Semi-Analytical Valuation of Basket Credit Derivatives in Intensity-Based Models

Allan Mortensen*

This version: January 13, 2006

Abstract

This paper presents a semi-analytical valuation method for basket credit derivatives in a flexible intensity-based model. Default intensities are modeled as correlated affine jump-diffusions. An empirical application documents that the model fits market prices of benchmark basket credit derivatives reasonably well, consistent with the observed correlation skew. Hence, I argue, contrary to comments in the literature, that intensity-based portfolio credit risk models can be both tractable and capable of generating realistic levels of default correlation.

*Goldman Sachs International, 133 Fleet Street, London EC4A 2BB, allan.mortensen@gs.com

1 Introduction

In recent years, the credit derivatives market has undergone a massive growth in terms of products as well as trading volume. The market is divided into single-name products and multi-name products, also referred to as *basket credit derivatives*, in which the cash flows are determined by the observed defaults and losses in an underlying pool of reference entities. The most traded single-name credit derivative is a credit default swap (CDS), whereas common basket credit derivatives include collateralized debt obligations (CDOs) and basket default swaps such as first-to-default and n'th-to-default swaps. In the pricing of basket products, the marginal default risks of the underlying names are usually implicitly given from the observed single-name CDS market quotes. Default correlation among the underlying names is therefore essentially the only remaining unobservable element in the valuation of basket credit derivatives, which are for that reason also referred to as correlation products.

The standard practice for the pricing and hedging of correlation products is currently the *copula approach*, which is a very convenient way of modeling default time correlation given the marginal default probabilities.[1] Factor copulas have become extremely popular in practice mainly due to their flexibility and tractability (efficient semi-analytical methods have been developed by Gregory and Laurent [2003], Andersen, Sidenius and Basu [2003], and Hull and White [2004]). The copula approach is however problematic for a couple of reasons. First, the choice of copula and the parameters of the chosen copula are usually difficult to interpret. The dependence structure is often exogenously imposed without a theoretical justification,[2] and results are very sensitive to the chosen copula family and parameters. Second, as argued in e.g. Duffie [2004], the standard copula approach does not offer a stochastic model for correlated credit spreads, which is important for realism and indeed necessary for valuing options on CDOs and CDS portfolios.

Given these problems, a natural alternative to copulas is a multivariate version of the *intensity-based models*,[3] where default is defined as the first jump of a pure jump process with a certain default intensity or hazard rate. Intensity-based credit risk models were introduced by Jarrow and Turnbull [1995], Lando [1994,1998], Schönbucher [1998] and Duffie and Singleton [1999] and have proven very useful in the single-name credit markets.[4] It is, however, often argued that intensity-based models are inappropriate for portfolio credit risk modeling. First, the range of default correlations that can be generated in these models is limited due to the fact that defaults occur independently conditional on the intensity processes (although as mentioned in Schönbucher [2003] this is less of a problem when jumps are included and when large pools are considered). Second, it is sometimes stated (e.g. in Schönbucher and Schubert [2001] and Hull and White [2004]) that intensity-based portfolio models are intractable and too time-consuming.

This paper, however, presents a semi-analytical valuation method in a multivariate intensity-based model, showing that these models can be both tractable and able to generate realistic correlations. Default intensities are modeled as correlated affine jump-diffusions decomposed into common and idiosyncratic parts. The intensity model is similar to the model proposed by Duffie and Gârleanu [2001] but this paper offers two extensions. First, a more flexible specification of the default intensities is proposed, allowing credit quality and correlation to be chosen independently. Second, heterogeneous default probabilities are allowed in the semi-analytical solution of this paper, whereas the analytical method of Duffie and Gârleanu [2001] requires homogeneity. This is important in practice since single-name CDS spreads often vary by several hundred basis points within the benchmark credit indices.

The semi-analytical solution is derived in two steps. In the first step, the distribution of the common factor is obtained by inversion of the characteristic function using fast Fourier transform (FFT) methods along with the powerful results of Duffie, Pan and Singleton [2000]. The characteristic function of an integrated affine process requires a small extension of their results, which is proved in the appendix. In the second step, heterogeneous default probabilities are handled using the recursive algorithm of Andersen, Sidenius and Basu [2003]. Semi-analytical valuation in affine jump-diffusion intensity models has previously been proposed by Gregory and Laurent [2003], although the complications of the common factor distribution in the first step are not addressed in that paper. Their method relies on an additional Fourier transform in the second step instead of the recursion used in this paper.

An empirical application documents that the intensity model fits the market prices of benchmark basket credit derivatives reasonably well. While the attainable levels of default correlation in an intensity model are limited, the model is in fact able to produce correlations consistent with the market-implied levels, at least for the large CDO pools considered in this study. The results also show that jumps are needed in the common component to obtain realistic correlations. In addition to the level, the model also fits the shape of correlations. That is, the model is able to generate pricing patterns consistent with the correlation skews observed in the standard model, which is the Gaussian copula. Skew-consistent pricing has previously been reported by Andersen and Sidenius [2004-5] in a Gaussian copula with random factor loadings and by Hull and White [2004] in a t copula. A comparison shows that the ability of the intensity model to match CDO market prices is at least as good as that of the copula models implemented in this paper.

Relative to the copula approach, the intensity-based model has the advantage that the parameters have economic interpretations and can be estimated, for example from CDS market data. The jump and correlation parameter values needed to generate the results of this paper are relatively high. The mere fact that we can discuss whether the parameters are reasonable or not testifies to the model's interpretability, which may also be useful for forming opinions on the absolute pricing levels of correlation products. Furthermore the model, by nature, delivers stochastic credit spreads and is therefore capable of pricing options on single-name CDSs, CDS indices and CDO tranches. If anything, these advantages come at the cost of additional computation time. Although semi-analytical, the model is slower than the fastest copula models but not necessarily so for some of the less tractable copulas, e.g. the t copula of Hull and White [2004] in which the marginal default probabilities are not known in closed form.

The remainder of the paper is organized as follows. Section 2 describes the intensity-based model as well as the semi-analytical solution. For comparison, the copula framework is outlined in Section 3. Section 4 presents the empirical application to basket credit derivatives pricing, and Section 5 concludes.

2 The intensity-based model

This section introduces the multivariate intensity-based model and describes the semi-analytical valuation method, but first some notation and a brief outline of the single-name intensity-based framework.

We consider an underlying pool of N equally-weighted entities over a finite time horizon T. For each entity $i = 1, \ldots, N$, τ_i denotes the time of default and $D_i(t) = \mathbf{1}_{\{\tau_i \leqslant t\}}$ is the default indicator up to time $t \in [0, T]$. A constant homogeneous recovery rate, $\delta \in [0, 1)$, is assumed throughout this paper, and consequently the focus is on the number of defaults, which is denoted by $D(t) = \sum_{i=1}^{N} D_i(t)$.[5] Throughout the paper, everything is done under the probability measure $\mathbb{Q}$. For the pricing analysis in Section 4, $\mathbb{Q}$ is assumed to be the risk-neutral measure.

2.1 The single-name setting

In intensity-based credit risk models, default is defined as the first jump of a pure jump process, and it is assumed that the jump process is represented by an intensity process. This implies that the probability of default occurring over a small interval of time (in the limiting sense) is proportional to the (default) intensity,

$$\lim_{\Delta \to 0} Pr\big(\tau \leqslant t + \Delta \big| \tau > t\big) = \lambda_t \Delta$$

Intensity-based models were first studied by Jarrow and Turnbull [1995] using constant default intensities, in which case default is the first jump of a Poisson process. With stochastic intensities, introduced by Lando [1994,1998], default is the first jump of a so-called Cox process or doubly stochastic process. Conditional on the path of the stochastic intensity, a Cox process is an inhomogeneous Poisson process. Hence, the conditional default probabilities are given by

$$Pr\big(\tau \leqslant t \big| \{\lambda_s\}_{0 \leqslant s \leqslant t}\big) = 1 - e^{-\int_0^t \lambda_s ds}$$

and the unconditional default probabilities by

$$Pr\big(\tau \leqslant t\big) = 1 - E\big[e^{-\int_0^t \lambda_s ds}\big]$$

As can be seen from the last expression, the relationship between survival probability and default intensity corresponds to the one between zero-coupon bond and short rate. Therefore, credit risk modeling in this framework is mathematically equivalent to interest rate modeling, and many well-known and useful techniques from this field can be applied for different specifications of the default intensity. A particularly tractable and flexible specification is the class of affine jump-diffusions, characterized and analyzed by Duffie, Pan and Singleton [2000]. This is the class of processes applied in the following.

2.2 The multi-name intensity model

In the multi-name model it is assumed that default of entity i is modeled as the first jump of a Cox process with a default intensity composed of a common and an idiosyncratic component in the following way:

$$\lambda_i(t) = a_i Y(t) + X_i(t) \tag{1}$$

where $a_i > 0$ is a constant and Y and X_i are independent affine processes specified below. As mentioned in Duffie and Gârleanu [2001], multiple sectors – interpreted as industries or geographic regions – could be incorporated in this setting as multiple common factors. The default intensity in (1) is a very simple modification of the specification in Duffie and Gârleanu [2001], which is the special case of $a_i = 1$. To see why the modification is relevant, consider a heterogeneous pool in the case $a_i = 1$ for $i = 1, \ldots, N$. Then obviously the common component must be smaller than the smallest default intensity in the pool (the components are non-negative, as we shall see). This implies that firms of low credit quality must have very low correlation with the common factor (high X_i relative to Y). Also, the firms of the highest credit quality will typically have relatively high correlation with the common factor. On the contrary, (1) imposes no implicit constraint on the combination of credit quality and correlation.[6]

Duffie, Pan and Singleton [2000] derive closed-form solutions – in terms of ordinary differential equations (ODEs) – to a range of relevant expectations involving affine processes. In some cases the ODEs have known explicit solutions – prominent examples include Ornstein-Uhlenbeck (Vasicek) and CIR processes, both of them in the pure diffusion case. Considering the empirical behavior of credit spreads, it appears appropriate to incorporate the possibility of jumps in the intensities. As in Duffie and Gârleanu [2001], we will accommodate that using the class of so-called *basic* affine jump-diffusions. More specifically, suppose the common component follows

$$dY(t) = \kappa_Y[\theta_Y - Y(t)]dt + \sigma_Y\sqrt{Y(t)}dW_Y(t) + dJ_Y(t)$$

and the idiosyncratic components follow

$$dX_i(t) = \kappa_i[\theta_i - X_i(t)]dt + \sigma_i\sqrt{X_i(t)}dW_i(t) + dJ_i(t)$$

for $i = 1, \ldots, N$, where

- $W_Y, W_1, \ldots, W_N$ are independent Wiener processes.
- $J_Y, J_1, \ldots, J_N$ are independent pure jump processes, independent of the Wiener processes.
- the jump times of $J_Y, J_1, \ldots, J_N$ are those of a series of Poisson processes with intensities $l_Y, l_1, \ldots, l_N$.
- the jump sizes of $J_Y, J_1, \ldots, J_N$ are independent of the jump times and follow exponential distributions with means $\mu_Y, \mu_1, \ldots, \mu_N$.

In short, we shall denote these as

$$Y \sim AJD(Y(0), \kappa_Y, \theta_Y, \sigma_Y, l_Y, \mu_Y)$$
$$X_i \sim AJD(X_i(0), \kappa_i, \theta_i, \sigma_i, l_i, \mu_i)$$

Note that a scaled AJD is an AJD with unchanged jump intensity but scaled jump size. The drift and diffusion parameters are scaled as usual in the CIR case, and thus

$$a_i Y \sim AJD(a_i Y(0), \kappa_Y, a_i \theta_Y, \sqrt{a_i}\sigma_Y, l_Y, a_i \mu_Y)$$

In the slightly more general class of affine processes, the volatility could have a constant term under the square root and the jump intensity could be an affine function of the state variable. This potentially added flexibility, however, seems very small and does not warrant the additional complications of estimating two more parameters in an already rich model. Also, the jump size distribution need not be exponential in the general affine class, but this choice ensures positive intensities and seems reasonable since most large jumps in credit spreads are positive. Moreover, this jump specification allows for an explicit solution to the ODEs for default probabilities.

Given the square root volatility structure and the positive jumps, the components – and thereby the default intensities – are strictly positive under the parameter restrictions $2\kappa_Y\theta_Y \geqslant \sigma_Y^2$ and $2\kappa_i\theta_i \geqslant \sigma_i^2$.

The sum of two AJDs with identical mean-reversion rates (κ's), volatilities and jump size parameters is again an AJD. Therefore, parameter restrictions could ensure that the default intensities also belong to the AJD class, although that is not needed. The independence of the two components ensures that the marginal default probabilities – needed for the calibration to the single-name CDS market – are known in closed form whether or not the intensity is AJD,

$$\begin{aligned} Pr\big(\tau_i \leqslant t\big) &= 1 - E\Big[e^{-\int_0^t \lambda_{i,s} ds}\Big] \\ &= 1 - E\Big[e^{-a_i \int_0^t Y(s)ds}\Big] \times E\Big[e^{-\int_0^t X_i(s)ds}\Big] \\ &= 1 - e^{A(t;\kappa_Y, a_i\theta_Y, \sqrt{a_i}\sigma_Y, l_Y, a_i\mu_Y) + B(t;\kappa_Y, \sqrt{a_i}\sigma_Y) a_i Y(0)} \\ &\qquad \times e^{A(t;\kappa_i,\theta_i,\sigma_i,l_i,\mu_i)+B(t;\kappa_i,\sigma_i)X_i(0)} \end{aligned}$$

where A and B are deterministic functions given explicitly in (9) of Appendix A.

The most general version of the model is very rich, with a total of $6 + 7N$ parameters. In many applications of the model it may be appropriate to reduce the parameter space without loosing the flexibility to control the most important quantities: CDS level, slope, volatility and correlation. Section 4.2 presents an empirical application of a parsimonious version of the model.

2.2.1 Semi-analytical loss distributions

For efficient computation of loss distributions in the model, define the common factor Z as the integrated common process,

$$Z(t) := \int_0^t Y(s)ds \tag{2}$$

Given the common factor, defaults occur independently across entities, and closed-form solutions of the default probabilities are given in the following form

$$\begin{aligned} p_i(t|z) := Pr\big(\tau_i \leqslant t \big| Z(t) = z\big) &= 1 - e^{-a_i z} E\big[e^{-\int_0^t X_i(s)ds}\big] \\ &= 1 - e^{-a_i z + A(t;\kappa_i,\theta_i,\sigma_i,l_i,\mu_i) + B(t;\kappa_i,\sigma_i)X_i(0)} \end{aligned} \tag{3}$$

again with A and B given explicitly in (9), Appendix A.

Unconditional joint default probabilities can be written as integrals of the conditional probabilities over the common factor distribution,

$$Pr\big(D(t) = j\big) = \int_{-\infty}^{\infty} Pr\big(D(t) = j \big| Z(t) = z\big)\, f_{Z(t)}(z)\, dz \tag{4}$$

where $f_{Z(t)}(\cdot)$ is the density function of the common factor. Once the density has been found, which is dealt with below, numerical integration can be done very efficiently using quadrature techniques since the integrand is a smooth function of the integrator.

Given the common factor, joint default probabilities can be calculated from the marginal default probabilities, (3), both with homogeneous and heterogeneous credit qualities.

In a homogeneous pool, the number of defaults given the common factor is binomially distributed,

$$Pr\big(D(t) = j \big| Z(t) = z\big) = \binom{N}{j} p_i(t|z)^j \left(1 - p_i(t|z)\right)^{N-j}$$

and the loss distribution, (4), is just a mixed binomial distribution.

In a heterogeneous pool, joint default probabilities can be obtained through the following recursive algorithm due to Andersen, Sidenius and Basu [2003]. Let $D^K(t)$ denote the number of defaults at time t in the pool consisting of the first K entities. Since defaults are conditionally independent, the conditional probability of observing j defaults in a K-pool can be written as

$$\begin{aligned} Pr\big(D^K(t) = j \big| Z(t) = z\big) &= Pr\big(D^{K-1}(t) = j \big| Z(t) = z\big) \times \big(1 - p_K(t|z)\big) \\ &\quad + Pr\big(D^{K-1}(t) = j-1 \big| Z(t) = z\big) \times p_K(t|z) \end{aligned} \tag{5}$$

for $j = 1, \dots, K$. For $j = 0$ the last term obviously disappears. The recursion starts from $Pr(D^0(t) = j|Z(t)) = \mathbf{1}_{\{j=0\}}$ and runs for $K = 1, \dots, N$ with $Pr(D(t) = j|Z(t)) = Pr(D^N(t) = j|Z(t))$. The intuition is that j defaults out of K names can be attained either by j defaults out of the first $K-1$ names and survival of the K'th name, or by $j-1$ defaults out of the first $K-1$ and a default of the K'th name. This method has previously been applied for semi-analytical valuation in copula models but is equally useful in the intensity-based models.

It only remains to find the distribution of the common factor. By definition, the characteristic function, $\varphi_{Z(t)}(\cdot)$, is given by

$$\varphi_{Z(t)}(u) := E[e^{iuZ(t)}] = \int_{-\infty}^{\infty} e^{iuz} f_{Z(t)}(z)\, dz$$

which is a Fourier transform of the density function. Thus, the density can be found by Fourier inversion,

$$f_{Z(t)}(z) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-iuz} \varphi_{Z(t)}(u)\, du$$

which can be computed very efficiently using fast Fourier transform (FFT) methods.[7]

The characteristic function of an integrated AJD is slightly outside the class of transforms considered in Duffie, Pan and Singleton [2000], but Proposition 1 in Appendix A proves that their result can be extended to cover this case as well.[8] As shown, the characteristic function is given by the exponential affine form

$$\varphi_{Z(t)}(u) = e^{\tilde{A}(0;t,u,\kappa_Y,\theta_Y,\sigma_Y,l_Y,\mu_Y)+\tilde{B}(0;t,u,\kappa_Y,\sigma_Y)Y(0)}$$

where $\tilde{A}$ and $\tilde{B}$ are complex-valued deterministic functions solving the ODEs in (10). The ODEs can be solved almost instantaneously, for example using Runge-Kutta methods.

In summary, the loss distribution is found by numerical integration of the conditional default distribution over the common factor density. The conditional default distribution is binomial in a homogeneous pool and is found by a simple recursion in a heterogeneous pool. The density of the common factor, in turn, is obtained through Fourier inversion of the characteristic function, which is known in closed form up to ODEs.

Relative to the semi-analytical factor copula solution in Andersen, Sidenius and Basu [2003], the only added computational complexity comes from the fact that the distribution of the common factor is more involved in the intensity model. If intensities were modeled as Ornstein-Uhlenbeck processes, the common factor would be Gaussian and the two methods would be equally tractable. The normal distribution is, however, not an appropriate description of default intensities – non-negativity is preferable and as we shall see jumps are needed to generate realistic correlation levels.

3 The copula approach

This section gives a brief outline of the copula approach to default correlation modeling and in particular of the three 1-factor copula models that will be compared to the intensity model in the next section.

In the copula approach, default before time t of entity i is defined as the event that a default variable (often interpreted as the asset value of entity i) is below some default boundary, $\chi_i \leqslant \xi_i(t)$. Default correlation is then modeled through the correlation of the χ_i's – a copula family is imposed on the joint distribution of the uniformly distributed transformations of the χ_i's. All three copulas considered here are 1-factor models where the default variable is a linear combination of a single market factor and an idiosyncratic factor. The loss distribution at t can then be computed from the marginal default probabilities and the copula function. Loss distributions at different time horizons are build using the same specification for the default variables, χ_i, but using different default boundaries (increasing with the horizon).

For each name, the default boundary, $\xi_i(t)$, is derived from the marginal default probabilities implied from a term structure of single-name CDS spreads.

This is usually done by backing out a deterministic default intensity curve – for example piecewise constant or linear between the observed maturities. In the copula applications of this paper, a piecewise constant intensity is assumed.

3.1 Gaussian copula

In the 1-factor Gaussian copula, the default variable is defined as

$$\chi_i = \beta_i \Upsilon + \sqrt{1-\beta_i^2}\, \epsilon_i$$

where $\Upsilon, \epsilon_1, \ldots, \epsilon_N$ are independent standard normal random variables, and the factor loadings are constant. The correlation between any pair of default variables, χ_i and χ_j, is then $\beta_i \beta_j$. For all the models in the comparison, homogeneous correlation parameters are assumed across all entities, i.e. $\rho = \beta_i \beta_j$. The conditional and unconditional default probabilities in the Gaussian copula are given in terms of the standard normal distribution function.

3.2 Gaussian copula with random factor loadings

In the random factor loadings (RFL) copula proposed by Andersen and Sidenius [2004-5], the factor loading is a function of the common factor. In the Gaussian version, the default variable is defined as

$$\chi_i = \beta_i(\Upsilon)\Upsilon + \alpha_{1i}\epsilon_i + \alpha_{2i}$$

where $\Upsilon, \epsilon_1, \ldots, \epsilon_N$ are independent standard normal variables, and α_{1i} and α_{2i} are constants chosen to ensure χ_i has zero mean and unit variance.

This paper applies the tractable two-point regime switching version with homogeneous correlations,

$$\beta_i(\Upsilon) = \begin{cases} \sqrt{\rho_1} & \text{for } \Upsilon \geqslant \nu \\ \sqrt{\rho_2} & \text{for } \Upsilon < \nu \end{cases}$$

The intuition is that the correlation is allowed to depend on the state of the economy, and there is empirical evidence suggesting that we should expect to see higher correlation in recessions ($\rho_2 > \rho_1$). The conditional and unconditional default probabilities in this version of the model are still known in closed form. See Andersen and Sidenius [2004-5] for the expressions for default probabilities as well as for α_{1i} and α_{2i}.

3.3 Double-t copula

In the double-t copula proposed by Hull and White [2004], the default variable is

$$\chi_i = \beta_i \frac{\Upsilon}{std.dev.(\Upsilon)} + \sqrt{1-\beta_i^2}\frac{\epsilon_i}{std.dev.(\epsilon_i)}$$

where the common and idiosyncratic factors follow t distributions with d degrees of freedom, and the loadings are constant. The homogeneous correlation between default variables is $\rho = \beta_i\beta_j$. The conditional default probabilities are known from the t distribution. The default boundaries, however, must be found by Monte Carlo simulation or numerical integration since the unconditional default probabilities are given in terms of the distribution of a sum of two t distributions which is unknown (not a t distribution). For more details, see Hull and White [2004].

4 Basket credit derivatives valuation

In this section, the intensity-based model will be applied for valuation of synthetic CDOs. After a description of the product, we shall test how the model conforms with market prices.

4.1 Synthetic CDOs

In a CDO, the credit risk of an underlying pool of bonds, loans or CDSs is passed through, according to some prioritization scheme, to a number of tranches with varying risk/return profiles. The so-called equity tranche is the most risky piece being exposed to the first defaults, whereas more senior tranches only incur losses if the portfolio loss exceeds certain thresholds. The structure is referred to as a *synthetic* CDO if the underlying credit risk is constructed synthetically through CDSs. For *cash* CDOs (or funded CDOs), the tranches resemble bonds with an up-front price in return for future interest, principal and recovery payments. This section considers *unfunded* synthetic CDOs with payoff structures closer to basket default swaps. The credit risk in buying a cash CDO tranche corresponds to selling loss protection on an interval of the underlying portfolio loss distribution in an unfunded CDO. The protection seller agrees to make potential future loss payments in return for periodic premium payments. The precise cash flows are described below.

The early days of the CDO market were characterized by low liquidity and non-standardization. Traditional cash CDOs have deal-specific documentation and underlying portfolio, and the payoff structure (the so-called CDO waterfall) is often complex and path-dependent. To improve liquidity, a number of investment banks introduced market making in standardized synthetic CDOs in 2003. The documentation and payoff structure were standardized as well as the underlying pool consisting of a benchmark CDS index. The product is therefore known as an *index tranche*, but it is sometimes also referred to as a *single-tranche CDO* due to the fact that it facilitates trading of a single CDO tranche as an OTC derivatives contract between two counterparties. This is much more flexible than the issuance process for traditional CDOs, where the originator has to set up a special purpose vehicle (SPV), obtain tranche credit ratings from the rating agencies, and find investors for the entire CDO capital structure. Moreover, the launch of market making in index tranches has enabled investors to take on both long and short positions in CDOs.

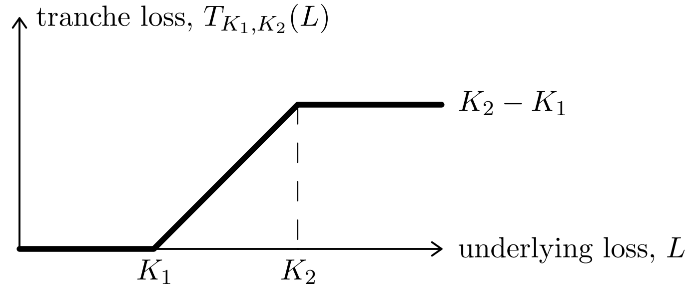

Exhibit 1: Index tranche loss. The loss on a tranche with attachment point K_1 and exhaustion point K_2 as a function of underlying portfolio loss.

4.1.1 Index tranche cash flows

Consider an index tranche covering a part of the portfolio loss distribution from an attachment point, K_1, to an exhaustion point, $K_2 > K_1$. The tranche loss as a function of portfolio loss L is then given by

$$T_{K_1,K_2}(L) := \max\left\{ \min\{L, K_2\} - K_1, 0\right\}$$

which has the familiar structure of a call spread option written on the portfolio loss, see Exhibit 1. In broad outline, the protection seller pays the observed tranche losses as they occur and receives premium payments on the remaining principal, which amortizes with the loss payments.

To be specific, suppose a T-year contract has an annual premium of S that is paid periodically in arrears at the discrete dates $t_1, t_2, \ldots, t_M = T$. The most liquid maturity is 5 years, and typical contracts specify quarterly payments which is also the case for the contracts considered in this paper. The cumulative percentage loss of the portfolio incurred up to time t is denoted by $L_t = (1-\delta)D(t)/N$. Furthermore, the time-t default-free short rate and s-year zero-coupon bond are denoted by r_t and $P(t, t+s)$, respectively.

The notional amount is some multiple, which is without loss of generality set to 1, of the tranche thickness, $K_2 - K_1$. The value of the protection leg is then

$$Prot(0,T) = E\Big[\int_0^T e^{-\int_0^t r_s ds} dT_{K_1,K_2}(L_t)\Big]$$

With quarterly premium payments, the set of premium payment dates provides a natural discretization of the integral. A more fine-grained partition could be applied with obvious notational changes in the following equations. If, in addition, losses on average occur in the middle of these intervals and interest rates are uncorrelated with losses, as is often assumed for these products,

$$Prot(0,T) = \sum_{j=1}^{M} P\big(0, \frac{t_j + t_{j-1}}{2}\big)\Big\{E\big[T_{K_1,K_2}(L_{t_j})\big] - E\big[T_{K_1,K_2}(L_{t_{j-1}})\big]\Big\}$$

where $t_0 = 0$.

The value of the premium leg is

$$Prem(0,T;S) = E\Big[\sum_{j=1}^{M} e^{-\int_0^{t_j} r_s ds} S(t_j - t_{j-1}) \int_{t_{j-1}}^{t_j} \frac{K_2 - K_1 - T_{K_1,K_2}(L_s)}{t_j - t_{j-1}} ds\Big]$$

where the integral represents the remaining principal over the interval t_{j-1} to t_j, which determines the premium payment a date t_j. With discretization and assumptions as above, we get

$$\begin{aligned} Prem(0,T;S) = S \sum_{j=1}^{M} &(t_j - t_{j-1}) P\big(0, t_j\big) \\ &\times \Big\{K_2 - K_1 - \frac{1}{2}\Big(E[T_{K_1,K_2}(L_{t_{j-1}})] + E[T_{K_1,K_2}(L_{t_j})]\Big)\Big\} \end{aligned}$$

The fair tranche premium (or spread) is then defined as the premium, S, that makes the value of the premium leg equal to the value of the protection leg.

The equity tranche spread is very high and therefore the timing of defaults becomes very important. To reduce the timing risk, the equity tranche premium leg is usually divided into an *up-front fee* plus a fixed running premium of 500 basis points (bps).[9] The up-front fee is quoted as a fraction, U, of tranche notional, i.e.

$$\begin{aligned} &Prem_{eq}(0,T;0.05,U) = UK_2 \\ &+ 0.05 \sum_{j=1}^{M} (t_j - t_{j-1}) P(0,t_j)\Big\{K_2 - \frac{1}{2}\Big(E[T_{0,K_2}(L_{t_{j-1}})] + E[T_{0,K_2}(L_{t_j})]\Big)\Big\} \end{aligned}$$

The fair equity tranche price is then quoted as the up-front fee, U, balancing the premium and protection legs.

All tranche prices are expressed in terms of expected tranche losses at different horizons, which are obtained from the portfolio loss distributions. In principle M loss distributions are needed, but very good approximations could be achieved by interpolating expected tranche losses from a lower number of loss distributions.

4.2 Empirical application

In the following, the intensity model is calibrated to a data set consisting of 5y index tranche market quotes as well as the CDS spreads of the underlying reference names. The CDO quotes are available for the five benchmark tranches trading on each of the two most liquid CDS indices: The Dow Jones iTraxx index, consisting of 125 European investment grade companies, with 0-3%, 3-6%, 6-9%, 9-12% and 12-22% tranches; and the Dow Jones CDX index, consisting of 125 North American investment grade companies, with 0-3%, 3-7%, 7-10%, 10-15% and 15-30% tranches. The tranche prices were obtained from Bloomberg, using the Bloomberg generic price source. The calibration is performed on two trading days: August 23, 2004, and December 5, 2005.

The individual default intensities are calibrated to 1y and 5y CDS spreads for the underlying names of the indices – the CDS data are summarized in Exhibit 2.[10] Relative to the CDX pool, the iTraxx pool is on average of better credit quality, and the spread levels in iTraxx are more homogeneous. All the (individual) CDS curves are upward-sloping, as expected for investment grade companies. The 3y CDS spread is usually also observed but matching the 1y and 5y points captures the shape of the curve sufficiently well. The calibration to the single-name CDS market is straight-forward, but for completeness Appendix B provides a few details.

The recovery rate is assumed constant at $\delta = 35\%$, consistent with the em-

CDS spreads (mid, in bps)	Average	Std. dev.	Min	Median	Max
August 23, 2004:					
DJ iTraxx 1y	15	10	2	12	57
DJ iTraxx 5y	39	21	11	36	127
DJ CDX 1y	23	36	3	16	349
DJ CDX 5y	67	70	17	48	460
December 5, 2005:					
DJ iTraxx 1y	10	9	2	8	76
DJ iTraxx 5y	37	32	9	29	284
DJ CDX 1y	15	51	1	9	584
DJ CDX 5y	51	65	6	35	633

Exhibit 2: Underlying CDS spread statistics.

pirical evidence for senior unsecured bonds reported by Moody's (Hamilton et al. [2004], Exhibit 15). A higher (lower) recovery rate would be offset by higher (lower) CDS-implied default intensities, and since index tranche spreads are primarily driven by the loss rates (i.e. $(1-\delta)$ times default intensity) in the underlying pool, the parameter value for the recovery rate is, within reasonable bounds, not too critical. A decrease in the recovery rate would assign slightly more mass to extreme losses (high and low) and would as such act as a small increase in correlation. That is however a second order effect, and the results are fairly robust to the recovery rate. Finally, the default-free interest rates are taken from the swap curves (EUR for iTraxx and USD for CDX) on the relevant trading days.

4.2.1 Calibration

As mentioned the general version of the model has a large number of parameters. This paragraph proposes a parsimonious version of the model that is appropriate for the present application and data set, while alternative parametrizations might be more empirically reasonable or useful in other applications. The purpose here is to test the performance of the intensity model in a specification that is parsimonious enough to be applicable in practise, yet flexible enough to match the individual credit curves. Furthermore, it is constructed such that the default intensity correlation is similar across names and can be controlled by a single correlation parameter. Similar correlation is accommodated by enforcing that, by some standards, the division of the intensities into systematic and idiosyncratic components is proportionally the same for all names.

Recall that $\lambda_i(t) = a_i Y(t) + X_i(t)$ and

$$\begin{aligned} a_i Y &\sim AJD(a_i Y(0), \kappa_Y, a_i\theta_Y, \sqrt{a_i}\sigma_Y, l_Y, a_i\mu_Y) \\ X_i &\sim AJD(X_i(0), \kappa_i, \theta_i, \sigma_i, l_i, \mu_i) \end{aligned}$$

To achieve similar correlation across the names, a_i is going to be a spread level parameter – a name with high (low) credit spreads will have a high (low) a_i – and $a_i = 1$ is going to represent an average name.

First, we make the following three assumptions

$$\kappa_i = \kappa_Y, \quad \sigma_i = \sqrt{a_i}\sigma_Y, \quad \mu_i = a_i\mu_Y$$

Without these, the sum of the systematic and idiosyncratic terms would not be an AJD. As mentioned this is not needed for the semi-analytical solution but it seems fairly reasonable to use a common mean-reversion rate as well as high volatilities and jump sizes for high spread names.

Second, assume that all names have the same aggregate jump intensity

$$\bar{l} := l_Y + l_i$$

To keep the notation consistent we also rename the three parameters

$$\overline{\kappa} := \kappa_Y, \quad \overline{\sigma} := \sigma_Y, \quad \overline{\mu} := \mu_Y$$

This way, the individual default intensities reduce to the following AJDs

$$\lambda_i \sim AJD(a_i Y(0) + X_i(0), \overline{\kappa}, a_i\theta_Y + \theta_i, \sqrt{a_i}\overline{\sigma}, \bar{l}, a_i\overline{\mu})$$

and the overline-parameters can be seen as AJD parameters for an average underlying name ($a_i = 1$).

Third, as in Duffie and Gârleanu [2001], it is assumed that the systematic parts of the mean-reversion level and of the jump intensity are identical and given by the same fraction for all underlying names,

$$\overline{w} := \frac{a_i \theta_Y}{a_i \theta_Y + \theta_i} = \frac{l_Y}{l_Y + l_i} \in [0, 1] \tag{6}$$

Thereby, $\overline{w}$ is a correlation parameter representing the systematic share of the intensities: a low $\overline{w}$ implies that most of the intensity is idiosyncratic, whereas a high $\overline{w}$ implies a high systematic share. In addition we ensure that the initial value of an average (hypothetical) name satisfies the same division,

$$\overline{w} := \frac{Y(0)}{Y(0) + X_{avg}(0)} \tag{7}$$

(The same division cannot be ensured for all individual initial values since the CDS slopes differ, but in practise this will almost be satisfied if the CDS shapes only differ slightly.)

These assumptions leave $\overline{\kappa}$, $\overline{\sigma}$, $\bar{l}$, $\overline{\mu}$ and $\overline{w}$ as free parameters. For each combination of the five free parameters, the procedure is the following:

1. For a hypothetical average name ($a_i = 1$), calibrate the initial level, $Y(0) + X_{avg}(0)$, and the mean-reversion level, $\theta_Y + \theta_{avg}$, to fit the 1y and 5y spreads of the average CDS curve. Using (7) and (6), $Y(0)$ and θ_Y are then given by a fraction $\overline{w}$ of the initial and mean-reversion levels, respectively.

2. For each underlying name, calibrate $X_i(0)$ and θ_i to fit the name-specific 1y and 5y CDS spreads. In doing so, the scaling factor, a_i, is always chosen to ensure that the systematic part of the mean-reversion level remains equal to $\overline{w}$ (from (6)), i.e. $a_i = \frac{\overline{w}\theta_i}{(1-\overline{w})\theta_Y}$.

The a_i's from this procedure will be spread around 1 with values close to the ratio of individual spread to average spread. If a homogeneous underlying pool is assumed, the second step is obviously not needed.

The model is fitted to index tranches by minimizing the tranche price deviations while maintaining exact calibration to the single-name CDS quotes. The criteria function is given by the root mean square price errors relative to bid/ask spreads,

$$\text{RMSE} = \sqrt{\frac{1}{5}\sum_{k=1}^{5}\Big(\frac{S_{k,\,market\ mid} - S_{k,\,model}}{S_{k,\,market\ ask} - S_{k,\,market\ bid}}\Big)^2} \tag{8}$$

where S_k is the spread of tranche k. The liquidity varies across tranches, and this way deviations from the most reliable prices get the highest weights.

4.2.2 Results from August 2004

The results for the iTraxx tranches on August 23, 2004, are reported in Exhibit 3. The calibrated intensity model fits the market prices very well, with an error measure of 0.34 and all model prices within the bid/ask range. The jump parameters are relatively high but not too implausible. To understand the magnitude, consider a typical ($a_i = 1$) company with the fitted parameter values. With initial value 0.02% and mean-reversion level 0.20%, the 1y and 5y CDSs are priced at the iTraxx average of 15bps and 39bps, respectively. A jump in the default intensity of 6.6% (the mean jump size) leads to a 5y spreads of 335bps. The median credit rating in the underlying pool is A, and a jump of that magnitude corresponds approximately to a rating migration from A to BB.[11] The observed 1-year frequencies for migrations from A to BB or worse are on average around 0.8% but vary a lot across observation periods. Therefore, an expected credit spread jump of 300bps once or twice in a century ($\bar{l}$ is 1.7%) does not seem excessive, at least not under the risk-neutral measure.

The correlation parameter is fitted at 83% implying that most of the variation in the default intensities stem from common fluctuations. This number means that, given that a jump occurs, with probability 83% the jump is in the common component. This appears very high considering the degree of co-movement observed in CDS premiums across time and considering the fact that sudden blow-ups in credit spreads often are idiosyncratic. Therefore, the combination of jump and correlation parameters seems unrealistic. As mentioned in the introduction, the fact that we can interpret parameters and discuss whether they are reasonable or not is exactly one of the advantages of the intensity-based model. This may also be helpful in forming opinions on the absolute pricing of correlation products (provided that the model is considered reasonable).

The criteria function is relatively flat in some regions of the parameter space since some of the parameters govern the same effects. For example, a combination of higher jump intensity and lower jump size might provide almost similar results. Also, higher volatility in the common intensity acts as higher correlation, and therefore a lower correlation parameter might be offset by higher values for the diffusion volatility or jump parameters.

The results for the CDX index on the same trading day are reported in Exhibit 4. Compared with the iTraxx results, the intensity model is fitted with higher jump intensity but lower jump size. Again, to investigate the jump impact we consider a company with the fitted parameters. A mean jump (4.7%) in the default intensity causes an increase in the 5-year CDS premium of around 250bps, which historically corresponds to a downgrade from A to BBB-BB. An expected frequency of 4.8% for such a migration is again relatively high. The correlation parameter is fitted at 77%, which is lower than for iTraxx but still rather high.

All the models under consideration show a worse fit to the CDX tranches but the relative performance of the models is the same. As expected, the results for the pure diffusion case (restricting the jump parameters to zero) show a very poor price fit. Diffusion intensities are unable to generate enough default correlation to explain market prices – too high equity spreads and too low senior spreads.

For comparison, the tables also reports the fit for the three copula models introduced in the previous section: the Gaussian copula, the random factor loading (RFL) copula proposed by Andersen and Sidenius [2004-5], and the double-t

copula of Hull and White [2004]. Unsurprisingly, the Gaussian copula provides a very poor description of the market prices. At the fitted value, a higher correlation is needed to match the equity tranche and the senior tranches, whereas much lower correlation is needed to match the 3-6% mezzanine tranche. The two alternative copulas match market prices much better, though not as good as the intensity model. We should note though that the intensity model has more free parameters than the two competing copulas, and introducing more parameters in the copulas may improve their fits. For both indices the RFL copula is fitted with higher correlation in recessions, as intuitively expected and as suggested by the numerical examples in Andersen and Sidenius [2004-5]. For example, the iTraxx correlation is 30.3% in the worst 2.7% states of the economy and 9.3% in the rest. The double-t copula obtains the closest fit to both indices with 4 degrees of freedom, which is consistent with the observations in Hull and White [2004].

The exercise has also been done under the assumption of homogeneous underlying names at the average spread levels. The results (not reported) show that for iTraxx tranches there is no added price fitting ability from the true heterogeneous pool relative to the approximating homogeneous pool. For CDX however, incorporating the 125 heterogeneous names does in fact improve the fit to the tranche prices considerably for all models under consideration. This might be explained by the fact that the CDX pool is much more diverse than the iTraxx pool in terms of credit quality.

Implied correlation skews

Recently, brokers and investment banks have been quoting tranche prices in terms of implied correlations through a standard model, analogous to the practice of quoting implied Black-Scholes volatilities in option markets. The standard model is the 1-factor Gaussian copula. For each tranche, the implied correlation is defined as the homogeneous correlation parameter needed in the standard model to reproduce the market price. This quotation device, sometimes referred to as *com-*

pound correlations, suffers from both existence and uniqueness problems (which are not encountered with Black-Scholes volatilities). Mezzanine tranche spreads are not monotone in correlation, and we may observe arbitrage-free market prices that are not attainable through any choice of compound correlation. An alternative quotation device is the so-called *base correlations*, defined as the implied correlations on a sequence of hypothetical equity tranches consistent with the market prices on the traded tranches. Base correlations are unique since equity prices are monotone, but existence is still not guaranteed and they can be very difficult to interpret.[12]

If the Gaussian copula was the true model, the implied compound correlations would be identical across tranches on the same underlying pool. As we saw in the calibration, however, in practice the Gaussian copula is not able to match all tranches with a single correlation parameter. Plotting the implied correlations for the different tranches gives rise to a so-called correlation skew.

Exhibits 5 and 6 illustrate the correlation skew in market prices as well as in the calibrated models. The model-implied correlations in AJD, RFL and double-t are fairly close to the market-implied correlations – especially for the iTraxx index. The correlations skew is very significant, and this pattern for the skew has been fairly stabile since market prices on index tranches became available. Note that the CDX 3-7% tranche has a market spread equivalent to a Gaussian copula correlation as low as 0.1%. As mentioned, an implied correlation need not exist for any observed arbitrage-free spread, and that is almost the case for this tranche.

Market-implied loss distributions

The shape of the market-implied loss distribution can be inferred from the models that have proven able to match the correlation skew. Exhibits 7 and 8 demonstrate the loss distributions for the two index pools, and the patterns are quite similar for the two indices although the iTraxx tranches have been fitted more closely by the models than the CDX tranches.

It appears that the market-implied loss density crosses the Gaussian copula standard model at three loss levels: relative to a fitted Gaussian copula, the market assigns lower probability to zero losses, higher probability to small-medium losses, lower probability to moderately high losses, and again higher probability to extreme loss scenarios.

The fat upper tail on the market-implied loss distribution is also evident from the low Gaussian copula spreads and high implied correlations for the most senior tranches. The lower end of the distribution is more complicated. As we saw, most of the models tend to overestimate the spread on the second-loss mezzanine tranche (iTraxx 3-6%, CDX 3-7%), indicating that the market assigns a relatively high probability mass to losses lower than, say, 3-4%. This is, however, only consistent with the market price of the equity tranche if the mass is skewed towards high equity losses with low probability of no defaults.

We also note from the lower panels of the figures that the RFL copula, in the regime-switching version applied here, produces a bimodal loss distribution, which may sometimes be difficult to explain intuitively.

4.2.3 Results from December 2005

From August 2004 to December 2005, the composition of the indices has changed, the underlying spread levels have decreased slightly and the spread dispersion has increased (Exhibit 2). Looking at the tranche market prices in Exhibits 9 and 10, it is evident that the implied correlations have decreased. Overall, the senior spreads have dropped to nearly half of the August 2004 levels, whereas the underlying spreads are only moderately lower. This indicates that lower correlation assumptions are priced into the tranches in December 2005. Also, higher equity tranche price in spite of unchanged or lower average spread may point to lower correlation, although this could also be explained by the fact that the worst underlying credits trade at higher spreads than in 2004. The liquidity of the index tranche market has improved through this period and bid/ask spread have dropped by a factor of 2-3, though not as significantly on a relative scale.

As expected, all models are fitted with lower correlation parameters. The AJD model and the RFL copula performs comparably and clearly better than the rest of the models. Again, the pure diffusion model and the Gaussian copula have very poor fitting ability. The correlation in the Gaussian copula is remarkably low due to the very low market spread on second-loss mezzanine tranches – in fact the CDX 3-7% market spread is below the lower spread bound in the Gaussian copula. The double-t copula is fitted with fatter tails (3 degrees of freedom) than in 2004. The correlation parameter in this model is very stabile across the two indices on both trading days, though not through time.

It is interesting to note see how the implied parameters in both AJD and RFL indicate that a catastrophe scenario is priced into the iTraxx tranches. In RFL, this is reflected by the boundary solution of 100% correlation in very deep recession outcomes (0.5%). In AJD, an extremely high jump size (18.7% mean) is implied, which on the other hand only is expected to occur at a frequency of once in 250 years.

In conclusion, the ability of the intensity-based model to fit CDO market prices appears to be at least comparable with alternative models proposed in the literature. The results should of course be interpreted with some caution since only two trading days are considered.

5 Conclusion

This paper illustrated a semi-analytical valuation method for basket credit derivatives in a multivariate intensity-based model. Analytical solutions are important for parameter estimation and calibration as well as for calculating sensitivities to parameters and inputs. This is particularly important in an intensity-based model since simulation of the model would require sampling of an entire intensity path for each underlying entity – as opposed to copulas where a default time can be sampled by drawing a single or a few random deviates depending on the relevant copula family.

The model fits the market prices of synthetic CDOs reasonably well. In other words, the model is able to generate pricing patterns consistent with the observed correlation skews. This allows for relative valuation of off-market correlation products from benchmark products in a fully consistent model, and thereby dispenses with the interpolation schemes based on implied correlations in the Gaussian copula, which are widely used in practice.

The basket credit derivatives market is still fairly immature, and in explaining the correlation skew it is difficult to rule out supply and demand effects caused by market segmentation or market inefficiency – such explanations are even being discussed for the steep volatility skews observed in the much more established and liquid equity option markets. Therefore, a perfect fit to the market prices should perhaps not be expected, and the choice of an appropriate model involves more than just searching for the best price fit. An interesting topic for future research is the hedging performance of the alternative correlation models. A significant amount of model risk is involved in the widespread delta-hedging of correlation products – different models suggest different hedge ratios. More light may be shed on this issue as the market matures and more market data become available.

Appendix A: Affine jump-diffusions

This appendix reports two useful results for affine jump-diffusion processes. The paper applies a sub-class (introduced by Duffie and Gârleanu [2001]) of the affine processes: a *basic affine process* is a process following a stochastic differential equation

$$dx_t = \kappa(\theta - x_t)dt + \sigma\sqrt{x_t}dW_t + dJ_t$$

initiated at x_0, where W is a Wiener process and J is a pure jump process, independent of the Wiener process, with jump times from a Poisson distribution with intensity l and jump sizes exponentially distributed with mean μ. We denote this by $x \sim AJD(x_0, \kappa, \theta, \sigma, l, \mu)$. For the mathematical foundation of affine processes in general, see Duffie, Filipović and Schachermayer [2003].

A.1 Default probabilities

The first result, used for computing default probabilities, is the following:

$$E\left[e^{-\int_0^T x_s ds}\right] = e^{A(T;\,\kappa,\theta,\sigma,l,\mu)+B(T;\,\kappa,\sigma)x_0}$$

where

$$\begin{aligned} A(T;\,\kappa,\theta,\sigma,l,\mu) &= \frac{\kappa\theta\gamma}{bc_1d_1}\,\log\Big(\frac{c_1 + d_1e^{bT}}{-\gamma}\Big) + \frac{\kappa\theta}{c_1}\,T \\ &\quad + \frac{l(c_2d_1/c_1 - d_2)}{bc_2d_2}\,\log\Big(\frac{c_2 + d_2e^{bT}}{c_2 + d_2}\Big) + \frac{l - c_2l}{c_2}\,T \qquad (9) \\ B(T;\,\kappa,\sigma) &= \frac{1 - e^{bT}}{c_1 + d_1e^{bT}} \end{aligned}$$

with

$$\gamma = \sqrt{\kappa^2 + 2\sigma^2}$$

$$c_1 = -(\gamma + \kappa)/2$$

$$d_1 = c_1 + \kappa$$

$$c_2 = 1 - \mu/c_1$$

$$d_2 = (d_1 + \mu)/c_1$$

$$b = d_1 + (\kappa c_1 - \sigma^2)/\gamma$$

This expression is a special case of a more general formula in Duffie and Singleton [2003], Appendix A.5.

A.2 Characteristic functions

The second relevant expectation is used for computing the characteristic function of an integrated AJD,

$$\varphi(u) = E\left[e^{iu\int_0^T x_t dt}\right]$$

This is slightly outside the class of transforms covered in Duffie, Pan and Singleton [2000], but the extension in the following proposition shows that the characteristic function of an integrated AJD also is given by an exponential affine form.

Proposition 1: *For a basic affine jump-diffusion process,* $x \sim AJD(x_0, \kappa, \theta, \sigma, l, \mu)$*,*

$$E\left[e^{iu\int_0^T x_t dt}\right] = e^{\tilde{A}(0;T,u,\kappa,\theta,\sigma,l,\mu)+\tilde{B}(0;T,u,\kappa,\sigma)x_0}$$

where $\tilde{A}(t), \tilde{B}(t): [0,T] \to \mathbb{C}$ *are complex-valued deterministic functions solving the following set of ODEs:*

$$\begin{aligned}
\tilde{A}'_{Re}(t) &= -\kappa\theta\tilde{B}_{Re}(t) - \frac{l\mu\big[\tilde{B}_{Re}(t) - \mu\tilde{B}_{Re}(t)^2 - \mu\tilde{B}_{Im}(t)^2\big]}{[1-\mu\tilde{B}_{Re}(t)]^2 + \mu^2\tilde{B}_{Im}(t)^2} \\
\tilde{A}'_{Im}(t) &= -\kappa\theta\tilde{B}_{Im}(t) - \frac{l\mu\tilde{B}_{Im}(t)}{[1-\mu\tilde{B}_{Re}(t)]^2 + \mu^2\tilde{B}_{Im}(t)^2} \\
\tilde{B}'_{Re}(t) &= \kappa\tilde{B}_{Re}(t) - \frac{\sigma^2}{2}\big[\tilde{B}_{Re}(t)^2 - \tilde{B}_{Im}(t)^2\big] \\
\tilde{B}'_{Im}(t) &= \kappa\tilde{B}_{Im}(t) - \sigma^2\tilde{B}_{Re}(t)\tilde{B}_{Im}(t) - u
\end{aligned} \tag{10}$$

with boundary conditions $\tilde{A}_{Re}(T) = \tilde{A}_{Im}(T) = \tilde{B}_{Re}(T) = \tilde{B}_{Im}(T) = 0$*. (For notational simplicity, the dependence on* T*,* u *and the AJD parameters has been suppressed in the ODEs.)*

Proof: The proof is similar in spirit to the proof of Proposition 1 in Duffie, Pan and Singleton [2000]. Define $Z(t) := \int_0^t x_s$ and

$$\Psi_t := e^{\tilde{A}(t)+\tilde{B}(t)x_t+iuZ(t)}$$

where $\tilde{A}$ and $\tilde{B}$ are complex functions of time with boundary conditions $\tilde{A}(T) = \tilde{B}(T) = 0$. If we can find such functions ensuring that Ψ is a martingale, we know that

$$e^{\tilde{A}(0)+\tilde{B}(0)x_0} = E\big[e^{\tilde{A}(T)+\tilde{B}(T)x_T+iuZ(T)}\big] = E\big[e^{iuZ(T)}\big]$$

and we are done.

From the general version of Itô's Lemma on $\Psi_t = \Psi(t, x_t, Z(t))$, we have

$$\Psi_T - \Psi_0 = \int_0^T \mu_\Psi(t,x_t)\Psi_t dt + \int_0^T \sigma_\Psi(t,x_t)\Psi_t dW_t + J_\Psi(T) \tag{11}$$

where

$$\begin{aligned}
\mu_\Psi(t,x) &= \tilde{A}'(t) + \tilde{B}'(t)x + \tilde{B}(t)\kappa(\theta - x) + iux + \frac{1}{2}\tilde{B}(t)^2\sigma^2 x \\
\sigma_\Psi(t,x) &= \tilde{B}(t)\sigma\sqrt{x} \\
J_\Psi(t) &= \sum_{s\leqslant t} \Psi_s - \Psi_{s-}
\end{aligned}$$

The second term on the right hand side of (11) is a martingale (the necessary integrability condition on $\sigma_\Psi(t,x)\Psi_t$ is satisfied in the basic affine class).

Let H denote the stochastic jump size, and define the jump transform, $h(\cdot) : \mathbb{C} \to \mathbb{C}$, by

$$h(c) := E\big[e^{cH}\big]$$

The jump size is exponentially distributed with mean μ. Thus,

$$h(c) = \frac{1}{\mu}\int_0^\infty e^{vc_{Re}+ivc_{Im}} e^{-v/\mu} dv$$

We shall evaluate the jump transform in $c = \tilde{B}$. We know $\tilde{B}_{Re}$ is a differentiable function of time ending in zero at time T. Thus, for $\tilde{B}_{Re}$ to take on strictly positive values, it would have to pass through zero with a negative slope. This is impossible since at $\tilde{B}_{Re}(t) = 0$ we have $\tilde{B}'_{Re}(t) = \frac{\sigma^2}{2}\tilde{B}_{Im}(t)^2 \geqslant 0$. Hence, $\tilde{B}_{Re}(t) \leqslant 0$ for $t \in [0,T]$.

Knowing that $c_{Re} < 1/\mu$, the integral can be calculated as

$$h(c) = \frac{(1-\mu c_{Re}) + i\mu c_{Im}}{(1-\mu c_{Re})^2 + \mu^2 c_{Im}^2}$$

Define

$$g(t) := l\big(h(\tilde{B}(t)) - 1\big)\Psi_t$$

From Lemma 1, Appendix A, in Duffie, Pan and Singleton [2000], $J_\Psi(t) - \int_0^t g(s)ds$ is a martingale (the necessary integrability condition on $g(t)$ is satisfied in the basic affine class).

Therefore, from (11), Ψ is a martingale if $\mu_\Psi(t,x)\Psi_t = -g(t)$ for all (t,x). Applying the matching principle, we see that this is fulfilled if

$$\tilde{B}'(t) - \tilde{B}(t)\kappa + iu + \frac{1}{2}\tilde{B}(t)^2\sigma^2 = 0$$

(from the x terms) and

$$\tilde{A}'(t) + \tilde{B}(t)\kappa\theta = -l\big(h(\tilde{B}(t)) - 1\big)$$

These two complex ODEs can be written out as the four deterministic ODEs in (10). □

Appendix B: Pricing Credit Default Swaps

This appendix gives a brief introduction to credit default swaps (CDSs). For more details, refer to e.g. Duffie [1999] or Hull and White [2000].

A CDS is an insurance contract between two counterparties written on the event of default of a third reference entity. In the event of default before maturity of the contract, the protection seller pays the loss given default to the protection buyer. That is, at default, the protection buyer delivers a defaulted bond to the protection seller in return for face value. To compensate for that, the protection buyer pays fixed premium payments periodically until default or maturity of the contract is reached.

Formally, with notation as in Sections 2 and 4, the value of the protection leg in a T-year CDS is

$$Prot(0,T) = E\Big[e^{-\int_0^\tau r_s ds}\mathbf{1}_{\{\tau\leqslant T\}}(1-\delta)\Big]$$

Suppose the CDS contract specifies that the annual premium, S, is paid in arrears at $t_1, t_2, \ldots, T_M = T$. Premium payments are made conditional on survival of the reference entity, and in the event of default, an accrual premium payment is made for the period since the previous payment date. Hence, the value of the premium leg is

$$Prem(0,T;S) = E\Big[\sum_{j=1}^{M} e^{-\int_0^{t_j} r_s ds}\mathbf{1}_{\{\tau>t_j\}}S(t_j - t_{j-1}) + e^{-\int_0^{\tau} r_s ds}\mathbf{1}_{\{t_{j-1}<\tau\leqslant t_j\}}S(\tau - t_{j-1})\Big]$$

With discretization and independence assumptions between recovery rates, interest rates and default events as in Section 4, the value of the protection leg is

$$Prot(0,T) = (1-\delta)\sum_{j=1}^{M} P(0,\frac{t_j + t_{j-1}}{2})\big[Pr(\tau\leqslant t_j) - Pr(\tau\leqslant t_{j-1})\big]$$

Similarly, the value of the premium leg is

$$Prem(0,T;S) = S\sum_{j=1}^{M}(t_j - t_{j-1})P(0,t_j)Pr(\tau > t_j) + \frac{t_j - t_{j-1}}{2}P(0,\frac{t_j + t_{j-1}}{2})\big[Pr(\tau \leqslant t_j) - Pr(\tau \leqslant t_{j-1})\big]$$

The fair CDS premium, S, is then given as the solution to $Prem(0,T;S) = Prot(0,T)$. In turn, given a CDS premium and a recovery rate, implied default intensity parameters can be found as the solution to the same equation.

DJ iTraxx index tranches	0-3%	3-6%	6-9%	9-12%	12-22%	RMSE	Fitted parameters
Market mid price	25.5%	146.0	60.3	36.3	19.3		
Bid/ask spread	1.3%	10.0	5.5	5.5	3.5		
Jump-diffusion intensities	25.0%	145.1	58.6	38.1	17.7	0.34	$\overline{\kappa} = 0.17$, $\overline{\sigma} = 0.058$, $\bar{l} = 0.017$, $\overline{\mu} = 0.066$, $\overline{w} = 0.83$
Pure diffusion intensities	30.0%	187.1	27.4	3.5	0.1	5.11	$\overline{\kappa} = 0.021$, $\overline{\sigma} = 0.064$, $\overline{w} = 1.00$
Gaussian copula	27.4%	222.3	52.5	13.8	1.6	4.58	$\rho = 0.145$
RFL Gaussian copula	25.3%	148.9	52.4	43.4	17.9	0.90	$\nu = -1.92$, $\rho_1 = 0.303$, $\rho_2 = 0.093$
Double-t copula	24.0%	153.4	56.5	32.4	16.4	0.84	$d = 4$, $\rho = 0.262$

Exhibit 3: Market and model prices for the DJ iTraxx 5-year index tranches on August 23, 2004. The market prices were obtained from Bloomberg and are quoted as running premiums in basis points, except for equity tranche prices which are upfront premiums. All the models have been fitted to the market prices by minimizing the root mean square errors relative to bid/ask spreads, RMSE as defined in (8). Interest rates are from the EUR swap curve, and the recovery rate is 35%.

DJ CDX index tranches	0-3%	3-7%	7-10%	10-15%	15-30%	RMSE	Fitted parameters
Market mid price	40.0%	312.5	122.5	42.5	12.5		
Bid/ask spread	2.0%	15.0	7.0	7.0	3.0		
Jump-diffusion intensities	46.9%	340.2	119.7	61.9	14.3	2.17	$\overline{\kappa} = 0.079$, $\overline{\sigma} = 0.058$, $\bar{l} = 0.048$, $\overline{\mu} = 0.047$, $\overline{w} = 0.77$
Pure diffusion intensities	49.3%	442.9	94.9	16.8	0.4	5.34	$\overline{\kappa} = 0.004$, $\overline{\sigma} = 0.085$, $\overline{w} = 1.00$
Gaussian copula	46.8%	474.4	131.8	36.9	2.9	5.30	$\rho = 0.167$
RFL Gaussian copula	48.6%	334.9	125.5	66.5	9.2	2.59	$\nu = -1.37$, $\rho_1 = 0.247$, $\rho_2 = 0.067$
Double-t copula	45.1%	367.0	114.9	54.9	20.0	2.44	$d = 4$, $\rho = 0.265$

Exhibit 4: Market and model prices for the DJ CDX 5-year index tranches on August 23, 2004. The market prices were obtained from Bloomberg and are quoted as running premiums in basis points, except for equity tranche prices which are upfront premiums. All the models have been fitted to the market prices by minimizing the root mean square errors relative to bid/ask spreads, RMSE as defined in (8). Interest rates are from the USD swap curve, and the recovery rate is 35%.

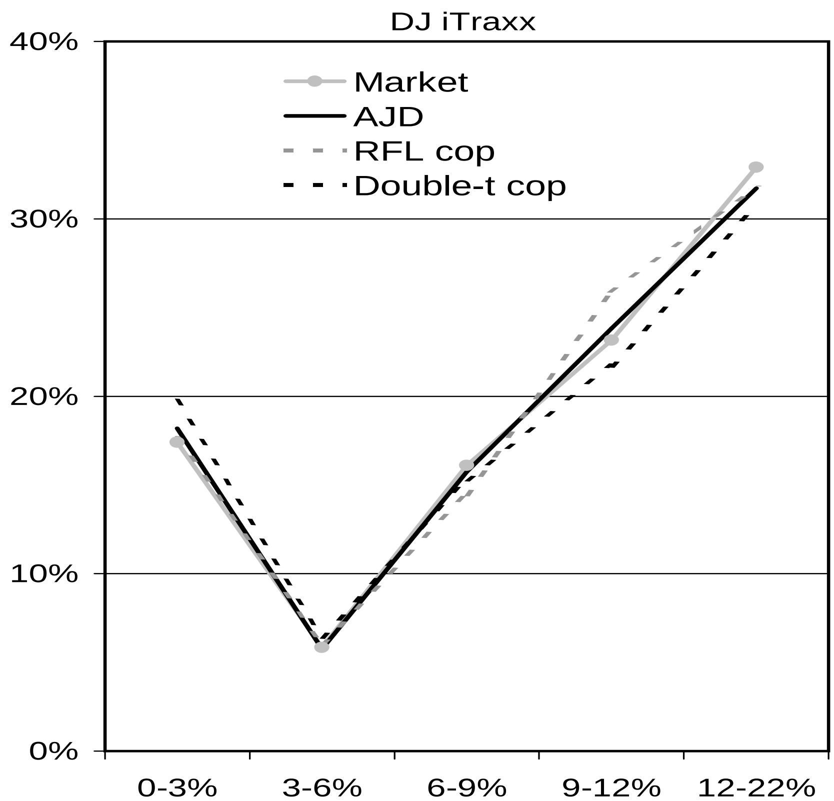

Exhibit 5: Implied correlations from market and model prices on the 5-year DJ iTraxx tranches. The market tranche prices were obtained from Bloomberg on August 23, 2004. The model tranche prices are from the fitted models with parameters as listed in Exhibit 3. The implied correlation is the Gaussian copula correlation parameter needed to match a given tranche price.

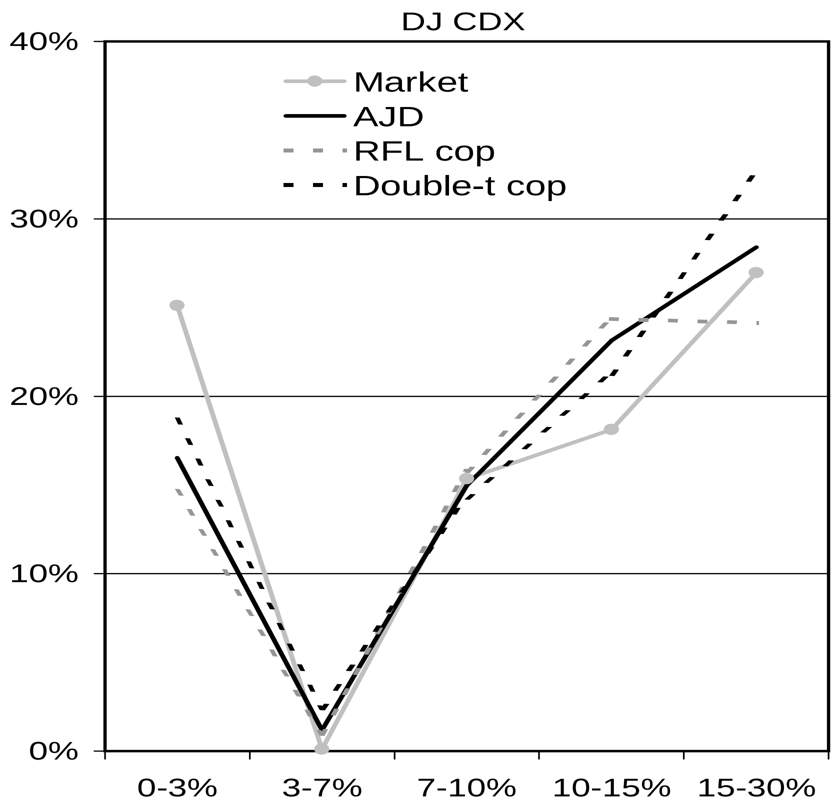

Exhibit 6: Implied correlations from market and model prices on the 5-year DJ CDX tranches. The market tranche prices were obtained from Bloomberg on August 23, 2004. The model tranche prices are from the fitted models with parameters as listed in Exhibit 4. The implied correlation is the Gaussian copula correlation parameter needed to match a given tranche price.

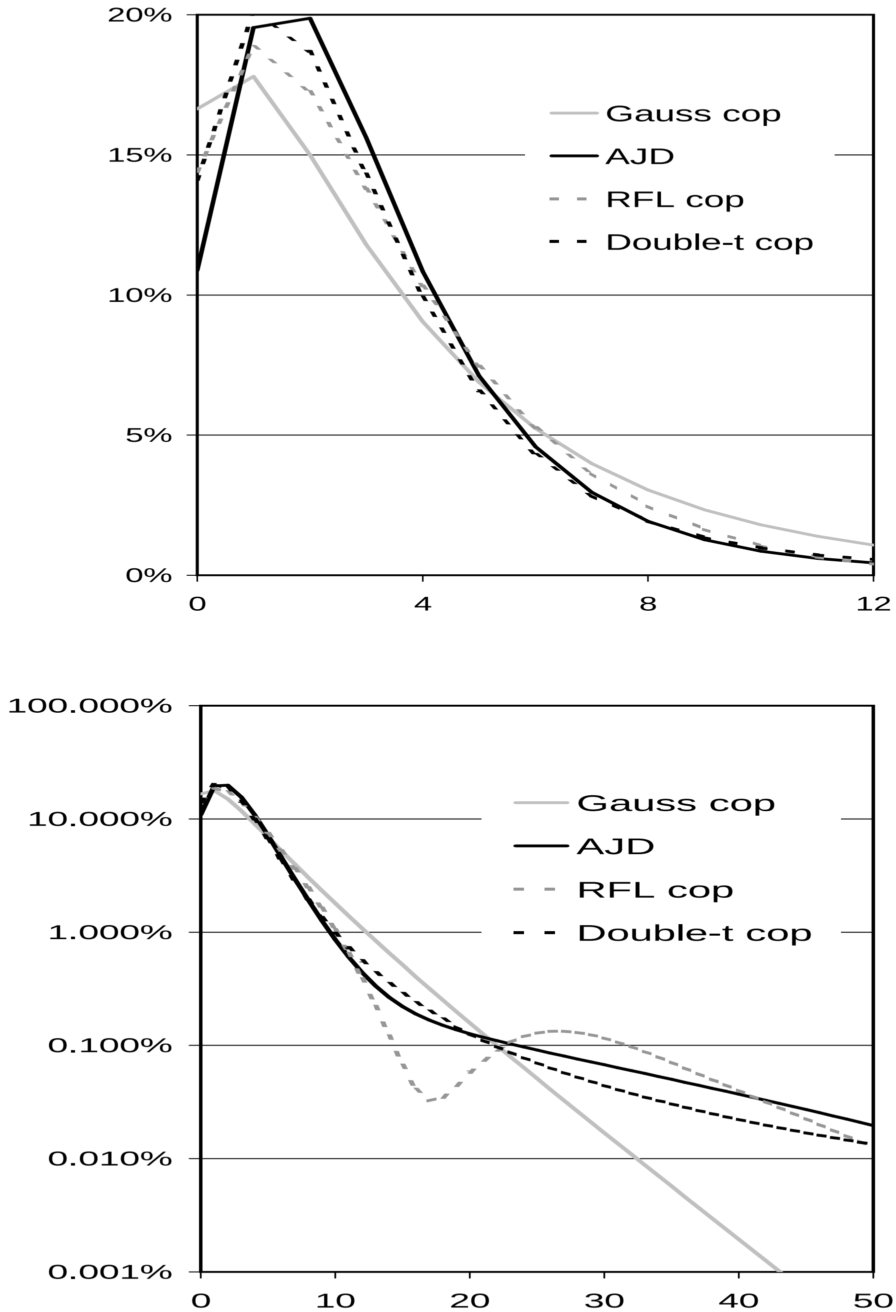

Exhibit 7: 5-year loss distributions for the DJ iTraxx pool. The graphs display model-implied probability functions for the number of defaults $0, 1, \ldots, 125$. The expected number of defaults is around 3.7 (expected portfolio loss 2.4%). The upper panel focuses on the low-loss outcomes, and the lower panel, with log-scale, shows the upper tail. The models have been calibrated to index tranche prices with parameters as listed in Exhibit 3.

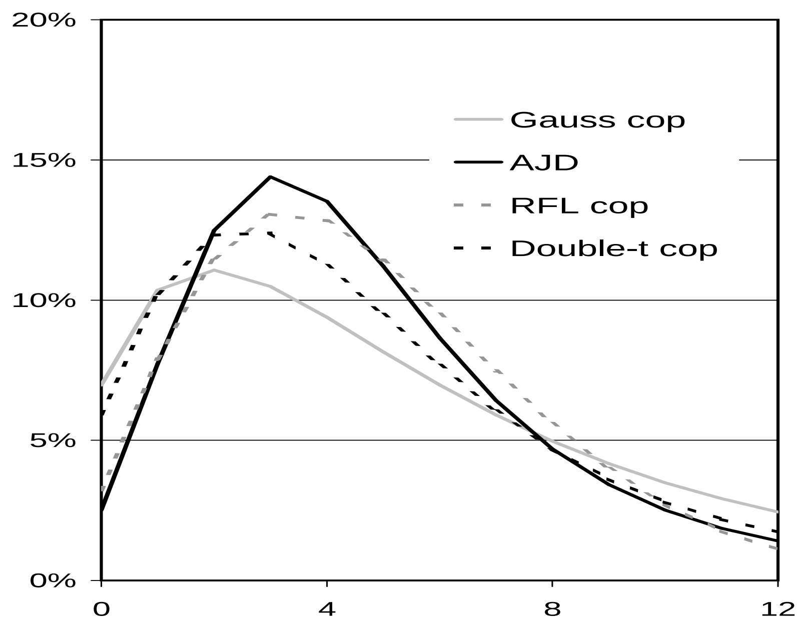

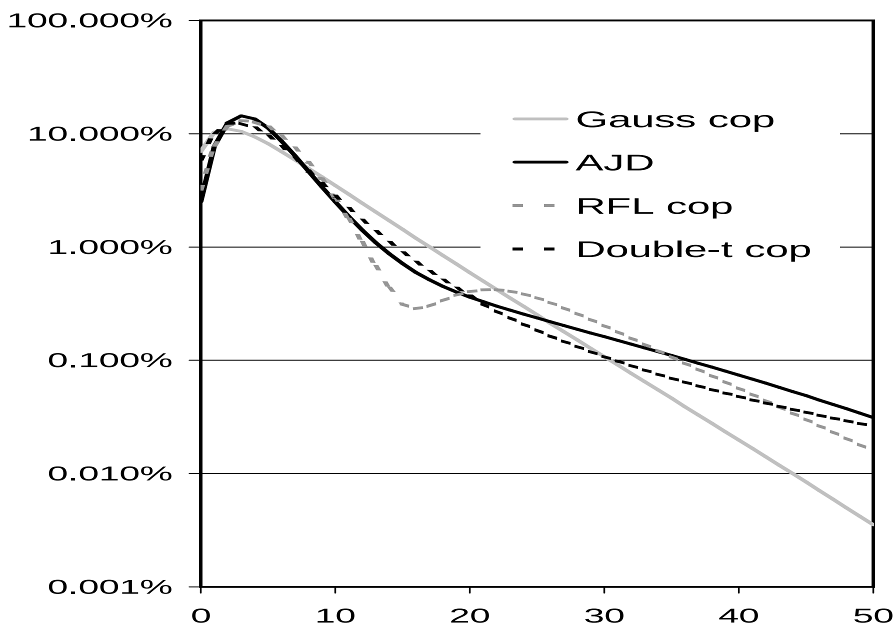

Exhibit 8: 5-year loss distributions for the DJ CDX pool. The graphs display model-implied probability functions for the number of defaults $0, 1, \ldots, 125$. The expected number of defaults is around 6.3 (expected portfolio loss 4.1%). The upper panel focuses on the low-loss outcomes, and the lower panel, with log-scale, shows the upper tail. The models have been calibrated to index tranche prices with parameters as listed in Exhibit 4.

DJ iTraxx index tranches	0-3%	3-6%	6-9%	9-12%	12-22%	RMSE	Fitted parameters
Market mid price	26.3%	80.6	23.1	10.3	5.8		
Bid/ask spread	0.6%	3.3	2.6	2.0	1.3		
Jump-diffusion intensities	28.7%	86.3	18.7	14.4	10.4	2.88	$\overline{\kappa} = 0.002$, $\overline{\sigma} = 0.045$, $\bar{l} = 0.004$, $\overline{\mu} = 0.187$, $\overline{w} = 0.60$
Pure diffusion intensities	32.5%	104.3	8.9	0.8	0.0	6.99	$\overline{\kappa} = 0.001$, $\overline{\sigma} = 0.065$, $\overline{w} = 0.38$
Gaussian copula	34.6%	99.9	2.9	0.1	0.0	8.44	$\rho = 0.035$
RFL Gaussian copula	27.0%	83.2	9.4	7.4	7.3	2.54	$\nu = -2.58$, $\rho_1 = 1.000$, $\rho_2 = 0.049$
Double-t copula	29.8%	101.1	24.4	13.2	6.6	3.99	$d = 3$, $\rho = 0.140$

Exhibit 9: Market and model prices for the DJ iTraxx 5-year index tranches on December 5, 2005. The market prices were obtained from Bloomberg and are quoted as running premiums in basis points, except for equity tranche prices which are upfront premiums. All the models have been fitted to the market prices by minimizing the root mean square errors relative to bid/ask spreads, RMSE as defined in (8). Interest rates are from the EUR swap curve, and the recovery rate is 35%.

DJ CDX index tranches	0-3%	3-7%	7-10%	10-15%	15-30%	RMSE	Fitted parameters
Market mid price	41.1%	117.5	32.9	15.8	7.9		
Bid/ask spread	0.8%	6.8	5.3	3.0	1.0		
Jump-diffusion intensities	43.2%	125.9	30.6	21.3	8.8	1.58	$\overline{\kappa} = 0.002$, $\overline{\sigma} = 0.024$, $\bar{l} = 0.012$, $\overline{\mu} = 0.088$, $\overline{w} = 0.52$
Pure diffusion intensities	47.3%	148.4	1.2	0.0	0.0	6.43	$\overline{\kappa} = 0.001$, $\overline{\sigma} = 0.036$, $\overline{w} = 0.48$
Gaussian copula	49.0%	154.4	1.1	0.0	0.0	7.08	$\rho = 0.018$
RFL Gaussian copula	43.3%	126.6	27.5	25.5	8.8	2.09	$\nu = -2.11$, $\rho_1 = 0.294$, $\rho_2 = 0.018$
Double-t copula	41.9%	176.1	32.3	15.3	6.0	4.00	$d = 3$, $\rho = 0.136$

Exhibit 10: Market and model prices for the DJ CDX 5-year index tranches on December 5, 2005. The market prices were obtained from Bloomberg and are quoted as running premiums in basis points, except for equity tranche prices which are upfront premiums. All the models have been fitted to the market prices by minimizing the root mean square errors relative to bid/ask spreads, RMSE as defined in (8). Interest rates are from the USD swap curve, and the recovery rate is 35%.

Endnotes

This work is based on Chapter 4 of the author's PhD dissertation, done at Copenhagen Business School. Ideas and views expressed in the paper do not in any way represent those of Goldman Sachs International. The author thanks Jens Christensen, Peter Feldhütter, David Lando, Niels Rom-Poulsen, Philipp Schönbucher, Søren Willemann, the Editor and an anonymous referee for useful discussions and comments.

[1] See e.g. Schönbucher [2003] for an introduction to copulas in credit risk modeling.

[2] In one of the first copula applications to portfolio credit risk, Li [2000] demonstrated that a one-period Gaussian copula with asset correlations as correlation parameters is equivalent to a multivariate Merton [1974] model. Due to this link, the copula correlation parameters are often interpreted as asset correlations. This is however only valid in a static one-period model, yet the Gaussian copula is often applied for modeling term structures of correlated default times. In the dynamic setting, a proper structural default definition is a barrier-hitting event along the lines of e.g. Black and Cox [1976], and not a sequence of Merton models. This is not just a theoretical concern – the dynamic loss properties implicit in these models are often very unrealistic.

[3] Intensity-based credit risk models have traditionally also been termed *reduced-form models*. Copula models are, however, even more reduced (the marginal default probabilities are specified in reduced form and the dependence structure is exogenously imposed), and therefore this label is not used here.

[4] Important empirical applications of intensity-based models include Duffee [1999] and Driessen [2004] on corporate bonds as well as Longstaff, Mithal and Neis [2004] on CDSs.

[5] Loss distributions for heterogeneous portfolio weights or recovery rates can also be computed semi-analytically in the intensity model. The recursive scheme used to build the loss distributions (to follow in equation (5)) can be refined to

cover discrete loss units instead of the number of defaults, as demonstrated in equation (10) of Andersen, Sidenius and Basu [2003] within the copula framework. The different loss amounts across the underlying names are then approximated by multiples of a chosen loss unit. The fact that recovery rates empirically are negatively correlated with default rates (see e.g. Hamilton et al. [2004]) is out of scope for this paper, but as in copula models this could be incorporated, at the cost of additional complication and computation time, by imposing some dependence of recovery rates on the common factor. However, numerical examples in Andersen and Sidenius [2004-5] suggest that stochastic recovery rates are not an essential modeling feature in fitting CDO market prices.

[6] Recent empirical evidence in Das et al. [2004] documents that high grade companies tend to have higher default probability correlations than low grade companies, but there are large variations in the correlation levels across companies within a given credit quality. This is captured by the firm-specific constants, a_i.

[7] See Černy [2004] for an introduction to FFT methods in derivatives pricing. Routines are available in many software packages and e.g. in Press et al. [1992].

[8] The Laplace transform of an integrated AJD does belong to the class of transforms covered in Duffie, Pan and Singleton [2000] and explicit ODE solutions are available (see Appendix A.1). The density function could then be obtained, using Mellin's inversion formula, by integration of the Laplace transform along a straight line in the complex plane. The density function is, however, more easily obtained by inversion of the characteristic function rather than the Laplace transform.

[9] This point can be illustrated in the following small example, which is deliberately simple and extreme to make the effect clearer. Consider a 5-year single-name CDS written on a very risky firm with a 5-year default probability of 80%. Suppose default occurs with a constant hazard rate λ, which must then be 32.19% ($= -\log(1-0.8)/5$). Assume the recovery rate is $\delta = 50\%$, interest rates are zero and the notional is 100. The value of the protection leg is then 40

(= $100(1-\delta)0.8$). Suppose a protection buyer has the choice between paying an up-front fee or a running premium. If the protection buyer pays 40 up-front and no running premium, her net profit is 10 in case of default before maturity and -40 otherwise. If the protection buyer instead chooses to pay no up-front fee and only running premium, the fair annual CDS premium is 1609bps (= $(1-\delta)\lambda$). Her net profit is then $50 - 16.09\tau$ in case of default before maturity and -80.47 otherwise (5 years of 16.09). The standard deviation of net profits is then 48.7 (the default time has exponential distribution), which could be reduced to 20 with an up-front fee. Also, the range of possible net profits can be narrowed from [-80.47, 50] to [-10, 40] with an up-front fee. Furthermore, it is clear that if the protection buyer only wants exposure to the event of default or not at some horizon and not to the timing of default, an up-front premium is preferred. In this example, buying protection without an upfront-fee can lead to a net loss even with default before maturity. If default occurs sufficiently late in the contract (after 3.11 years), the aggregate premium payments dominate the loss compensation.

[10] Alternatively the underlying credits could be calibrated to market prices of the CDS indices. The indices do not trade exactly at, though usually fairly close to, their intrinsic values given the individual credit curves.

[11] See e.g. the average CDS statistics in Table 2 of Longstaff, Mithal and Neis [2004].

[12] Willemann [2004] discusses some problems with base correlations.

References

Andersen, L. and J. Sidenius. "Extensions to the Gaussian Copula: Random Recovery and Random Factor Loadings," *Journal of Credit Risk*, 1 (2004-5), pp. 29-70

Andersen, L., J. Sidenius, and S. Basu. "All your hedges in one basket." *Risk*, November (2003), pp. 67-72.

Black, F. and J. Cox. "Valuing Corporate Securities: Some Effects of Bond Indenture Provisions." *Journal of Finance*, 31 (1976), pp. 351-367.

Černy, A. "Introduction to Fast Fourier Transform in Finance." *Journal of Derivatives*, 12 (2004), pp. 73-88.

Das, S., L. Freed, G. Geng, and N. Kapadia. "Correlated Default Risk." Working Paper, Santa Clara University, 2004.

Driessen, J. "Is Default Event Risk Priced in Corporate Bonds?" Working Paper, University of Amsterdam, 2004; forthcoming in *Review of Financial Studies.*

Duffee, G. "Estimating the Price of Default Risk." *Review of Financial Studies*, 12 (1999), pp. 197-226.

Duffie, D. "Credit Swap Valuation." *Financial Analysts Journal*, 55 (1999), pp. 73-87.

Duffie, D. "Time to adapt copula methods for modelling credit risk correlation." Comment in *Risk*, April (2004).

Duffie, D., D. Filipović, and W. Schachermayer. "Affine Processes and Applications in Finance." *Annals of Applied Probability*, 13 (2003), pp. 984-1053.

Duffie, D. and N. Gârleanu. "Risk and Valuation of Collateralized Debt Obligations." *Financial Analysts Journal*, 57 (2001), pp. 41-59.

Duffie, D., J. Pan, and K. Singleton. "Transform Analysis and Asset Pricing for Affine Jump Diffusions." *Econometrica*, 68 (2000), pp. 1343-1376.

Duffie, D. and K. Singleton. "Modeling Term Structures of Defaultable Bonds." *Review of Financial Studies*, 12 (1999), pp. 687-720.

Duffie, D. and K. Singleton. "*Credit Risk: Pricing, Measurement and Management.*" Princeton University Press, 2003.

Gregory, J. and J.-P. Laurent. "I Will Survive." *Risk*, June (2003), pp. 103-107.

Hamilton, D., P. Varma, S. Ou, and R. Cantor. "Default & Recovery Rates of Corporate Bond Issuers." Moody's Investors Service, New York, January 2004.

Hull, J. and A. White. "Valuing Credit Default Swaps I: No counterparty default risk." *Journal of Derivatives*, 8 (2000), pp. 29-40.

Hull, J. and A. White. "Valuation of a CDO and an n'th to Default CDS

Without Monte Carlo Simulation." *Journal of Derivatives*, 12 (2004), pp. 8-23.

Jarrow, R. and S. Turnbull. "Pricing Options on Financial Securities Subject to Credit Risk." *Journal of Finance*, 50 (1995), pp. 53-85.

Lando, D. "Three Essays on Contingent Claims Pricing." PhD Dissertation, Cornell University, 1994.

Lando, D. "On Cox processes and Credit Risky Securities." *Review of Derivatives Research*, 2 (1998), pp. 99-120.

Li, D. X. "On Default Correlation: A Copula Function Approach." *Journal of Fixed Income*, 9 (2000), pp. 43-54.

Longstaff, F. A., S. Mithal, and E. Neis. "Corporate Yield Spreads: Default Risk or Liquidity? New Evidence from the Credit-Default Swap Market." Working Paper, UCLA, 2004; forthcoming in *Journal of Finance.*

Merton, R. C. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates." *Journal of Finance*, 29 (1974), pp. 449-470.

Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. Flannery. "*Numerical Recipes in C: The Art of Scientific Computing.*" Cambridge University Press, 2nd edition, 1992.

Schönbucher, P. "Term Structure Modelling of Defaultable Bonds." *Review of Derivatives Research*, 2 (1998), pp. 161-192.

Schönbucher, P. "*Credit Derivatives Pricing Models: Models, Pricing and Implementation.*" John Wiley and Sons Ltd, 2003.

Schönbucher, P. and D. Schubert. "Copula-Dependent Default Risk in Intensity Models." Working Paper, Bonn University, 2001.

Willemann, S. "An Evaluation of the Base Correlation Framework for Synthetic CDOs." Working Paper, Aarhus School of Business, 2004.