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

Allan Mortensen†

This version: January 31, 2005

Abstract

This paper presents a semi-analytical valuation method for basket credit derivatives in a flexible intensity-based model. Default intensities are modelled as heterogeneous, 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.

JEL classification: G13
Keywords: credit derivatives, CDOs, default correlation, intensity-based models, affine jump-diffusions

*I am grateful to Jens Christensen, Peter Feldhütter, David Lando, Niels Rom-Poulsen and Søren Willemann for useful discussions and comments. All errors are of course my own.

†Department of Finance, Copenhagen Business School, am.fi@cbs.dk

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 modelling default time correlation given the marginal default probabilities.[1] Factor copulas have become extremely popular in practice, mainly due to the tractability of the semi-analytical methods 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 usually very sensitive to the 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 the CDS index options launched recently.

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

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

[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 modelling 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. Extensive calculations (not reported) indicate that, for a given set of marginal default probabilities, the Gaussian copula using asset correlations tends to overestimate default correlation slightly relative to a multivariate dynamic structural model.

[3] Intensity-based credit risk models have traditionally 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.

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 modelling. 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 modelled 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 quotes often vary by several hundred basis points within the benchmark CDS indices.[5]

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

[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]Besides that, standard implementations of the analytical solution in Duffie and Gârleanu (2001) may be numerically instable when the pool is large (more than, say, 50 entities), which is problematic since pools in practice often consist of 100 names or more. The analytical method proposed here is numerically stable for large pools.

prices of benchmark basket credit derivatives reasonably well. While the attainable levels of default correlation in an intensity-based model are limited, the model is able to produce correlations consistent with the market-implied levels, at least for the large CDO pools considered in this study. For small pools, default correlations in the model may be too low for some cases. 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 reasonably well. That is, the model is able to generate pricing patterns fairly 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) in a Gaussian copula with stochastic 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 comparable with that of the stochastic factor loading copula but not quite as good as that of the t copula.

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 but do not seem excessive. 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 well-suited for the pricing of 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, needed for calibration to the single-name CDSs, 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. Section 3 presents the empirical application to basket credit derivatives pricing, and Section 4 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 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 recovery rate, $\delta \in [0, 1)$, is assumed throughout this paper, and therefore the focus is on the number of defaults, which is denoted by $D_t = \sum_{i=1}^{N} D_i(t)$.[6]

Throughout the paper, everything is done under the probability measure $\mathbb{Q}$, which could in principle be any probability measure. For the pricing analysis in Section 3, $\mathbb{Q}$ is assumed to be the risk-neutral measure. For risk management, $\mathbb{Q}$ could be taken as the real-world measure.

2.1 The single-name setting

In the intensity-based credit risk model, default is defined as the first jump of a pure jump process, and it is assumed that the jump process has an intensity process. More formally, it is assumed that a non-negative process λ exists such that the process

$$M(t) := \mathbf{1}_{\{\tau \leqslant t\}} - \int_0^t \mathbf{1}_{\{\tau > s\}} \lambda_s \, ds \tag{1}$$

is a martingale.

In this setting, default is an unpredictable event (an inaccessible stopping time), and the probability of default over a small interval of time (in the limiting sense) is proportional to the intensity

$$\lim_{\Delta \to 0} Pr\big(\tau \leqslant t + \Delta \big| \tau > t\big) = \lambda_t \Delta \tag{2}$$

These 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. For 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

$$Pr\big(\tau \leqslant t \big| \{\lambda_s\}_{0 \leqslant s \leqslant t}\big) = 1 - e^{-\int_0^t \lambda_s ds} \tag{3}$$

and the unconditional default probabilities by

$$Pr\big(\tau \leqslant t\big) = 1 - E\big[e^{-\int_0^t \lambda_s ds}\big] \tag{4}$$

[6]Loss distributions for heterogeneous recovery rates or face values can still be computed semi-analytically in the model using a discrete bucketing procedure (a refinement of the recursion in (14)) along the lines of Andersen, Sidenius, and Basu (2003) or Hull and White (2004). 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 it could be incorporated, at the cost of additional complication and computation time, by imposing some dependence of recovery rates on the common factor. Numerical examples in Andersen and Sidenius (2004) suggest that stochastic recovery rates are not an essential modelling feature in fitting CDO market prices.

As can be seen from the last expression, default risk modelling in this framework is mathematically equivalent to interest rate modelling. Therefore, many well-known and useful techniques from this field can be applied for different specifications of the default intensity. A particularly tractable and flexible family of models is the class of affine jump-diffusions, characterized and analyzed by Duffie, Pan, and Singleton (2000). This paper applies the affine class in a multi-name setting.

2.2 The multi-name model

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

$$\lambda_{i,t} = \alpha_i x_t + x_{i,t} \tag{5}$$

where $\alpha_i > 0$ is a constant and the two processes, x and x_i, are independent affine processes. As shown 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 (5) is a very simple modification of the specification in Duffie and Gârleanu (2001), which is the special case of $\alpha_i = 1$. To see why the modification is relevant, consider a heterogeneous pool in the case $\alpha_i = 1$ 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 x). Also, the firms of the highest credit quality will typically have relatively high correlation with the common factor. On the contrary, (5) imposes no implicit constraint on the combination of credit quality and correlation.[7]

More specifically, suppose the common component follows

$$dx_t = \kappa(\theta - x_t)dt + \sigma\sqrt{x_t}dW_t + dJ_t \tag{6}$$

and the idiosyncratic components follow

$$dx_{i,t} = \kappa_i(\theta_i - x_{i,t})dt + \sigma_i\sqrt{x_{i,t}}dW_{i,t} + dJ_{i,t} \tag{7}$$

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

- $W, W_1, \ldots, W_N$ are independent Wiener processes.

[7]Recent empirical evidence in Das et al. (2004) documents that high grade companies tend to have higher default 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, α_i.

- $J, J_1, \ldots, J_N$ are independent pure-jump processes, independent of the Wiener processes.
- the jump times are those of a series of Poisson processes with intensities $l, l_1, \ldots, l_N$.
- the jump sizes are exponentially distributed with means $\mu, \mu_1, \ldots, \mu_N$, independent of the jump times.

The restriction to positive jumps seems reasonable since most large jumps in credit spreads are positive. Given the square root volatility structure and the positive jumps, the components – and thereby the default intensities – are strictly positive under the parameter restrictions $2\kappa\theta \geqslant \sigma^2$ and $2\kappa_i\theta_i \geqslant \sigma_i^2$.

These processes are *basic affine processes* (AJD), as defined by Duffie and Gârleanu (2001). Appendix A provides a few useful results for AJDs, and using the notation introduced in the appendix, the processes under consideration can, in short, be denoted as

$$\begin{aligned} x &\sim AJD(x_0, \kappa, \theta, \sigma, l, \mu) \\ \alpha_i x &\sim AJD(\alpha_i x_0, \kappa, \alpha_i\theta, \sqrt{\alpha_i}\sigma, l, \alpha_i\mu) \\ x_i &\sim AJD(x_{i,0}, \kappa_i, \theta_i, \sigma_i, l_i, \mu_i) \end{aligned} \tag{8}$$

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.

The tractable specifications of x and x_i are needed to allow for a semi-analytical solution of the portfolio loss distribution. 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 but that is not needed. The independence of the two components automatically ensures that the marginal default probabilities – needed for the calibration to the single-name CDS market – are known in closed form,

$$\begin{aligned} Pr\big(\tau_i \leqslant t\big) &= 1 - E\big[e^{-\int_0^t \lambda_{i,s} ds}\big] \\ &= 1 - E\big[e^{-\alpha_i \int_0^t x_s ds}\big] \times E\big[e^{-\int_0^t x_{i,s} ds}\big] \\ &= 1 - e^{A(t;\,\kappa,\alpha_i\theta,\sqrt{\alpha_i}\sigma,l,\alpha_i\mu)+B(t;\,\kappa,\sqrt{\alpha_i}\sigma)\alpha_i x_0+A(t;\,\kappa_i,\theta_i,\sigma_i,l_i,\mu_i)+B(t;\,\kappa_i,\sigma_i)x_{i,0}} \end{aligned} \tag{9}$$

where A and B are deterministic functions given explicitly in (30) in Appendix A.1.

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, which can be done in a reasonable way without loosing the flexibility to fit the most important name-specific quantities: CDS level, slope, volatility and correlation. Section 3.2 presents an empirical application of a very 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,

$$Z_t := \int_0^t x_s ds \tag{10}$$

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

$$\begin{aligned} p_i(t|z) := Pr\big(\tau_i \leqslant t \big| Z_t = z\big) &= 1 - e^{-\alpha_i z} E\big[e^{-\int_0^t x_{i,s} ds}\big] \\ &= 1 - e^{-\alpha_i z + A(t;\kappa_i,\theta_i,\sigma_i,l_i,\mu_i) + B(t;\kappa_i,\sigma_i) x_{i,0}} \end{aligned} \tag{11}$$

again with A and B given explicitly in Appendix A.1.

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

$$Pr\big(D_t = j\big) = \int_{-\infty}^{\infty} Pr\big(D_t = j \big| Z_t = z\big)\, f_{Z_t}(z)\, dz \tag{12}$$

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 relatively smooth function of the integrator.

Given the common factor, joint default probabilities can be calculated from the marginal default probabilities, (11), both with homogeneous and heterogeneous credit qualities.

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

$$Pr\big(D_t = j \big| Z_t = z\big) = \binom{N}{j} p_i(t|z)^j \big(1 - p_i(t|z)\big)^{N-j} \tag{13}$$

and the loss distribution, (12), 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_t^K 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

$$\begin{aligned} Pr\big(D_t^K = j \big| Z_t = z\big) = {} & Pr\big(D_t^{K-1} = j \big| Z_t = z\big) \times \big(1 - p_K(t|z)\big) \\ & + Pr\big(D_t^{K-1} = j-1 \big| Z_t = z\big) \times p_K(t|z) \end{aligned} \tag{14}$$

for $j = 1, \ldots, K$. For $j = 0$ the last term obviously disappears. The recursion starts from $Pr(D_t^0 = j|Z_t) = \mathbf{1}_{\{j=0\}}$ and runs for $K = 1, \ldots, N$ with $Pr(D_t = j|Z_t) = Pr(D_t^N = j|Z_t)$. 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

$$\varphi_{Z_t}(u) := E[e^{iuZ_t}] = \int_{-\infty}^{\infty} e^{iuz} f_{Z_t}(z)\, dz \tag{15}$$

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

$$f_{Z_t}(z) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-iuz} \varphi_{Z_t}(u)\, du \tag{16}$$

which can be computed very efficiently using fast Fourier transform (FFT) methods.[8]

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.2 proves that their result can be extended to cover this case as well.[9] As shown, the characteristic function is given by the exponential affine form

$$\varphi_{Z_t}(u) = e^{\tilde{A}(0;t,u,\kappa,\theta,\sigma,l,\mu)+\tilde{B}(0;t,u,\kappa,\sigma)x_0} \tag{17}$$

where $\tilde{A}$ and $\tilde{B}$ are complex-valued deterministic functions solving the ODEs in (32). 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.

[8] 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).

[9] 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.

Relative to the semi-analytical factor copula solution in Andersen, Sidenius, and Basu (2003), the only added computational complexity in the intensity-based model is that the distribution of the common factor is more involved. If intensities were modelled 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 Basket credit derivatives valuation

This section applies the intensity-based model for valuation of synthetic CDOs. After a short description of the product, we shall test how the model conforms with market prices.

3.1 Synthetic CDOs

In a CDO, the credit risk of an underlying pool of bonds, loans or CDSs is passed through to a number of tranches according to some prioritization scheme. The structure is referred to as a *synthetic* CDO when 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.[10] 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

[10] The terminology is similar to the single-name market, where the credit risk in buying a corporate bond roughly corresponds to selling a credit default swap.

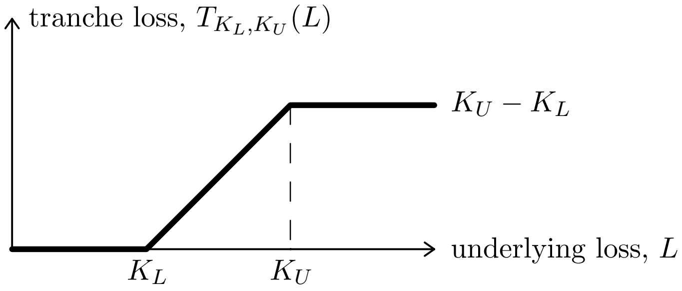

Figure 1: Index tranche loss. The loss on a tranche with attachment point K_L and detachment point K_U as a function of underlying portfolio loss.

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.

3.1.1 Index tranche cash flows

Consider an index tranche covering a part of the portfolio loss distribution from a lower attachment point, K_L, to an upper detachment point, K_U. The tranche loss as a function of portfolio loss L is defined as

$$T_{K_L,K_U}(L) := \max\{L - K_L, 0\} - \max\{L - K_U, 0\} \tag{18}$$

which has the familiar structure of a call spread option written on the portfolio loss, see Figure 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 premium payments in arrears with a frequency of f per year. Typical contracts specify quarterly payments, $f = 4$, which is the case for the contracts considered in this paper. Denote the annual premium by S and the payment dates by $t_j = j/f$, for $j = 1, \ldots, Tf$. 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, $K_U - K_L$. The value of the protection leg is then

$$Prot(0, T) = E\Big[\int_0^T e^{-\int_0^t r_s ds} dT_{K_L,K_U}(L_t)\Big] \tag{19}$$

With quarterly premium payments, as in this paper, 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,

$$Prot(0,T) = \sum_{j=1}^{Tf} P\big(0, t_j - \frac{1}{2f}\big)\Big\{E\big[T_{K_L,K_U}(L_{t_j})\big] - E\big[T_{K_L,K_U}(L_{t_{j-1}})\big]\Big\} \qquad (20)$$

The value of the premium leg is

$$Prem(0,T;S) = E\Big[\sum_{j=1}^{Tf} e^{-\int_0^{t_j} r_s ds}\frac{S}{f}\int_{t_{j-1}}^{t_j}\frac{K_U - K_L - T_{K_L,K_U}(L_s)}{t_j - t_{j-1}}ds\Big] \qquad (21)$$

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 similar to those for the protection leg we get

$$\begin{aligned} Prem(0,T;S) = \frac{S}{f}\sum_{j=1}^{Tf} P\big(0,t_j\big) \qquad (22) \\ \times \Big\{K_U - K_L - \frac{1}{2}\Big(E^Q[T_{K_L,K_U}(L_{t_{j-1}})] + E^Q[T_{K_L,K_U}(L_{t_j})]\Big)\Big\} \end{aligned}$$

The fair tranche premium 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 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).[11] The up-front fee is quoted as a fraction, y, of tranche notional,

[11] 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% ($= -\log(1-0.8)/5$) to match the default probability. 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 (using the fact that the default time is exponentially distributed), 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 in the event of default prior to maturity. If default occurs sufficiently late in the contract (after 3.11 years), the premium payments dominate the loss compensation.

i.e.

$$
\begin{aligned}
&Prem_{eq}(0,T;0.05,y) = yK_U \\
&+ \frac{0.05}{f} \sum_{j=1}^{Tf} P(0,t_j)\Big\{K_U - \frac{1}{2}\Big(E^Q[T_{0,K_U}(L_{t_{j-1}})] + E^Q[T_{0,K_U}(L_{t_j})]\Big)\Big\}
\end{aligned} \tag{23}
$$

The fair equity tranche price is then defined as the up-front fee, y, 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 Tf loss distributions are needed, but very good price approximations could be achieved by interpolating expected tranche losses from a lower number of loss distributions.

3.2 Empirical application

As mentioned the general version of the model is very rich, and of course the chosen parametrization in any application of the model should reflect the data set at hand and the context of the application. In the following, a parsimonious version of the model is calibrated to a data set consisting of market quotes on synthetic CDOs and underlying CDSs for the most liquid 5-year maturity. The CDO quotes are available on 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. All quotes were obtained from Bloomberg, using the Bloomberg generic price source, as of August 23, 2004.

The 125 iTraxx CDS (mid) quotes range from 11bps to 127bps with average and median quotes at 39.1bps and 36.0bps, respectively. The index is quoted slightly below the average at 38.8bps (mid). The CDX pool is highly heterogeneous with quotes ranging from 18bps to 468bps, and an average and median of 67.1bps and 48.0bps, respectively. The CDX index is quoted substantially below the average at 61.5bps.

3.2.1 Calibration

A number of restrictions need to be imposed to arrive at a specification with a single name-specific parameter, which can be calibrated to the single name-specific input, and a single correlation parameter, which can be varied without affecting the marginal distributions.

Recall that $\lambda_{i,t} = \alpha_i x_t + x_{i,t}$ and

$$\begin{aligned} \alpha_i x &\sim AJD(\alpha_i x_0, \kappa, \alpha_i \theta, \sqrt{\alpha_i}\sigma, l, \alpha_i \mu) \\ x_i &\sim AJD(x_{i,0}, \kappa_i, \theta_i, \sigma_i, l_i, \mu_i) \end{aligned} \tag{24}$$

First of all, given only a single point on the CDS curves, the slopes are neutralized by initializing the processes at their mean-reversion levels, $x_0 = \theta$ and $x_{i,0} = \theta_i$. The CDS curves are then approximately flat – the implied 1- to 5-year CDS premiums are upward-sloping within 10bps for the relevant credit qualities.[12] If the model was calibrated to term structures of CDS spreads, instead of just the 5-year points, the initial values should be left among the free parameters in the fitting of the shape of the curves.

Next, we assume that the mean-reversion levels of the default intensities are common across entities apart from the scaling factors, i.e. $\alpha_i \overline{\theta}$ where $\overline{\theta}$ is a common constant. Also, we assume a common jump intensity of $\overline{l}$. These assumptions correspond to picking $\theta_i = \alpha_i(\overline{\theta} - \theta)$ and $l_i = \overline{l} - l$. Furthermore, we assume

$$\kappa_i = \kappa, \quad \sigma_i = \sqrt{\alpha_i}\sigma, \quad \mu_i = \alpha_i \mu \tag{25}$$

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

$$\lambda_i \sim AJD(\alpha_i \overline{\theta}, \kappa, \alpha_i \overline{\theta}, \sqrt{\alpha_i}\sigma, \overline{l}, \alpha_i \mu) \tag{26}$$

As in Duffie and Gârleanu (2001), it is assumed that the ratio of systematic to idiosyncratic jump intensity is identical to the ratio of systematic to idiosyncratic mean-reversion level. This allows default correlation to be controlled by a single correlation parameter,

$$w := \theta \,/\, \overline{\theta} = l \,/\, \overline{l} \in [0, 1] \tag{27}$$

representing the systematic part. Note that $w = 0$ implies independent intensity processes and $w = 1$ implies perfectly correlated intensity processes, which of course does not mean perfectly correlated default events since defaults still occur independently given the intensity processes. Importantly, the correlation parameter does not affect the marginal default intensities and can therefore be chosen freely after the model has been calibrated to single-name CDS quotes.

One final restriction is imposed because both $\overline{\theta}$ and the α_i's are level parameters of the credit spreads. A higher (lower) $\overline{\theta}$ would be offset by lower (higher) α_i's in the calibration, and therefore I propose to fix $\overline{\theta}$ such that the CDS index or average is matched for $\alpha_i = 1$. For a homogeneous pool, this implies $\alpha_i = 1$ for all entities. For a heterogeneous pool, the α_i's are spread around 1.

This parametrization leaves five free parameters: κ, σ, $\overline{l}$, μ and w. In summary, the procedure is the following. For each combination of the first four parameters:

[12] If instead the processes were initialized at their long-run means $x_0 = \theta + l\mu/\kappa$ and $x_{i,0} = \theta_i + l_i \mu_i / \kappa_i$, as proposed in Duffie and Gârleanu (2001), the CDS curves would be downward-sloping, which is counterfactual for most investment grade companies.

1. Calibrate $\overline{\theta}$ to fit the CDS index or average level.

2. Calibrate the α_i's to fit the heterogeneous CDS quotes. (If homogeneous CDS quotes are assumed, $\alpha_i = 1$ for all entities.)

3. Vary w.

The calibration to the single-name CDS market in the first two points is straightforward, but for completeness Appendix B provides a few details.

The model is fitted to the index tranche market prices by minimizing the root mean square price errors relative to bid/ask spreads,

$$\text{RMSE} = \sqrt{\frac{1}{5}\sum_{j=1}^{5}\Big(\frac{S_{j,\,market\;mid} - S_{j,\,model}}{S_{j,\,market\;ask} - S_{j,\,market\;bid}}\Big)^2} \tag{28}$$

where S_j is the spread of tranche j. The liquidity varies across tranches, and this way price errors on the most reliable prices get the highest weights.

In line with empirical levels, we assume a constant riskless interest rate of $r = 0.03$ and a recovery rate of $\delta = 0.4$. The CDS and index tranche spreads are both relatively insensitive to the level of riskless interest rates. Furthermore, a lower (higher) recovery rate would be offset by lower (higher) CDS-implied marginal default intensities, and since index tranche spreads are primarily driven by the loss rates (loss given default times default intensity) in the underlying pool, the parameter value for the recovery rate is, within reasonable bounds, not too critical. Consequently, the results are fairly robust to these two specifications.

3.2.2 Results

The model is calibrated both with the full list of heterogeneous CDS spreads and with homogeneous CDS spreads at the average level, which can be seen as an approximation with a simpler hypothetical pool.

The results for the iTraxx index are reported in Table 1. The calibrated intensity model fits the market prices very well with an error measure of 0.74 based on the heterogeneous CDS spreads. Surprisingly, the model fits slightly better (RMSE of 0.67) when the marginal distributions are calibrated to a homogeneous CDS spread (the pool average of 39.1bps). The number of free parameters is the same in both cases, but one might expect that incorporating the information in the 125 different CDS quotes would lead to a better fit to CDO prices. One reason why heterogeneity does not improve the fit to the iTraxx tranches may be that the 125 quotes are relatively homogeneous (11-127bps) – as we shall see, the heterogeneous model does fit better for the much more heterogeneous CDX index (18-468bps).

The jump parameters are relatively high but not too implausible, since the impact of a jump in the instantaneous default intensity is moderated by the mean-reversion tendency of the process. To understand the magnitude, consider a typical ($\alpha_i = 1$) company with the fitted parameter values $\kappa = 0.27$, $\sigma = 5\%$, $\bar{l} = 1.7\%$ and $\mu = 7.8\%$ (the heterogeneous case). With initial value and mean-reversion level of the default intensity at $\overline{\theta} = 0.46\%$, the 5-year CDS is priced at 39.1bps (the iTraxx average). A jump in the default intensity of 780bps (the mean jump size) increases the 5-year CDS spread to 307bps – i.e. the credit spread only jumps by 268bps (of course short maturity spreads jump more, around 400bps for a 1-year CDS). 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.[13] The observed 1-year frequencies for migrations from A to BB or worse (B or CCC) 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 around 91-93%, which implies that most of the variations in the default intensities are common fluctuations. The correlation parameter does seem very high and higher than expected considering the degree of co-movement observed in CDS premiums across time. As illustrated in the table, restricting the correlation parameter at a lower value (for example 70%) could however still fit prices reasonably well with slightly higher jump parameters. The table also shows that pure diffusion intensities, as expected, generate too low default correlation – too high equity spreads and too low senior spreads.

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 be helpful in forming opinions on the absolute pricing of correlation products (provided that the model is considered reasonable).

For comparison, the table also reports the fit of three copula models: (i) the 1-factor Gaussian copula, (ii) the Gaussian copula with stochastic factor loadings proposed by Andersen and Sidenius (2004), and (iii) the double-t copula of Hull and White (2004). A very brief outline of the three copula models is provided for completeness in Appendix C. For more details, refer to the individual papers.

As expected, 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 most senior tranches, whereas much less correlation is needed to match the most junior mezzanine tranche (3-6%). The two alternative copulas fit market prices much better. The fit of the intensity model is slightly better than that of the copula with stochastic factor loadings, whereas the double-t copula provides the best fit to the market prices. The basket credit derivatives market is still very immature and it is difficult to rule out supply and demand effects caused by market segmentation or market inefficiency. Therefore, a perfect

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

fit to the market prices should perhaps not be expected, and the model with the best fit may not be the most appropriate one. As mentioned in the introduction, the properties of the models are very different.

The results for the CDX index are reported in Table 2. Compared with the iTraxx results, the AJD model is fitted with higher jump intensities but lower jump sizes. Again, to investigate the jump impact we consider a company with the fitted parameters $\kappa = 0.2$, $\sigma = 5.4\%$, $\bar{l} = 3.7\%$ and $\mu = 6.7\%$ (the heterogeneous case), for which a 5-year CDS premium of 67.1bps (the CDX average) is reached with $\bar{\theta} = 0.73\%$. A default intensity jump of 670bps causes an increase in the 5-year CDS premium to 210bps – a jump of 143bps, corresponding approximately to a downgrade by one and a half rating categories (A to $BBB-$ or $BB+$). An expected frequency of 3.7% for such a migration is again relatively high but certainly not unrealistic.

All the models show a worse fit to the CDX index tranches, but the relative performance of the models is about the same. This could be interpreted as evidence that the iTraxx market is more efficient than the CDX market. The fits of the intensity model and the stochastic factor loading copula are comparable, whereas the double-t copula still fits better. For the CDX index, incorporating the 125 heterogeneous CDS spreads improves the ability to match the CDO prices for all the models as we would expect. As mentioned, this was not the case for the iTraxx index which may be explained by the fact that the range of underlying spreads in the CDX index is much wider than the range of iTraxx spreads.

In conclusion, the ability of the intensity-based model to fit CDO market prices is comparable with alternative models proposed in the literature. The results should of course be interpreted with some caution since only a single trading day is considered. The interesting questions of performance and parameter stability over time of the various models are left for further research.

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.[14] For each tranche, the implied correlation is defined as the homogeneous correlation parameter that produces a standard model

[14] To complete the standard model, assumptions on recovery rates, marginal default probabilities and interest rates are needed in addition to the correlation specification. As mentioned in Finger (2005), complete agreement with respect to a full specification of the standard model has not yet been reached across all market participants. The standard model of this paper is, as in Hull and White (2004), based on marginal default probabilities derived from constant and homogeneous default intensities. Furthermore, as in the rest of the paper, $r = 3\%$ and $\delta = 40\%$.

price identical to the market price.[15] If the Gaussian copula was the true model, the implied correlations would be identical across tranches on the same underlying pool. As we saw in the calibration, in practice they are not, and a very significant so-called correlation skew is observed.

Figure 2 illustrates the correlation skews generated by the calibrated AJD model. The model-implied correlations are fairly close to the market-implied correlations – especially for the iTraxx index. Although this is just another way of representing tranche prices, minimizing the price errors is not necessarily the same objective as getting the nicest fit to the correlation skew. Figure 3 shows similar results for the Gaussian copula with stochastic factor loadings and the double-t copula.

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 implied correlation skew. Figures 4 and 5 illustrate the loss distributions for the two index pools, and the patterns are very similar although the iTraxx CDO tranches have been fitted much 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: the market assigns lower probability to zero losses, higher probability to small-medium losses, lower probability to moderately high losses, and higher probability to very high losses.

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. The standard model tends to overvalue the spread on the second-loss mezzanine tranche (iTraxx 3-6%, CDX 3-7%). This indicates 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 and low probability of no defaults.

We also note from the lower panels of the figures that the stochastic factor loading copula, in the regime-switching version applied here, produces a bimodal loss distribution, and therefore the sensitivity of mezzanine prices to the parameters of the model can be very unpredictable.

[15] This is referred to as *compound correlations*. This quotation device, however, suffers from both existence and uniqueness problems, which are not encountered with Black-Scholes volatilities. Tranche spreads are not monotone in compound correlation (except for equity tranches), and we may observe arbitrage-free market prices that are not attainable by a choice of 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 spreads are monotone, but existence is still not guaranteed and they can be very difficult to interpret (Willemann (2004) discusses some problems with base correlations).

4 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 the single-name CDS quotes of the underlying reference entities. Furthermore, even if simulation of cash flows is needed, e.g. in traditional cash CDOs with path-dependent waterfalls or in CDO-squared or -cubed, the analytical methods are still useful for variance reduction using control variate techniques. Analytical solutions are 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 fairly 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 problematic interpolation schemes based on implied correlations in the Gaussian copula, which are widely used in practice.

An interesting topic for future research is the hedging performance of the alternative default correlation models proposed in the literature. 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. For this type of process, Duffie, Pan, and Singleton (2000) derive closed-form solutions – in terms of deterministic ordinary differential equations (ODEs) – to a wide range of expectations relevant in e.g. derivatives pricing. In some cases the ODEs have known explicit solutions.

This paper applies the basic affine sub-class, as defined by Duffie and Gârleanu (2001). A *basic affine process*, which I denote $x \sim AJD(x_0, \kappa, \theta, \sigma, l, \mu)$, is a stochastic 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.

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 complexity of estimating two more parameters in an already rich model. Also, the jump size distribution need not be exponential in the general class, but this ensures positive default intensities and seems reasonable since most large jumps in credit spreads are positive. Moreover, this jump specification allows for the explicit solution given below.

A.1 Default probabilities

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

$$E\left[e^{-\int_0^T x_s ds}\right] = e^{A(T;\,\kappa,\theta,\sigma,l,\mu)+B(T;\,\kappa,\sigma)x_0} \tag{29}$$

where

$$\begin{aligned} A(T;\,\kappa,\theta,\sigma,l,\mu) &= \frac{\kappa\theta\gamma}{bc_1d_1}\log\left(\frac{c_1+d_1e^{bT}}{-\gamma}\right) + \frac{\kappa\theta}{c_1}T \\ &\quad + \frac{l(ac_2-d_2)}{bc_2d_2}\log\left(\frac{c_2+d_2e^{bT}}{c_2+d_2}\right) + \frac{l-c_2l}{c_2}T \\ B(T;\,\kappa,\sigma) &= \frac{1-e^{bT}}{c_1+d_1e^{bT}} \end{aligned} \tag{30}$$

with

$$\begin{aligned}
\gamma &= \sqrt{\kappa^2 + 2\sigma^2} \\
c_1 &= -(\gamma + \kappa)/2 \\
d_1 &= c_1 + \kappa \\
c_2 &= 1 - \mu/c_1 \\
d_2 &= (d_1 + \mu)/c_1 \\
b &= d_1 + (\kappa c_1 - \sigma^2)/\gamma \\
a &= d_1/c_1
\end{aligned}$$

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, $x \sim AJD(x_0, \kappa, \theta, \sigma, l, \mu)$,*

$$E\left[e^{iu\int_0^T x_t dt}\right] = e^{\tilde{A}(0;T,u,\kappa,\theta,\sigma,l,\mu)+\tilde{B}(0;T,u,\kappa,\sigma)x_0} \tag{31}$$

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

$$\begin{aligned}
\tilde{A}'_{Re}(t) &= -\kappa\theta\tilde{B}_{Re}(t) - \frac{l\mu\left[\tilde{B}_{Re}(t) - \mu\tilde{B}_{Re}(t)^2 - \mu\tilde{B}_{Im}(t)^2\right]}{[1-\mu\tilde{B}_{Re}(t)]^2 + \mu^2\tilde{B}_{Im}(t)^2} \\
\tilde{A}'_{Im}(t) &= -\kappa\theta\tilde{B}_{Im}(t) - \frac{l\mu\tilde{B}_{Im}(t)}{[1-\mu\tilde{B}_{Re}(t)]^2 + \mu^2\tilde{B}_{Im}(t)^2} \\
\tilde{B}'_{Re}(t) &= \kappa\tilde{B}_{Re}(t) - \frac{\sigma^2}{2}\left[\tilde{B}_{Re}(t)^2 - \tilde{B}_{Im}(t)^2\right] \\
\tilde{B}'_{Im}(t) &= \kappa\tilde{B}_{Im}(t) - \sigma^2\tilde{B}_{Re}(t)\tilde{B}_{Im}(t) - u
\end{aligned} \tag{32}$$

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). Recall that $Z_t = \int_0^t x_s$ and define

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

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

$$\Psi_T - \Psi_0 = \int_0^T \mu_\Psi(t, x_t)\Psi_t dt + \int_0^T \sigma_\Psi(t, x_t)\Psi_t dW_t + J_\Psi(T) \tag{33}$$

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 (33) 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,

$$\begin{aligned}
h(c) &= \frac{1}{\mu}\int_0^\infty e^{vc_{Re}+ivc_{Im}}e^{-v/\mu}dv \\
&= \frac{1}{\mu}\int_0^\infty e^{v(c_{Re}-\frac{1}{\mu})}[cos(vc_{Im}) + isin(vc_{Im})]dv
\end{aligned}$$

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 $c_{Re} < 1/\mu$, we get

$$\begin{aligned}
h(c) &= \frac{\mu}{(\mu c_{Re}-1)^2+\mu^2 c_{Im}^2} \times \\
&\quad \Big[e^{v(c_{Re}-\frac{1}{\mu})}\Big\{(c_{Re}-\frac{1}{\mu})cos(vc_{Im}) + c_{Im}sin(vc_{Im}) \\
&\qquad + i(c_{Re}-\frac{1}{\mu})sin(vc_{Im}) - ic_{Im}cos(vc_{Im})\Big\}\Big]_{v=0}^{\infty} \\
&= \frac{(1-\mu c_{Re})+i\mu c_{Im}}{(1-\mu c_{Re})^2+\mu^2 c_{Im}^2}
\end{aligned}$$

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 (33), 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 (32). □

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.[16] To compensate for that, the protection

[16] Often, the CDS contract offers the protection buyer a cheapest-to-deliver option – i.e. the option to choose the delivered bond from a list of eligible bonds. This option has a non-negative impact on the CDS premium (the value of the protection leg is potentially increased), but the effect is small and, as in most studies, it is ignored in this paper. In recent empirical studies, Guha (2003) and Acharya, Bharath, and Srinivasan (2004) find that the recovery values of different bonds of a defaulted issuer usually are very similar across maturities and coupons. This finding supports the assumption of recovery of face value at the time of default, and with this assumption the delivery option is worthless.

buyer pays fixed premium payments periodically until default or maturity of the contract is reached.

Formally, with notation as in Sections 2 and 3, 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 premium, S, is paid in arrears at a frequency f (i.e. f payments of S/f each year), typically quarterly. 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

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

where $t_j = j/f$ for $j = 1, \ldots, fT$.

With discretization and independence assumptions between recovery rates, interest rates and default events as in Sections 2 and 3, the value of the protection leg is

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

Similarly, the value of the premium leg is

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

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 parameters can be found as the solution to the same equation.

Appendix C: Three copula models

For completeness, this appendix gives a very short outline of the three 1-factor copula models used for comparison with the intensity model.

In all three copula models, default up to time t of entity i is defined as the event that a default variable is below some default boundary, $X_i \leqslant c_i$, where the default variable is a linear combination of a market factor and an idiosyncratic factor.

The default boundary, c_i, is derived from the marginal default probability implied from single-name CDS spreads. This gives a loss distribution for the horizon t. Loss distributions at different time horizons are build using the same specification for the default variables, X_i, but using different default boundaries (increasing with the horizon). The marginal default probabilities at different horizons are obtained by backing out deterministic default intensities from CDS spreads. In the copula application of this paper, a constant intensity is backed out from the 5-year CDS spreads.

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

$$X_i = a_i Z + \sqrt{1 - a_i^2}\, \epsilon_i$$

where $Z, \epsilon_1, \ldots, \epsilon_N$ are independent standard normal random variables and the factor loadings are constant. The correlation between any pair of default variables, X_i and X_j, is then $\rho_{ij} = a_i a_j$. For all the models in the comparison, homogeneous correlation parameters are assumed across all entities, $\rho = a_i a_j$. The conditional and unconditional default probabilities in the Gaussian copula are given in terms of the standard normal distribution function.

(ii) In the stochastic factor loading copula, the default variable is defined as

$$X_i = a_i(Z) Z + v_i \epsilon_i + m_i$$

again with independent standard normal common and idiosyncratic factors. The factor loadings are stochastic, and v_i and m_i are constants chosen to ensure zero mean and unit variance. This paper applies the tractable two-point regime switching version with homogeneous correlations,

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

The intuition is that correlations increase ($\rho_2 > \rho_1$) in bad states of the economy (represented by $Z < \nu$). The conditional and unconditional default probabilities in this version of the model are still known in closed form. For more details, see Andersen and Sidenius (2004).

(iii) In the double-t copula, the default variable is

$$X_i = a_i \frac{Z}{std.dev.(Z)} + \sqrt{1 - a_i^2} \frac{\epsilon_i}{std.dev.(\epsilon_i)}$$

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 $\rho = a_i a_j$. For more details, see Hull and White (2004). 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 by the distribution of a sum of two t distributions, which is unknown (not a t distribution).

References

Acharya, V. V., S. T. Bharath, and A. Srinivasan (2004): "Understanding the Recovery Rates on Defaulted Securities," Working Paper, London Business School, Univ. of Michigan and Univ. of Georgia.

Andersen, L. and J. Sidenius (2004): "Extensions to the Gaussian Copula: Random Recovery and Random Factor Loadings," Working Paper, Bank of America; forthcoming in *Journal of Credit Risk*.

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

Black, F. and J. Cox (1976): "Valuing Corporate Securities: Some Effects of Bond Indenture Provisions," *Journal of Finance*, 31(2):351–367.

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

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

Duffee, G. (1999): "Estimating the Price of Default Risk," *Review of Financial Studies*, 12(1):197–226.

Duffie, D. (1999): "Credit Swap Valuation," *Financial Analysts Journal*, 55(1):73–87.

——— (2004): "Time to adapt copula methods for modelling credit risk correlation," *Risk*, April:41–59.

Duffie, D. and N. Gârleanu (2001): "Risk and Valuation of Collateralized Debt Obligations," *Financial Analysts Journal*, 57(1):41–59.

Duffie, D., J. Pan, and K. Singleton (2000): "Transform Analysis and Asset Pricing for Affine Jump Diffusions," *Econometrica*, 68(6):1343–1376.

Duffie, D. and K. Singleton (1999): "Modeling Term Structures of Defaultable Bonds," *Review of Financial Studies*, 12(4):687–720.

——— (2003): *Credit Risk: Pricing, Measurement and Management,* Princeton University Press.

Finger, C. C. (2005): "Issues in the pricing of synthetic CDOs," *Journal of Credit Risk*, 1(1):113–124.

GREGORY, J. AND J.-P. LAURENT (2003): "I Will Survive," *Risk*, June:103–107.

GUHA, R. (2003): "Recovery of Face Value at Default: Empirical Evidence and Implications for Credit Risk Pricing," Working Paper, London Business School.

HAMILTON, D., P. VARMA, S. OU, AND R. CANTOR (2004): "Default & Recovery Rates of Corporate Bond Issuers," Moody's Investors Service, New York, January 2004.

HULL, J. AND A. WHITE (2000): "Valuing Credit Default Swaps I: No counterparty default risk," *Journal of Derivatives*, 8(1):29–40.

——— (2004): "Valuation of a CDO and an n'th to Default CDS Without Monte Carlo Simulation," *Journal of Derivatives*, 12(2):8–23.

JARROW, R. AND S. TURNBULL (1995): "Pricing Options on Financial Securities Subject to Credit Risk," *Journal of Finance*, 50(1):53–85.

LANDO, D. (1994): "Three Essays on Contingent Claims Pricing," PhD Dissertation. Cornell University.

——— (1998): "On Cox processes and Credit Risky Securities," *Review of Derivatives Research*, 2(2-3):99–120.

LI, D. X. (2000): "On Default Correlation: A Copula Function Approach," *Journal of Fixed Income*, 9(4):43–54.

LONGSTAFF, F. A., S. MITHAL, AND E. NEIS (2004): "Corporate Yield Spreads: Default Risk or Liquidity? New Evidence from the Credit-Default Swap Market," Working Paper, UCLA; forthcoming in *Journal of Finance.*

MERTON, R. C. (1974): "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance*, 29(2):449–470.

PRESS, W. H., S. A. TEUKOLSKY, W. T. VETTERLING, AND B. FLANNERY (1992): *Numerical Recipes in C: The Art of Scientific Computing,* Cambridge University Press, 2nd edition.

SCHÖNBUCHER, P. (1998): "Term Structure Modelling of Defaultable Bonds," *Review of Derivatives Research*, 2(2-3):161–192.

——— (2003): *Credit Derivatives Pricing Models: Models, Pricing and Implementation,* John Wiley and Sons Ltd.

SCHÖNBUCHER, P. AND D. SCHUBERT (2001): "Copula-Dependent Default Risk in Intensity Models," Working Paper, Bonn University.

Černy, A. (2004): "Introduction to Fast Fourier Transform in Finance," *Journal of Derivatives*, 12(1):73–88.

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

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:							
Homogeneous CDS	26.8%	144.2	62.7	41.7	19.2	0.67	$\kappa = 0.37$, $\sigma = 0.059$, $\bar{l} = 0.016$, $\mu = 0.091$, $w = 0.91$
Heterogeneous CDS	27.0%	142.5	62.8	41.3	18.1	0.74	$\kappa = 0.27$, $\sigma = 0.050$, $\bar{l} = 0.017$, $\mu = 0.078$, $w = 0.93$
Jump-diffusion intensities: $w \leqslant 0.7$							
Homogeneous CDS	27.4%	134.0	65.9	42.6	17.7	1.12	$\kappa = 0.40$, $\sigma = 0.056$, $\bar{l} = 0.026$, $\mu = 0.081$, $w = 0.70$
Heterogeneous CDS	27.4%	133.9	65.9	42.2	16.9	1.15	$\kappa = 0.39$, $\sigma = 0.055$, $\bar{l} = 0.026$, $\mu = 0.081$, $w = 0.70$
Pure diffusion intensities:							
Homogeneous CDS	35.6%	150.0	12.6	0.9	0.0	6.51	$\kappa = 0.48$, $\sigma = 0.079$, $w = 1.00$
Heterogeneous CDS	35.5%	147.0	11.8	0.7	0.0	6.54	$\kappa = 0.48$, $\sigma = 0.079$, $w = 1.00$
Gaussian copula:							
Homogeneous CDS	28.8%	226.5	55.3	15.0	1.8	4.74	$\rho = 0.150$
Heterogeneous CDS	28.5%	225.6	55.3	15.1	1.8	4.69	$\rho = 0.159$
Gauss. cop. w/stochastic factor loadings:							
Homogeneous CDS	25.8%	150.3	53.3	44.6	18.7	0.91	$\nu = -1.86$, $\rho_1 = 0.102$, $\rho_2 = 0.309$
Heterogeneous CDS	25.6%	150.1	53.3	44.5	18.3	0.91	$\nu = -1.86$, $\rho_1 = 0.109$, $\rho_2 = 0.324$
Double-t copula:							
Homogeneous CDS	25.0%	150.7	57.8	34.2	18.0	0.41	$d = 4$, $\rho = 0.268$
Heterogeneous CDS	25.0%	154.1	58.7	34.0	17.2	0.53	$d = 4$, $\rho = 0.272$

Table 1: Market and model prices for the DJ iTraxx 5-year index tranches. The market prices were obtained from Bloomberg on August 23, 2004. 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 (28). Interest rates are constant at 3%, and the recovery rate is 40%.

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:							
Homogeneous CDS	51.3%	349.7	124.6	66.1	16.5	3.20	$\kappa = 0.25$, $\sigma = 0.059$, $\bar{l} = 0.048$, $\mu = 0.059$, $w = 0.79$
Heterogeneous CDS	49.6%	343.8	122.9	65.7	15.3	2.80	$\kappa = 0.20$, $\sigma = 0.054$, $\bar{l} = 0.037$, $\mu = 0.067$, $w = 0.93$
Jump-diffusion intensities: $w \leqslant 0.7$							
Homogeneous CDS	51.9%	341.8	126.6	66.0	15.3	3.22	$\kappa = 0.28$, $\sigma = 0.059$, $\bar{l} = 0.060$, $\mu = 0.057$, $w = 0.70$
Heterogeneous CDS	51.1%	329.1	127.2	63.8	12.1	2.89	$\kappa = 0.31$, $\sigma = 0.060$, $\bar{l} = 0.060$, $\mu = 0.065$, $w = 0.70$
Pure diffusion intensities:							
Homogeneous CDS	58.6%	444.5	65.4	7.4	0.1	7.38	$\kappa = 0.30$, $\sigma = 0.082$, $w = 1.00$
Heterogeneous CDS	58.3%	416.5	51.1	4.4	0.0	7.53	$\kappa = 0.30$, $\sigma = 0.082$, $w = 1.00$
Gaussian copula:							
Homogeneous CDS	49.7%	485.6	134.1	36.9	2.7	5.84	$\rho = 0.150$
Heterogeneous CDS	47.8%	464.8	132.9	38.9	3.3	5.10	$\rho = 0.186$
Gauss. cop. w/stochastic factor loadings:							
Homogeneous CDS	51.2%	342.0	130.0	70.5	9.8	3.25	$\nu = -1.38$, $\rho_1 = 0.064$, $\rho_2 = 0.229$
Heterogeneous CDS	49.2%	336.9	128.8	67.2	9.4	2.75	$\nu = -1.37$, $\rho_1 = 0.091$, $\rho_2 = 0.265$
Double-t copula:							
Homogeneous CDS	47.8%	351.9	115.0	58.2	22.8	2.83	$d = 4$, $\rho = 0.242$
Heterogeneous CDS	46.4%	350.7	113.3	55.4	20.6	2.41	$d = 4$, $\rho = 0.278$

Table 2: Market and model prices for the DJ CDX 5-year index tranches. The market prices were obtained from Bloomberg on August 23, 2004. 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 (28). Interest rates are constant at 3%, and the recovery rate is 40%.

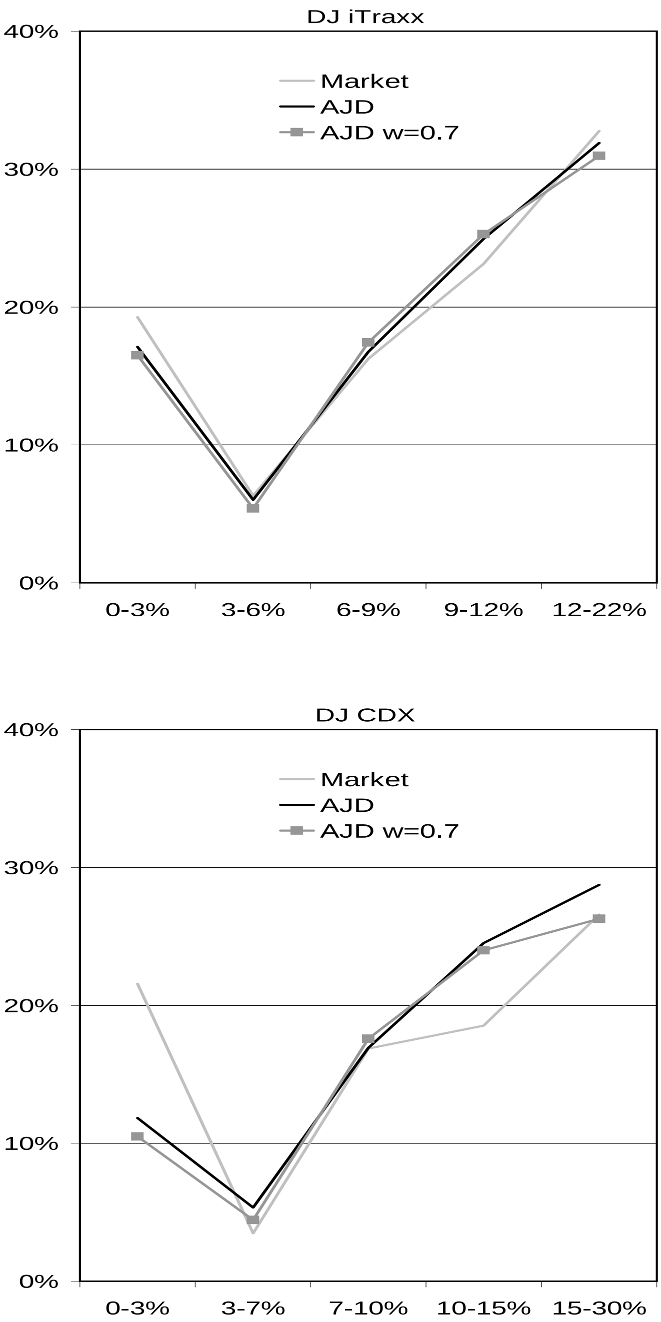

Figure 2: Implied correlations from the market and the AJD model on the 5-year DJ iTraxx and CDX tranches. The market prices were obtained from Bloomberg on August 23, 2004. The model prices are from the fitted AJD models with parameters as listed in Tables 1 and 2 (calibrated to heterogeneous CDS spreads). The implied correlations are calculated in a homogeneous Gaussian copula with an interest rate of 3% and a recovery rate of 40%.

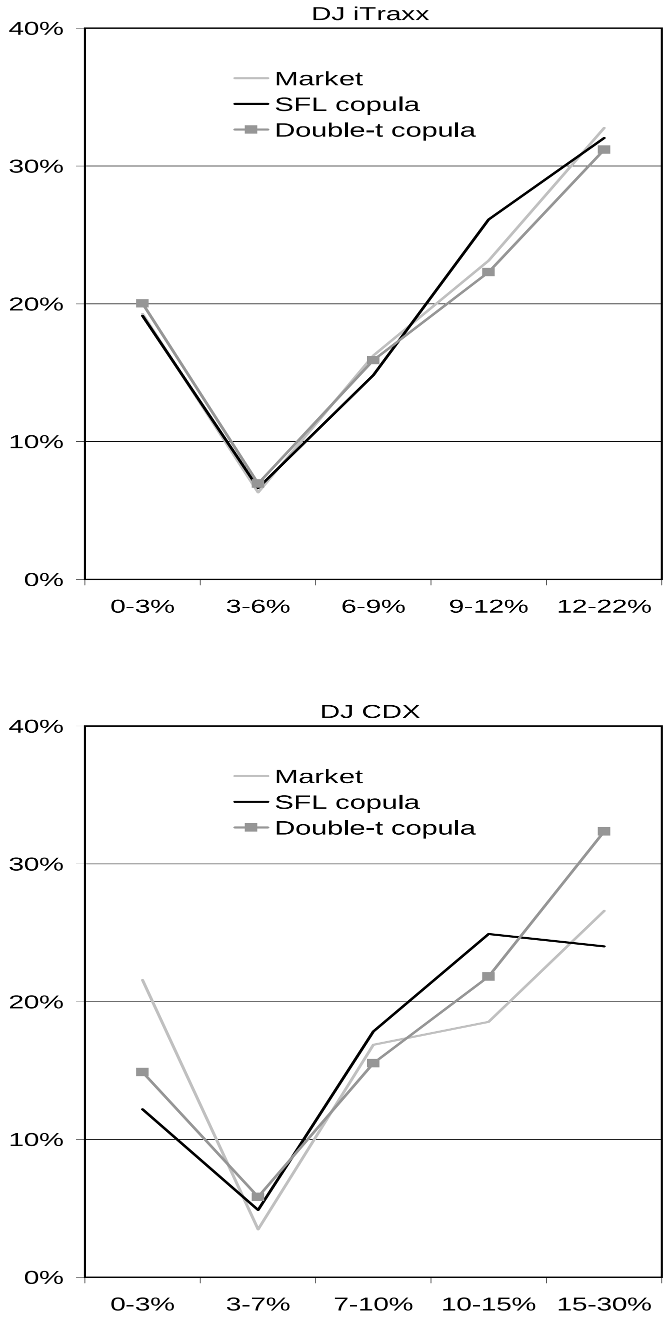

Figure 3: Implied correlations from the market and the copula models on the 5-year DJ iTraxx and CDX tranches. The market prices were obtained from Bloomberg on August 23, 2004. The model prices are from the fitted double-t copula and the fitted Gaussian copula with stochastic factor loadings with parameters as listed in Tables 1 and 2 (calibrated to heterogeneous CDS spreads). The implied correlations are calculated in a homogeneous Gaussian copula with an interest rate of 3% and a recovery rate of 40%.

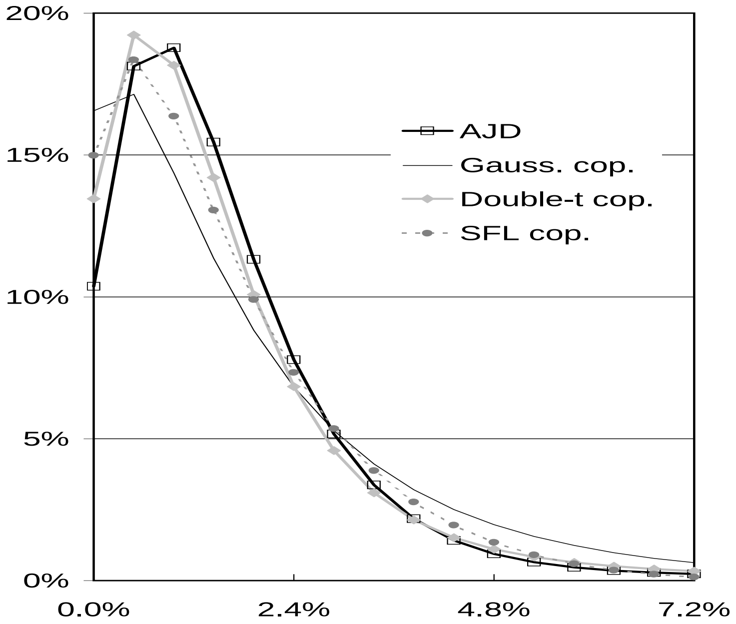

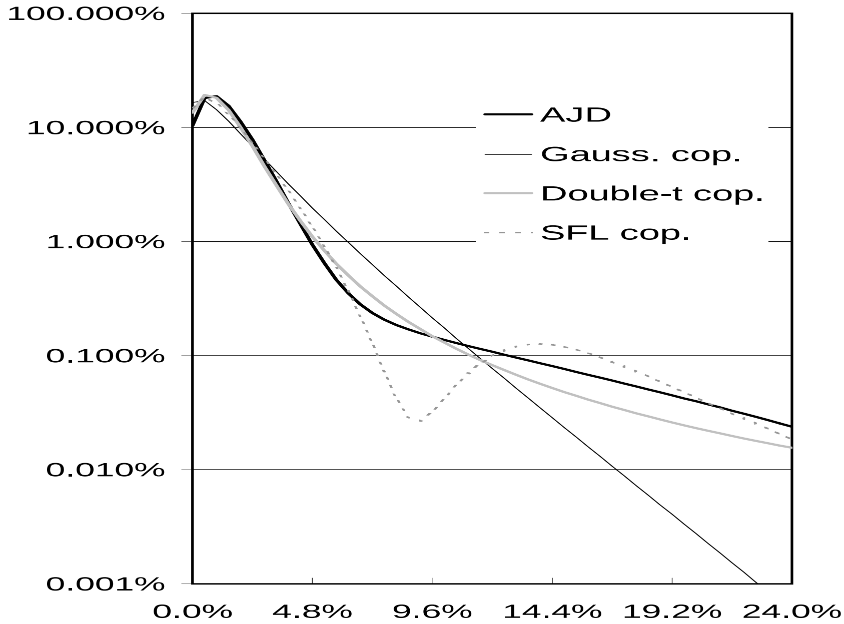

Figure 4: 5-year loss distributions for the DJ iTraxx pool. The graphs display probability functions for the loss percentage with support $0\%, 0.48\%, \ldots, 60\%$, corresponding to $0, 1, \ldots, 125$ defaults and recovery 40%. The expected loss is around 2.4%. The upper panel focuses on the low-loss probabilities, 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 Table 1 (the heterogeneous case).

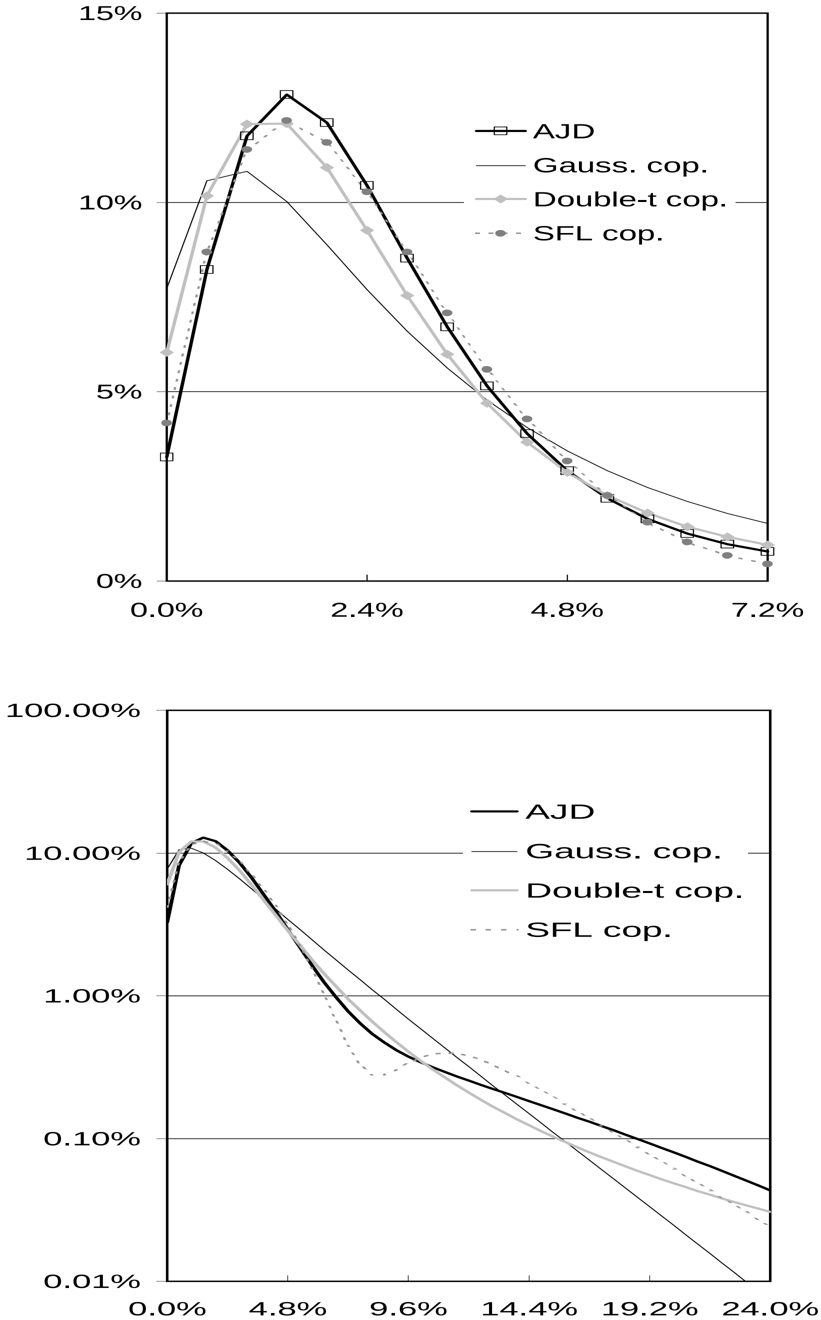

Figure 5: 5-year loss distributions for the DJ CDX pool. The graphs display probability functions for the loss percentage with support $0\%, 0.48\%, \ldots, 60\%$, corresponding to $0, 1, \ldots, 125$ defaults and recovery 40%. The expected loss is around 3.9%. The upper panel focuses on the low-loss probabilities, 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 Table 2 (the heterogeneous case).