of one tranche to take an offsetting position in another tranche in order for their position to be neutral to changes in that parameter.

Table 5: Sensitivity of tranche premia to a 1% increase in the parameters of the becalmed-catastrophe gamma model [(a)]

Tranche	ϕ	γ	Catastrophe	Becalmed
0%-3%	-100.092	-0.385	-3.611	38.236
3%-7%	-24.886	-2.106	-8.742	-3.041
7%-10%	-8.001	0.856	-3.482	-2.484
10%-15%	3.120	1.267	-0.406	-0.293
15%-30%	1.258	0.366	-0.311	-0.071
30%-100%	0.6109	-0.169	0.746	-0.018

(a) Basis-point increase in the premia of each tranche in response to a 1% increase in the value of each parameter.

The sensitivity analysis in Table 5 also offers an insight into the dynamics of the model. The ϕ parameter shifts risk from the junior tranches to the senior, decreasing and increasing their premia respectively. Increasing ϕ therefore increases the overall level of codependence, and can be thought of as producing a rise in the base correlation curve (shown in Chart 3). An increase in the γ parameter moves risk from the junior and senior tranches towards the mezzanine tranches. It reduces the level of jump-to-default risk (recall that *decreasing* γ increases jump-to-default risk) and makes the correlation structure 'more normal' and akin to that of the Gaussian copula. γ can be thought of as controlling the steepness of the base correlation curve (Chart 7) under the Gaussian model, with an infinite value of gamma corresponding to an entirely normal correlation structure found in the Gaussian copula (in the absence of the becalmed-catastrophe extension), and a flat correlation curve. As might be expected, the catastrophe and becalmed parameters increase the premia on the senior and junior tranches respectively.

Bespoke pricing

The gamma model also simplifies the pricing of non-standard 'bespoke' structured credit products, which typically consist of a certain tranche of exposure to a unique basket of credits that differ in their attachment points to those of mainstream indices such as the CDX. The pricing of these under the Gaussian copula is problematic as it is impossible to determine the correlation parameter to apply to this tranche, with its unique attachments. Correlations are quoted and provided only for traded tranches of main indices. Under a coherent model that fits the entire tranche structure with a single set of parameters, such as those examined here, the consistent pricing of bespoke products becomes mechanical. The new attachment points can be inputted into the model and a value of the bespoke tranche immediately results.

There remains the need to calibrate the parameter values to tranches of a traded index, such as the CDX, whose traded composition matches that of the bespoke product. There may be no traded index with comparable underlying firms, in which case a combination of parameters obtained from different indices could be used. But uncertainty may be reduced to the extent that all that is required is to find the right set of parameters.

5 Conclusion

Structured credit products can provide policymakers with information on the extent of defaults on their underlying credits, as perceived by the market. The relative value of their tranches allows inferences to be drawn on the nature and extent of codependence between defaults, and offers an insight into the 'tail risk' of multiple firms defaulting simultaneously. The presentation offered here aims to capture this codependence between defaults and subsequently match the traded premia of the tranches of structured credit products. It offers a refinement of the current 'industry-standard' Gaussian copula model that cannot capture the degree of codependence of defaults implied by the market values of structured credit. In doing so it offers an intuitive economic explanation for the movements in the prices of structured credit instruments witnessed during the recent crisis.

This work has three broad outputs. First, it allows policymakers to draw inferences as to the market-implied probability distributions of the scale of defaults. Second, the parameters of the model themselves offer an insight into the drivers of underlying defaults, and in particular the degree to which defaults are driven by systemic shocks to firms' balance sheets. Their evolution over time explains the varying nature of the codependence of defaults perceived by market participants. Finally, models of this sort may be of use to those who trade structured credit products, in that they may facilitate improved risk management.

Technical appendix

For a given set of parameters (ϕ in the case of the Gaussian model; ϕ and γ in the case of the gamma model; ϕ, γ, *b*, and *c* in the case of the becalmed-catastrophe gamma) we seek a method of computing the corresponding premia on each tranche of the CDX index. With this method in hand, it is a simple matter to then iterate around different parameter values to find those that best fit the prices on a given day.

As outlined in Section 3, the approach used here is as follows:

1. Estimate the probability of an individual firm defaulting, conditional on the value of the global factor.
2. Derive an analytical approximation for the loss on each tranche, conditional on the value of the global factor.
3. Integrating this over possible values of the global factor then yields the unconditional loss on each tranche, and hence its premia.
4. We repeat steps 1 to 3 for different parameter values in order to find those that best fit the value observed in the market.

Only step 1 differs between the Gaussian and gamma models. The remainder of the procedure is identical.

<u>Probabilities of default, conditional on the global factor</u>

Recall from Section 2.1 that under the Gaussian model

$$X_i = \sqrt{\rho}W_g + \sqrt{1-\rho}W_i . \qquad \textbf{(4)}$$

Firm *i* defaults by time *t* with some exogenous probability $p_i(t)$ (calculated from its CDS premium) by setting it to default if $X_i = \sqrt{\rho}W_g + \sqrt{1-\rho}W_i < \theta(t)_i$ **(5)**

where $\theta_i(t) = \Phi^{-1}(p_i(t))$. **(6)**

The probability of default, conditional on this global factor <u>W_g</u>, is then

$$p_i(t \mid W_g) = P(default\_by\_t \mid W_g) = \Phi\left(\frac{\theta_i(t) - \sqrt{\rho}W_g}{\sqrt{1-\rho}}\right) = \Phi\left(\frac{\Phi^{-1}(p_i(t)) - \sqrt{\rho}W_g}{\sqrt{1-\rho}}\right). \qquad \textbf{(7)}$$

The case of the gamma model is very similar. The weighted sum of random variables is now

$$X_i = -G_g^{\phi\gamma,\lambda} - G_i^{(1-\phi)\gamma,\lambda}, \qquad \textbf{(8)}$$

where $G^{a,b}$ is a gamma-distributed random variable with intensity *a* and shape parameter *b*. The summation property of the intensity parameter gives that X_i is also a gamma-distributed random variable, with parameters γ and λ. It follows that

$$p_i(t \mid G_g) = P(default\_by\_t \mid G_g) = P\left(X_i < -\Gamma^{\gamma^{-1}}(1-p_i(t))\right) = P\left(-G_i < G_g - \Gamma^{\gamma^{-1}}(1-p_i(t))\right)$$

$$= 1 - \Gamma^{(1-\phi)\gamma}\left(\Gamma^{\gamma^{-1}}(1-p_i(t)) - G_g\right). \qquad \textbf{(9)}$$

The calculation of conditional expected losses

On default, each loan is assumed to cause a loss of proportion $f_i(1-r_i)$ to the total portfolio of CDS contracts, where f_i is the fraction of the total portfolio that the loan constitutes, (1/125 the case of the CDX), and r_i is the recovery rate, assumed to be 40%.

The random loss caused by the default of loan *I* as $L_i = f_i(1-r_i)I_i$ where $I_i = I_{\{loan\_i\_defaulted\}}$, which takes the value 1 if the firm defaults, and 0 otherwise. This random loss at time *t* are Bernoulli random variables, which are independent *conditional on the value of the global factor*. Its mean and variance are given by:

$$E\left(L_i(\mathrm{t}) \mid W_g\right) = f_i(1-r_i)p^i(\mathrm{t}) \qquad \textbf{(10)}$$

and

$$Var\left(L_i(\mathrm{t}) \mid W_g\right) = f_i^2(1-r_i)^2 p^i(\mathrm{t})\left(1-p^i(\mathrm{t})\right).$$

$$\ldots\text{as } p^i(t) = P(default\_by\_t \mid G_g). \qquad \textbf{(11)}$$

By the central limit theorem, the conditional distribution of the portfolio loss *L* at *t* can be approximated by a normal distribution with mean

$$M_t = E\left(L(\mathrm{t}) \mid W_g\right) = \sum_i E\left(L_i(\mathrm{t}) \mid W_g\right) = \sum_i f_i(1-r_i)p_i(t \mid W_g) \qquad \textbf{(12)}$$

and variance

$$V_t = Var\left(L(t) \mid W_g\right) = \sum_i Var\left(L_i(t) \mid W_g\right) = \sum_i f_i^2(1-r_i)^2 p_i(t \mid W_g)\left(1-p_i(t \mid W_g)\right). \qquad \textbf{(13)}$$

The loss on a tranche with attachment/detachment points *a* and *b*, *T*(*a*,*b*), is given by:

$$L^{T(a,d)}(t) = \max\left(0, L(t)-a\right) - \max\left(0, L(t)-d\right)$$

It follows that

$$E\left(L^{T(a,d)}(t)\mid W_g\right)=\int_a^{\infty}(x-a)\phi\left(\frac{x-M_t}{V_t}\right)dx-\int_d^{\infty}(x-d)\phi\left(\frac{x-M_t}{V_t}\right)dx\,,$$

which, after a little algebra, yields:

$$E\left(L^{T(a,d)}(t)\mid W_g\right)=(M_t-a)\Phi\left(\frac{M_t-a}{V_t}\right)-(M_t-d)\Phi\left(\frac{M_t-d}{V_t}\right)+\sqrt{V_t}\phi\left(\frac{M_t-a}{V_t}\right)-\sqrt{V_t}\phi\left(\frac{M_t-d}{V_t}\right)$$

. **(14)**

Tranche premia

Integration over all possible values of the global factor then yields the unconditional loss on the portfolio. In the case of the Gaussian copula, this global factor is Gaussian distributed, and

$$E\left(L^{T(a,b)}(t)\right)=\int E\left(L^{T(a,b)}(t)\mid W_g\right)dW_g\,. \quad \textbf{(15)}$$

In the case of the gamma copula this integral is with respect to the gamma distribution. In both cases it is computed numerically.

Finally, the premia on each tranche is given by the ratio of the discounted expected loss on the tranche to its discounted expected remaining:

$$premium_{(a,b)}=\frac{\sum_{\tau_i}e^{-r\tau_i}\left\{E\left(L^{T(a,b)}(\tau_i)\right)-E\left(L^{T(a,b)}(\tau_i^{-})\right)\right\}}{\sum_{\tau_i}e^{-r\tau_i}\left\{(d-a)I_{\left\{E\left(L^{T(a,b)}(\tau_i)\right)\in(0,a\right\}}+(d-E\left(L^{T(a,b)}(\tau_i)\right))I_{\left\{E\left(L^{T(a,b)}(\tau_i)\right)\in(a,d\right\}}\right\}} \quad \textbf{(16)}$$

where both sums are taken over a set of times, $\{\tau_i\}_{i=1}^{N}$,when premium payments are made.

Iterate to fit the market tranche premium

Finally we repeat this procedure for different parameter values in order to find those which produce modelled premia that best fit those observed in the market, subject to a fitting criterion. We use a direct search 'simplex' style algorithm.

Becalmed/catastrophe extensions

The inclusion of the becalmed/catastrophe states presents no significant further challenge. All that changes is the distribution of the global factor, W_g, whose cumulative probability density is now given by:

$$F(x)=\begin{cases} p_m & x\in(-k,0) \\ p_m+(1-p_m-p_c)\Gamma(x) \quad \text{for} & x\in[0,k) \\ 1 & x=k \end{cases}$$

where k is arbitrarily large, as in equation **(4)**.

This must be accounted for in the integral in equation **(15)**, and in the calculation of the inverse cumulative density function for the W_g used in equations **(9)** and **(6)** to determine the default barrier. All other calculations remain as before.

References

Albrecher, H, Ladoucette, S and Schoutens, W (2007), 'A generic one-factor Lévy model for pricing synthetic CDOs', *Advances in Mathematical Finance*, pages 259-78, Birkhaüser, Boston.

Anderson, N and Sleath, J (2001), 'New estimates of the UK real and nominal yield curves', *Bank of England Working Paper No. 126*.

Applebaum, D (2009), *Levy processes and stochastic calculus*, Cambridge University Press.

Bahra, B (1996), 'Probability distributions of future asset prices implied by option prices', *Bank of England Quarterly Bulletin*, Vol. 36, pages 299–311.

Bahra, B (1997), 'Implied risk-neutral probability density functions from option prices: theory and application', *Bank of England Working Paper No. 66*.

Baxter, M (2007), 'Gamma process dynamic modelling of credit', *Risk*, October, pages 98-101.

Belsham, T, Vause, N and Wells, S (2005), 'Credit correlation: interpretation and risks', *Bank of England Financial Stability Review*, December, pages 103-15.

Black, F and Scholes, M (1973), 'The pricing of options and corporate liabilities', *Journal of Political Economy,* Vol. 81, pages 637–54.

Breeden, D and Litzenberger, R (1978), 'Prices of state-contingent claims implicit in option prices', *Journal of Business*, Vol. 51, pages 621-51.

Cariboni, J and Schoutens, W (2007), 'Pricing credit default swaps under levy models', *Journal of Computational Finance*, Vol. 10, pages 1-21.

Clews, R, Panigirtzoglou, N and Proudman, J (2000), 'Recent developments in extracting information from options markets', *Bank of England Quarterly Bulletin*, February, pages 50-60.

Cox, D R and Miller, H D (1965), *The theory of stochastic processes*, Chapman & Hall.

Fama, E (1970), 'Efficient capital markets: a review of theory and empirical work', *Journal of Finance,* Vol. 25, pages 383–417.

Hull, J (2007), *Options, futures and other derivatives,* New Delhi: Prentice-Hall of India.

Li, D (2000), 'On default correlation: a copula function approach', *Journal of Fixed Income*, Vol. 9, pages 43–54.

Li, D and Liang, M (2005), 'CDO2 pricing using gaussian mixture model with transformation of loss distribution', Barclays Capital, Quantitative Analytics, Global Credit Derivatives, working paper.

Madan, D and Seneta, E (1990), 'The variance gamma (VG) model for share market returns', *The Journal of Business*, Vol. 63, pages 511-24.

Masol, V and Schoutons, W (2008) 'Comparing some alternative Levy base correlation models for pricing and hedging CDO tranches', *Technical Report 08-01,* Section of Statitsics, K U Leuven.

McGinty, L, Beinstein, E, Ahluwalia, R and Watts, M (2004), 'Introducing base correlations', JPMorgan Chase and Co.

Merton, R (1974), 'On the pricing of corporate debt: the risk structure of interest rates', *Journal of Finance*, Vol. 29, pages 449–70.

Nelsen, R (1999), *An introduction to copulas*, New York: Springer.

Panigirtzoglou, N and Scammel, R (2002), 'Analysts' earnings forecasts and equity valuations', *Bank of England Quarterly Bulletin*, Spring, pages 59-66.

Rogoff, B and Ursino, L (2007), 'The ABCs of CDX', Lehman Brothers, Fixed Income Research, Credit Strategy.