

An Empirical Investigation of an Intensity-Based Model for Pricing CDO Tranches *

Peter Feldhütter †

This draft: May 7, 2008
First draft: March 22, 2007

Abstract

Using an extensive data set of 15,600 CDS and CDO tranche spreads on the North American Investment Grade CDX index I conduct an empirical analysis of a Duffie and Gârleanu (2001) intensity-based model for correlated defaults. I examine the model with respect to model assumptions, pricing in both the cross section and time series dimension, and hedging ability. The results show that the model assumptions are reasonable and that average prices are matched well. In addition, the model accurately tracks the prices over time of the more risky tranches. Finally, the model sensitivity of the most risky tranches to underlying CDS spreads match actual sensitivities better than those implied by the commonly used Gaussian copula. The last result suggests that the model is well-suited for hedging the equity tranche.

*I would like to thank Andreas Eckner, David Lando, Jesper Lund, Allan Mortensen and seminar participants at the International Financial Research Forum 2008 in Paris for helpful discussions and comments.

†Department of Finance, Copenhagen Business School, Solbjerg Plads 3, DK-2000 Frederiksberg, Denmark (e-mail: pf.fi@cbs.dk).

1 Introduction

The market for credit derivatives has experienced a massive growth in recent years and investors can now take on exposure to specific segments of the default loss distribution of an index of firms by trading in standardized Collateral Debt Obligation (CDO) tranches. Each tranche has a different sensitivity to the default correlation and credit risk of firms in the index, and numerous models specifying default and correlation dynamics have been proposed in the last years.

A good model of multi-name default should ideally have the following properties. First, the model should be able to match actual tranche prices consistently such that for a fixed set of model parameters, prices are matched over a period of time. This is important for pricing non-standard tranches in a market where prices are available for standard tranches. Second, CDO tranche prices should exhibit correct sensitivities to changes in the underlying firms' default probabilities, such that this 'spread risk' can be hedged. Often, market participants take a view on correlation by investing in a CDO tranche and hedge the 'spread risk' by entering an offsetting position in underlying CDS contracts. The accuracy of the offsetting position depends on the model's ability to replicate the CDO tranche price sensitivities with respect to the underlying CDS contracts. Third, the model should have parameters that are economically interpretable, such that parameter values can be discussed and critically evaluated. If a bespoke CDO tranche on a non-standard pool of underlying firms needs to be priced and parameters cannot be inferred from existing market prices, economic interpretability provides guidance in choosing parameters. Fourth, credit spreads and their correlation should be modelled dynamically such that options on multi-name products can be priced. And fifth, since market makers quote prices at any given time, pricing formulas should not be too time-consuming to evaluate.

The standard model for pricing and hedging CDO tranches is the Gaussian copula. The model is a static one-period model (applied repeatedly to cash flows at different dates) where the default environment for all firms is determined by a single normal variable. Since there is no dynamic behavior for default risk present in the model, it lacks realism and does not have the ability to value CDO derivatives such as options. Also, holding parameters fixed it is not possible to consistently price CDO tranches over time or even price tranches consistently on a given day. Matching different tranche prices on a single day requires different correlation parameters, a phenomenon called the

'correlation skew'. This is problematic since it is not clear which correlation parameter to use when pricing non-standard tranches and simple interpolation schemes can lead to a number of problems such as negative prices (see Willemann (2005)). Alternatives to the Gaussian copula have been proposed such as the t -copula (Mashal and Naldi (2002)), double- t copula (Hull and White (2004)), and the Clayton copula to name just a few and while some of the models match CDO tranche spreads better they are in essence still static models.

In single-name default modeling the stochastic intensity-based framework introduced in Lando (1994) and Duffie and Singleton (1999) has proven very successful and is widely used¹. Default of a firm in an intensity-based model is determined by the first jump of a pure jump process with a stochastic default intensity. Duffie and Gârleanu (2001) extend the single-name intensity framework to a multi-name setting by letting firms' default intensities be the sum of an idiosyncratic component and a common component that affects the default of all firms. In this setting, pricing of options is possible, parameters are interpretable, and hedge ratios can be calculated. While the model has many attractive properties it has not been much used due to two reasons. First, it has been argued that it is too time-consuming to price correlation products in intensity models (Hull and White (2004)). However, Mortensen (2006) prices CDO tranches semi-analytically and finds that the model is as fast as for example the t -copula of Hull and White (2004). Second, some argue that intensity-based models cannot generate enough default correlation to match market prices. This is true in a pure diffusion, but with jumps in the intensity process market prices can indeed be matched.

This paper estimates a multi-name intensity model as proposed in Duffie and Gârleanu (2001) and modified in Mortensen (2006) to allow for heterogeneous default intensity dynamics. While Mortensen (2006) calibrates the model to CDS and CDO tranche spread on a single day, I use an extensive panel set of 15,600 CDS and CDO spreads on the Dow Jones North American Investment Grade Index and examine the properties of the model along several dimensions: assumptions regarding the underlying firms' default intensity dynamics are thoroughly examined by estimating the default intensity dynamics of all the 125 underlying firms separately. Also, the ability to price CDO tranches is carefully examined in both the cross section and

¹Examples of empirical applications are Duffie and Singleton (1997), Duffee (1999), and Longstaff, Mithal, and Neis (2005).

time series dimension. In other words, I examine whether the model is able to simultaneously match both average CDO spreads and the dynamics of CDO spreads over time. Eckner(2007a, 2007b) implements a similar model but recalibrates the model parameters daily and is focused on risk premia on the CDO market. Finally, I compare the hedging ability of the model with that of a standard Gaussian copula by seeing how well the models match the sensitivity of CDO tranche prices to changes in underlying CDS spreads. To the best of my knowledge such a comprehensive evaluation of a CDO pricing model has not been done in the literature yet. Longstaff and Rajan (2006) and Bhansali, Gingrich, and Longstaff (2008) use time series of CDO tranche prices to estimate a portfolio loss model, but since they model only portfolio loss they cannot evaluate whether the loss dynamics of the model is consistent with underlying firms' default intensity dynamics and the model does not produce deltas with respect to underlying CDS spreads. Houdain and Guegan (2006) look at the ability of several copula models to hedge one CDO tranche with another, but they do not discuss pricing.

The empirical results show that the model is able to match average CDO tranche spreads across time. In addition, the model captures the spread variation through time in the most risky tranches well, while the model exhibits too little price variation in senior tranches. Finally, the model matches the sensitivity of the equity tranche, the most risky tranche, to underlying CDS spreads better than a Gaussian copula. This implies that the model hedges spread risk in the equity tranche more accurately than the Gaussian copula.

The remainder of the paper is organized as follows. Section 2 formulates the multi-name default model and derives CDO tranche pricing formulas. Section 3 explains the estimation methodology and section 4 describes the data. Section 5 examines the modeling assumptions and pricing results while section 6 report hedging results. Section 7 concludes.

2 Intensity-Based Default Risk Model

This section explains the model framework for single-name and multi-name defaults. For pricing single-name credit default swaps the intensity-based framework introduced in Lando (1994) and Duffie and Singleton (1999) is used and default intensities are assumed to be affine. For multi-name collateralized debt obligation valuation I follow Duffie and Gârleanu (2001) in modeling the default intensity of each issuer as the sum of an idiosyncratic

and a common affine process. Default correlation is created through the joint dependence of default intensities on the common factor and loss distributions are calculated semi-analytically as in Mortensen (2006).

2.1 Default Modeling

Default of a single issuer, τ , is supposed to arrive at intensity $\lambda_t \geq 0$ which implies that the conditional probability at time t of defaulting within a small period of time Δt is approximately

$$P(\tau \leq t + \Delta t | \tau > t) \approx \lambda_t \Delta t.$$

Unconditional default probabilities are given by

$$P(\tau \leq t) = 1 - E[e^{-\int_0^t \lambda_s ds}]$$

which shows that default probabilities in an intensity-based framework can be calculated using techniques from interest rate modeling.

Assume there are P different issuers. To model correlation between issuers I follow Mortensen (2006) and assume that the intensity of each issuer can be written as the sum of an idiosyncratic component and a scaled common component

$$\lambda_{i,t} = a_i Y_t + X_{i,t} \tag{1}$$

where $a_i > 0, i = 1, \dots, P$ are constants and Y, X_1, X_2, \dots, X_P are independent. The common factor Y creates dependence in defaults of different issuers and may be viewed as governing economic performance of the economy while X_i governs the idiosyncratic default risk of firm i . a_i is the sensitivity of firm i to the performance of the economy and is relevant because in the case of $a_i = 1, i = 1, \dots, P$, which is assumed in Duffie and Gârleanu (2001), the common factor would have to be smaller than the smallest intensity thereby restricting the amount of default correlation the model can generate.

Both the common factor and idiosyncratic factors are assumed to follow affine jump diffusions under the risk-neutral measure

$$d\xi_t = (\kappa_0 + \kappa_1 \xi_t) dt + \sigma \sqrt{\xi_t} dW_t^Q + dJ_t^Q \tag{2}$$

where W^Q is a Brownian motion, jump times (independent of W^Q) are those of a Poisson process with intensity $l \geq 0$, and jump sizes are independent

of the jump times and follow an exponential distribution with mean $\mu > 0$. The process is well-defined if $\kappa_0 > 0$. As a special case, if the jump intensity is equal to zero the intensity follows a CIR process. In short this process is denoted

$$AJD(\xi_0, \kappa_0, \kappa_1, \sigma, l, \mu).$$

As a useful result note that $a_i \xi$ is again $AJD(a_i \xi_0, a_i \kappa_0, \kappa_1, \sqrt{a_i} \sigma, l, a_i \mu)$.

2.2 Risk Premium

For the basic affine process in equation (2) I assume an essentially affine risk premium for the diffusive risk and constant risk premia for the risk associated with the timing and sizes of jumps. Cheridito, Filipovic, and Kimmel (2007) propose an extended affine risk premium as an alternative to an essentially affine risk premium, which would allow the parameter κ_0 to be adjusted under P in addition to the adjustment of κ_1 . However, extended affine models require the Feller condition to hold as discussed in Feldhütter (2006), and this restriction is violated in 111 out of the 125 single-name estimations as well as in the CDO estimation in the empirical section². As I do not wish to impair the pricing ability of the model because of the risk premium assumption, I choose an essentially affine risk premium.

This leads to the following dynamics for the factor under the historical measure P

$$d\xi_t = (\kappa_0 + \kappa_1^P \xi_t) dt + \sigma \sqrt{\xi_t} dW_t^P + dJ_t^P \quad (3)$$

where the jump times are those of a Poisson process with intensity l^P and the jump sizes are independent of the jump times and follow an exponential distribution with mean $\mu^P > 0$.

2.3 Loss Distribution

The loss distribution at time t for P issuers is found semi-analytically by calculating it conditional on the common factor and then integrating out the

²To illustrate why the Feller condition is necessary in extended affine models consider the simple diffusion case, $dX_t = (\kappa_0^Q + \kappa_1^Q X_t) dt + \sigma \sqrt{X_t} dW^Q$. The risk premium $\Lambda_t = \frac{\lambda_0}{\sqrt{X_t}} + \lambda_1 \sqrt{X_t}$ keeps the process affine under P but the risk premium explodes if $X_t = 0$. To avoid this, the Feller restriction $\kappa_0 > \frac{\sigma^2}{2}$ under both P and Q ensures that X_t is strictly positive.

common factor. It is therefore convenient to define the integrated common process

$$I_t = \int_0^t Y_s ds.$$

Conditional on I_t defaults are independent and given as

$$p_i(t|z) = P(\tau_i \leq t | I_t = z) = 1 - e^{-a_i z + A_i(t) + B_i(t) X_{i,0}}$$

where the functions A_i and B_i have closed-form solutions derived in Duffie and Gârleanu (2001) and restated in Appendix B. The conditional probability of observing j defaults among K issuers is found by the recursive algorithm

$$P(D_t^K = j|z) = P(D_t^{K-1} = j|z)(1 - p_K(t|z)) + P(D_t^{K-1} = j-1|z)p_K(t|z)$$

which is due to Andersen, Sidenius, and Basu (2003) (the last term disappears if $j = 0$). The unconditional loss distribution is

$$P(D_t^K = j) = \int_0^\infty P(D_t^K = j|z) f(z) dz \quad (4)$$

where f is the density function of I_t . The density function is found by Fourier inversion of the characteristic function $\phi_{I_t}(u)$ of I_t

$$f(z) = \frac{1}{2\pi} \int_{-\infty}^\infty e^{-iuz} \phi_{I_t}(u) du. \quad (5)$$

Finally, the characteristic function of I_t is exponentially-affine in z and the results in Duffie and Gârleanu (2001) give an explicit expression.³

2.4 Synthetic CDO Pricing

CDOs began to trade frequently in the mid-nineties and in the last decade issuance of CDOs has experienced a massive growth. In a CDO the credit risk of a portfolio of debt securities is passed on to investors by issuing CDO

³Duffie and Gârleanu (2001) derive an explicit solution for $E[\exp(q \int_0^t X_s ds)]$ where X is an AJD and q is a real number. As noted by Eckner (2007a) the formula works for q complex as well.

tranches written on the portfolio. The tranches have varying risk profiles according to their seniority. A synthetic CDO is written on CDS contracts instead of actual debt securities. To illustrate the cash flows in a synthetic CDO an example that reflects the data used in this paper is useful.

Consider a CDO issuer, called A , who sells protection with notional \$0.8 million in 125 5-year CDS contracts for a total notional of \$100 million. Each CDS contract is written on a specific corporate bond, and A receives quarterly a CDS premium until the contract expires or the bond defaults. In case of default, A receives the defaulted bond in exchange for face value. The loss is the difference between face value and market value⁴.

A issues a CDO tranche on the first 3% losses in his CDS portfolio and B "buys" this tranche. No money is exchanged at time 0 but the principal on the tranche is \$3 million. If the premium on the tranche is, say, 2,000 basis points, A pays a quarterly premium of 500 basis points to B on the remaining principal. If a default occurs on any of the underlying CDS contracts, the loss is covered by B and his/her principal is reduced accordingly. B continues to receive the premium on the remaining principal until either the CDO tranche matures or the remaining principal is exhausted. Since the first 3% portfolio losses are covered by this tranche it is called the 0-3% tranche. Similarly, A sells 3 – 7%, 7 – 10%, 10 – 15%, 15 – 30%, and 30 – 100% tranches such that the total principal equals the principal in the CDS contracts. For a tranche covering losses between K_1 and K_2 , K_1 is called the attachment point and K_2 the exhaustion point.

Next, I find the fair spread on a specific tranche. Consider a CDO tranche maturing at time T , with quarterly payments $t_1, t_2, \dots, t_M = T$, and covering portfolio losses from K_1 to K_2 . The coupon rate for the tranche is found by equating the value of the protection and premium leg and solving for the coupon rate. Denoting the total portfolio loss in percent as L_t , the tranche loss is given as

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

and the value of the protection leg in a CDO tranche with maturity T is

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

⁴Pricing CDS contracts is explained in Appendix A.

while the value of the premium leg with premium S is

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

where r_s refers to the riskless rate and $\int_{t_{j-1}}^{t_j} \frac{K_2 - K_1 - T_{K_1, K_2}(L_s)}{t_j - t_{j-1}} ds$ is the remaining principal during the period t_{j-1} to t_j . All expectations are under the risk-neutral measure. Following Mortensen (2006) I discretize the integral at premium payment dates and assume that riskfree rates are uncorrelated with losses and defaults occur halfway between premium payments. With these assumptions the value of the protection leg is

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

where $t_0 = 0$ while the value of the premium leg is

$$Prem(0, T; S) = S \sum_{j=1}^M (t_j - t_{j-1}) P(0, t_j) \left(K_2 - K_1 - \frac{E[T_{K_1, K_2}(L_{t_{j-1}})] + E[T_{K_1, K_2}(L_{t_j})]}{2}\right).$$

The tranche premium S solves $Prot(0, T) = Prem(0, T; S)$.

2.5 A Parsimonious Model of Single- and Multi-Name Default

The model for default probabilities and correlations outlined in section 2.1 has in its most general form 126 latent processes with associated parameters, so it is necessary to simplify the model considerably in order to estimate the parameters and latent variables. In the following a more parsimonious version of the general model is outlined and in the section 5 these assumptions are thoroughly examined.

I impose the restrictions

$$\kappa_{1,i} = \kappa_{1,Y} \tag{6}$$

$$\sigma_i = \sqrt{a_i} \sigma_Y \tag{7}$$

$$\mu_i = a_i \mu_Y \tag{8}$$

such that the default intensity of issuer i , $\lambda_{i,t} = a_i Y_i + X_{i,t}$, collapses to a single AJD

$$AJD(a_i Y_0 + X_{i,0}, a_i \kappa_{0,Y} + \kappa_{0,i}, \kappa_{1,Y}, \sqrt{a_i} \sigma_Y, l_Y + l_i, a_i \mu_Y).$$

Also, I assume that the size of the systematic parts of the constant term in the drift and of the jump intensity are the same for a given issuer and across issuers such that

$$\omega = \frac{a_i \kappa_{0,Y}}{a_i \kappa_{0,Y} + \kappa_{0,i}} \quad (9)$$

$$= \frac{l_Y}{l_Y + l_i} \quad (10)$$

where $\omega \in [0, 1]$. Restrictions (6)-(10) are identical to those in Mortensen (2006) who generalizes the model in Duffie and Gârleanu (2001) (the a_i s are all equal to one in Duffie and Gârleanu (2001)). The restrictions imply that the default intensity of issuer i is an affine jump diffusion with parameters

$$AJD(a_i Y_t + X_{i,t}, \frac{a_i \kappa_{0,Y}}{\omega}, \kappa_{1,Y}, \sqrt{a_i} \sigma_Y, \frac{l_Y}{\omega}, a_i \mu_Y).$$

As already mentioned, if all a_i s are set to 1 this implies that Y has to be smaller than the smallest intensity among the P issuers at any point in time, seriously restricting the size of Y . Therefore, issuers with small intensities should have small a_i s and I assume that for a given issuer, a_i is equal to the average 5-year CDS spread for this issuer across the estimation period divided by the average 5-year CDS spread across time and issuers. Consequently, if a_i is smaller than one, issuer i has a smaller average 5-year CDS spread compared to the complete pool of P issuers and vice versa. The recovery rate is fixed at 35% and $X_i(t)$ is for each date and issuer chosen such that the 5-year CDS spread is fitted perfectly. As noted by Mortensen (2006) 35% is consistent with historical evidence. However, the precise value is not very critical since a higher recovery rate is offset by a higher default intensity, and only the product of recovery and default intensity is important in the valuation of CDO tranches.

With these assumptions the parameters (under the risk-neutral measure) to be estimated are $\kappa_{0,Y}, \kappa_{1,Y}, \sigma_Y, \mu_Y, l_Y, \omega$ as well as the path the common factor Y_1, \dots, Y_T where T is the number of days in the sample.

To summarize, for each parameter combination of $\kappa_{0,Y}, \kappa_{1,Y}, \sigma_Y, \mu_Y, l_Y, \omega, Y_1, \dots, Y_T$ the procedure for calculating CDO tranche prices is the following:

- Find a_i as the average CDS spread for issuer i in the estimation period divided by the total average CDS spread for all issuers. The a_i s are independent of the specific parameter values and are calculated only once.
- For the first day in the sample, $t = 1$, find for each issuer $X_{i,t}$ such that this issuer's 5-year CDS spread is fitted perfectly. This yields values of the idiosyncratic factors $X_{1,t}, X_{2,t}, \dots, X_{P,t}$ and CDO tranche prices are calculated for this day.
- Repeat the previous step for the rest of the days in the sample, $t = 2, \dots, T$.

3 Estimation

In the empirical section parameters of the intensity-model are estimated using the Bayesian approach MCMC⁵. Parameters and the path and jumps of the latent basic affine process are estimated on basis of a panel data set of CDS premia and CDO tranche prices by writing the model on state space form. To assess model misspecification under the full model, 'marginal models' are also separately estimated on panel data sets of CDS premia for each issuer. The estimation principle is the same for both cases apart from the specific pricing formulas and priors and estimation method are therefore explained without reference to the specific data set.

3.1 MCMC Steps

In order to write the model on state space form, the continuous-time specification in equation (3) is approximated using an Euler scheme

$$\xi_{t+1} - \xi_t = (\kappa_0 + \kappa_1^P \xi_t) \Delta_t + \sigma \sqrt{\Delta_t \xi_t} \epsilon_{t+1}^\xi + J_{t+1} Z_{t+1} \quad (11)$$

where Δ_t is the time between two observations and

$$\begin{aligned} \epsilon_{t+1}^\xi &\sim N(0, 1) \\ Z_{t+1} &\sim \exp(\mu^P) \\ P(J_{t+1} = 1) &= l^P \Delta_t. \end{aligned}$$

⁵For a general introduction to MCMC see Robert and Casella (2004) and for a survey of MCMC methods in financial econometrics see Johannes and Polson (2003).

To simplify notation in the following, I denote $\Theta^Q = (\kappa_0, \kappa_1, l, \mu, \sigma, \omega)$, $\Theta^P = (\kappa_1^P, l^P, \mu^P)$, and $\Theta = (\Theta^Q, \Theta^P)$.

At time $t = 1, \dots, T$ N prices are recorded and they are stacked in the $N \times 1$ vector S_t ⁶. S denotes the $N \times T$ matrix with S_t in the t 'th column. Prices are assumed to be observed with measurement error, so the observation equation is

$$S_t = f(\Theta^Q, \xi_t) + \epsilon_t, \quad \epsilon_t \sim N(0, \Sigma_\epsilon) \quad (12)$$

and f is the pricing formula.

The interest lies in samples from the target distribution $p(\Phi, \Sigma_\epsilon, \xi, J, Z|S)$. The Hammersley-Clifford Theorem (Hammersley and Clifford (1970) and Besag (1974)) implies that samples are obtained from the target distribution by sampling from a number of conditional distributions. Effectively, MCMC solves the problem of simulating from a complicated target distribution by simulating from simpler conditional distributions. If one samples directly from a full conditional the resulting algorithm is the Gibbs sampler (Geman and Geman (1984)). If it is not possible to sample directly from the full conditional distribution one can sample by using the Metropolis-Hastings algorithm (Metropolis et al. (1953)). I use a hybrid MCMC algorithm that combines the two since not all conditional distributions are known. Specifically, the MCMC algorithm is given by (where $\Theta_{\setminus \theta_i}$ is defined as the parameter vector Θ without parameter θ_i)⁷

⁶Specifically, S_t records 5 CDO tranche prices with a maturity of 5 years on each day when estimating the full model while S_t records 3 CDS prices of maturities 1, 3, and 5 year when estimating parameters of individual CDS issuers.

⁷All random numbers in the estimation are draws from Matlab 7.0's generator which is based on Marsaglia and Zaman (1991)'s algorithm. The generator has a period of almost 2^{1430} and therefore the number of random draws in the estimation is not anywhere near the period of the random number generator.

$$\begin{aligned}
p(\theta_i|\Theta_{\setminus\theta_i}^Q, \Theta^P, \Sigma_\epsilon, \xi, J, Z, S) &\sim \text{Metropolis-Hastings} \\
p(\kappa_1^P|\Theta^Q, \Theta_{\setminus\kappa_1^P}^P, \Sigma_\epsilon, \xi, J, Z, S) &\sim \text{Normal} \\
p(l^P|\Theta^Q, \Theta_{\setminus l^P}^P, \Sigma_\epsilon, \xi, J, Z, S) &\sim \text{Beta} \\
p(\mu^P|\Theta^Q, \Theta_{\setminus\mu^P}^P, \Sigma_\epsilon, \xi, J, Z, S) &\sim \text{Inverse Gamma} \\
p(\Sigma_\epsilon|\Theta, \xi, J, Z, S) &\sim \text{Inverse Wishart} \\
p(\xi|\Theta, \Sigma_\epsilon, J, Z, S) &\sim \text{Metropolis-Hastings} \\
p(J|\Theta, \Sigma_\epsilon, \xi, Z, S) &\sim \text{Bernoulli} \\
p(Z|\Theta, \Sigma_\epsilon, \xi, J, S) &\sim \text{Exponential or Restricted Normal}
\end{aligned}$$

Details in the derivations of the conditionals and proposal distributions in the Metropolis-Hastings steps are given in Appendix C. Both the parameters and the latent processes are subject to constraints and if a draw is violating a constraint it can simply be discarded (Gelfand et al. (1992)).

4 Data

This paper uses daily quotes from MarkIt Group Limited. MarkIt receives data from more than 50 global banks and each contributor provides pricing data from its books of record and from feeds to automated trading systems. These data are aggregated into composite numbers after filtering out outliers and stale data and a price is published only if minimum three contributors provide data⁸.

The focus in the empirical part of the paper is on the 5-year Dow Jones CDX North America Investment Grade Index in the period March 21, 2006 to September 20, 2006. The index is updated semiannually and the index for this period is called CDX NA IG Series 6. The maturity of the index is June 20, 2011. Daily CDO tranche prices for the 0–3%, 3–7%, 7–10%, 10–15%, and 15–30% tranches are not available for the first 7 days, so the daily data used in estimation is from March 30, 2006 to September 20, 2006. There are holidays on April 14, April 21, June 3, July 4, and September 4 leaving in total 120 days with prices available. The quoting convention for the equity

⁸Other papers using MarkIt data are Jorion and Zhang (2006), Zhu (2006), and Micu, Remolona, and Wooldridge (2004).

tranche(the 0 – 3% tranche) differs from that of other tranches. Instead of quoting a running premium, the equity tranche is quoted in terms of an up-front fee. Specifically, an up-front fee of 30% means that the investor receives 30% of the tranche notional at time 0 plus a premium of 500 basis points per year paid quarterly.

In addition to CDO tranche prices, 5-year CDS spreads for the 125 underlying constituents are used in the estimation of the multi-name default model.⁹ The number of observations in the estimation of the multi-name default model is therefore 15,600: 125 5-year CDS spreads and 5 CDO tranche prices observed on 120 days.

Also, one-factor affine jump-diffusions are fitted to a panel data set of CDS spreads for each issuer. The panel data used in this estimation is based on daily 1-, 3-, and 5-year CDS premia from October 27, 2003 to October 26, 2006.¹⁰

Table 1 shows summary statistics of the CDS and CDO data.

[Table 1 about here.]

For riskless rates I use LIBOR and swap rates since Feldhütter and Lando (2007) show that swap rates are a more accurate proxy for riskless rates than Treasury yields. 3-, 6-, and 9-months riskless zero coupon bonds are calculated from 3-, 6-, and 9-months LIBOR rates (taking into account money

⁹The underlying contracts for 5-year CDS spreads in the period March 21, 2006 to June 19, 2006 have maturity June 20, 2011, consistent with the maturity of 5-year Dow Jones CDX North America Investment Grade Index whose tranche prices are quoted in the same period. However, for the period June 20-September 19, 2006, the maturity of the underlying 5-year CDS contracts is September 20, 2011. To correct for the maturity mismatch between the CDO and the underlying CDS contracts, all CDS prices are adjusted in the following way: On each date, all CDS spreads are adjusted with the same factor such that the average CDS spread match the CDX NA IG Index level reported by MarkIt. The adjustment factor is almost constant at 0.94 for the period June 20-September 19. For September 20, 2006, the adjustment factor is 0.8843 since the maturity for CDS contracts on this date is December 20, 2011.

¹⁰A few issuers have missing data points in the period, and the missing data points are filled in by linearly interpolating from the nearest dates where CDS premia matching the maturity is available. An alternative to interpolation would be to augment the data set by replacing missing data with simulated data as proposed in Tanner and Wong (1987). However, since the number of missing observations is small, the simpler interpolation procedure is preferred. This results in 52 interpolated data points out of a total of 96,631 data points. Also, four companies do not have data for the full estimation period.

market quoting conventions). For longer maturities riskless bonds are bootstrapped from swap rates. 1-, 2, 3-, 4-, and 5-year swap rates are collected and par rates at semiannual intervals are interpolated using cubic spline. The discount curve at these semiannual intervals is found by bootstrapping the par rates. Finally, discount curve values at other maturities are found by interpolating zero coupon bond prices using again cubic spline.

5 Results

5.1 Model Misspecification

The parsimonious model outlined in section 2.5 is based on a number of simplifying assumptions. An important implication is that CDS spreads of an individual issuer can be adequately described by a one-factor affine jump-diffusion. To see, whether this assumption is reasonable a one-factor model is fit to a panel data set of 1-, 3-, and 5-year CDS spreads over a period of three years for each issuer (the data is described in the previous section 4). In these 125 estimations, the measurement errors in observed CDS spreads are assumed to be uncorrelated across time and maturity and with identical variances, such that Σ_ϵ in (12) is diagonal with a common variance s^2 in the diagonal. In addition, the jump intensity under the risk-neutral measure is set to 1 since it is very difficult to identify l and μ separately. Also, the number of simulations in the burn-in period is 30,000 and in the estimation period 120,000. Every 120'th simulation in the estimation is saved. A summary of parameter estimates are given in Table 2¹¹.

[Table 2 about here.]

A goodness-of-fit measure of the one-factor model for each issuer is given by the estimated parameter s which measures the standard deviation of pricing errors. A low s implies a good fit while a high s implies a bad fit. Since the average CDS bid-ask spreads are around 5-6 basis points, a value of s around

¹¹For a few issuers the estimated mean jump size μ^P under the historical measure is very large. This occurs because there were no jumps in the sample period for these issuers, and therefore there is no information in the data about the size of jumps. Using a flat prior on the mean jump size implies that there is no prior information on the jump size either, resulting in arbitrary estimates.

3 - half the bid-ask spread - indicates a very good fit¹². As seen in Table 2, the average value of s is 3.72 across all estimations and therefore a one-factor affine jump diffusion model provides a good fit to the term structure of CDS spreads of the issuers underlying the CDX NA IG Index. However, there are a few issuers where the one-factor model provides an inaccurate description of the CDS term structure over time. The firm with the worst goodness-of-fit is Harrah's Operating Company, a company within the casino entertainment industry, since the estimated parameter $s = 11.1$ is highest among all firms.

The actual and fitted CDS spreads along with jump probabilities are shown in Figure 1. We see that CDS spreads are fitted reasonably well, but on October 2, 2006, a jump in CDS spreads occurred and they more than doubled on this day. This jump is due to an acquisition proposal by Apollo Management and Texas Pacific Group, an offer that was ultimately accepted on December 19. After the acquisition the debt in Harrah's increased and foreseeing this debt increase Fitch on October 2, 2006 downgraded Harrah's from investment grade to speculative grade.

The one-factor model is incapable of matching CDS spreads after the acquisition offer, which suggests that a one-factor model for default risk is too simple a model for a firm that switches from investment grade to speculative grade. Consistent with this view, Duffee (1999) finds that default risk is more explosive under the equivalent martingale measure for low-rated firms than for high-rated firms. However, as seen in Figure 1 spreads are reasonably matched before the offer - including the period for which the CDO estimation is done.

[Figure 1 about here.]

American International Group, a major American insurance company, is an example of an issuer where the one-factor affine jump-diffusion model is a sufficiently rich model to capture the dynamics of CDS spreads as seen in Figure 2. There are several jumps in spreads and periods of both high and low spreads, but in all periods the dispersion in spreads across maturities are well-matched.

[Figure 2 about here.]

¹²On February 9 2007, the average bid-ask spread on 5-year CDS contracts underlying the CDX NA IG Series 7 was 5.44 basis points according to CreditFlux.

The evidence in the tables along with the two examples suggest that a one-factor jump-diffusion model is a good description of the majority of issuers underlying the Dow Jones North American Investment Grade Index.

In addition to the assumption that CDS spreads can be modelled as one-factor processes section 2.5 outlined further assumptions regarding the parameters governing the risk-neutral dynamics of the processes. Several assumptions relate to the choice of issuer specific a_i 's defined in equation (1) and chosen as the issue specific average 5-year CDS spread over the estimation period divided by the total average 5-year CDS spread. They are:

- *Jump intensities l_i are identical across issuers.* Assumption (10) implies that the jump intensities of all issuers are the same. As mentioned previously it is possible to identify the product of jump intensity and size, $l_i \times \mu_i$, but difficult to identify each component separately. Therefore, it is in practice no restriction on the dynamics of default intensities to restrict the jump intensities to be the same.
- *The parameter $\kappa_{0,i}$ is linear in a_i .* Assumption (9) states that $\kappa_{0,i}$ is linear in a_i and the lower-right graph in Figure 3 shows that the assumption is a reasonable. This is underpinned by the fact that the confidence bands of $\kappa_{0,i}$ contain the fitted line for 98 out of the 125 issuers.
- *Jump product $l_i \times \mu_i$ is linear in a_i .* Assumptions (8) and (10) imply that the jump intensity and size products $l_i \times \mu_i$ are linear in the a_i s, so an issuer with a high $l_i \times \mu_i$ should have a high a_i and vice versa. The product $l_i \times \mu_i$ is plotted against a_i for each issuer in the top-right graph in Figure 3 along with a fitted line illustrating the linear relationship. Although there are some issuers for which their product $l_i \times \mu_i$ is too low relative to their a_i 's the graph shows a linear relationship and the confidence bands of $l_i \times \mu_i$ contain the fitted line for 89 out of the 125 issuers.
- *Diffusion volatility σ_i^2 is linear in a_i .* The diffusion volatility is linear in the a_i 's according to assumption (7) such that issuers with high default probabilities have high diffusion volatility. The top-left graph in Figure 3 plots σ_i^2 against a_i for each issuer. The graph shows that high default probability implies high diffusion volatility although the linear relationship is not perfect. The confidence bands of σ_i^2 contain the fitted line for 61 out of 125 issuers.

- *Mean reversion $\kappa_{1,i}$ is constant across issuers.* Finally, the mean reversion coefficient of all issuers is the same according to assumption (6) and the bottom-left graph in Figure 3 plots $\kappa_{1,i}$ against a_i for each issuer. A fitted horizontal line is drawn to illustrate that $\kappa_{1,i}$ and a_i are unrelated for each issuer. The graph shows that the mean-reversion coefficient is unrelated to a_i . Also, 114 out of 125 issuers have a positive mean reversion coefficient implying an explosive process. Since the sign of the mean reversion coefficient is largely the same across issuers and the coefficient is in the range of around 0-0.6 there is a large degree of homogeneity regarding the mean reversion parameter. However, the confidence bands of $\kappa_{1,i}$ contain the fitted line for only 10 out of 125 issuers. Overall, it is reasonable to assume that the mean reversion coefficients have the same sign and they lie in the range of 0-0.6, but they are estimated with high precision and statistically they are different from each other.

[Figure 3 about here.]

Overall, this section shows that a one-factor affine jump-diffusion model provides a good fit to investment grade single-name issuers and that most of the assumptions outlined in section 2.5 are reasonable. The most critical assumption is that the mean reversion coefficients of all issuers are identical: Although more than 90% of the issuers have a positive coefficient of 0 – 0.6, implying explosive processes, they are statistically clearly different.

5.2 CDO Parameter Estimates and Pricing Results

The multi-name default model is estimated on basis of panel data set of 120 daily CDO tranche prices and CDS contracts as described in section 4. The measurement error matrix Σ_ϵ in (12) is assumed to be diagonal and the prior on the measurement errors is that they are equal to the bid-ask spreads of the tranches¹³. The number of simulations in the burn-in period is 20,000 and in the estimation period 10,000 and every 10'th simulation is saved.

¹³Bid-ask spreads in Mortensen (2006) are used: 0.8% for the 0-3% tranche, 6.8bp for the 3-7% tranche, 5.4bp for the 7-10% tranche, 3bp for the 10-15% tranche, and 1bp for the 15-30% tranche. The prior is chosen to be inverse Wishart distributed with parameters $V = \text{diag}(0.008^2, 0.068^2, 0.054^2, 0.03^2, 0.01^2)$ and $m = 10,000$. Since $m = 10,000$ the prior is weighted strongly compared to the data. Further estimation details are given in C.3.

Parameter estimates are given in Table 3. In the table we see that $\sigma = 3.67$ in the CDO estimation is close to firm averages of 3.56. We also see that $\kappa_1 = 0.465$ which implies that the default intensity is explosive consistent with the underlying firms' default dynamics. The value is higher than the average value of 0.164 for the individual firms, but is not unreasonably high as the estimates in Table 2 shows. Since the average value of a_i is 1, the values of κ_0 , $l\mu$, and σ^2 should match the average parameter values of the underlying firms. $\kappa_0 = 1.59 \times 10^{-5}$ is estimated at a somewhat low value compared to the firm average of 4.9×10^{-4} but Table 2 shows that a number of firms have a lower value of κ_0 than the value implied from the CDO estimation. The product $l\mu = 3.92 \times 10^{-3}$ is estimated at a higher value than the firm average of 7.29×10^{-4} and it is higher than the estimated product for any of the underlying issuers.

The parameter estimates in the CDO estimation are largely consistent with the estimates from the univariate CDS premium estimations. However, the common factor in the CDO estimation is more explosive and has a higher product of jump intensity and jump size than implied from the individual CDS estimations. Since both estimations match the 5-year CDS premia, it must be that the initial value of the common factor is lower in the CDO estimation. For maturities up to five years the average CDS premia in the CDO estimation are consistent with actual average CDS premia - the lower value of the common factor is offset by a higher product of jump intensity and jump sizes. For maturities larger than five years the implied CDS premia in the CDO model are higher than those implied from the univariate estimations because the mean aversion coefficient $\kappa_1 = 0.465$ plays a dominant role at longer maturities. Consequently, it is likely that CDO tranche prices at maturities more than five years would be overestimated.

[Table 3 about here.]

The pricing ability of the affine model can be examined by means of average pricing errors and RMSEs, which are given in Table 4. We see that on average the model has a tendency to produce spreads that are slightly too low for the 7–10% tranche but for the other four tranches the model matches average spreads well. Comparing the RMSEs with the only other study that has conducted a time series analysis of CDO tranche prices, Longstaff and Rajan (2006), it appears that the affine model prices senior tranches more accurately than the three-factor portfolio model in their paper while

lower tranches are priced more accurately in their model¹⁴. Since spreads of the five tranches are very different in size, a more relevant measure is the relative pricing errors and RMSEs which are also given in table 4. The relative mean errors show that the 10 – 15% tranche has the worst relative fit and the relative RMSEs show in contrast to absolute RMSEs that the most risky tranches are fitted better than the senior tranches. The conclusions from absolute pricing error results are more or less reversed when looking at relative pricing errors. Figure 4 shows actual and fitted tranche prices over time and the figure shows why relative pricing errors for senior tranches are high: the model prices the first two tranches accurately over time while there is not enough variability in model-implied senior spreads to match actual senior spreads. This aspect of the model will also show up in the hedging results in the next section.

[Table 4 about here.]

[Figure 4 about here.]

Overall, the model has a good fit to the prices of the first two tranches, underestimates the prices of the third tranche slightly, and captures the average price level of the two most senior tranches. However, the model misses the price variation in the two senior tranches. What makes it possible for the model to fit the average senior spreads is the inclusion of jumps. As noted in Mortensen (2006), without jumps the model is a pure diffusion model and is unable to generate enough default correlation to match senior tranche spreads. Therefore, it is the jump parameters that govern senior tranche prices; basically it is the product $l \times \mu$. Since both jump intensity l and expected jump size μ are constant, the senior tranche spreads are almost constant.

The model can in a parsimonious way be changed such that the time series properties of senior spreads are better captured. If jump intensities are made state dependent, such that the jump intensity for the process in equation (2) is changed from $l > 0$ to $l_0 + l_1\xi$, with $l_0 > 0$ and $l_1 > 0$, the model is still within the affine framework. This would generate more time variation in senior spreads since the spreads would be positively correlated

¹⁴For the 15 – 30% tranche the RMSE in this paper is 0.74 while the average RMSE in their paper is 3.59. For the 10 – 15% tranche the RMSE in this paper is 2.09 while the average RMSE in their paper is 4.42.

with the realization of the common factor through l^1 . However, senior spreads are negatively correlated with the common factor in the current model, and a second common factor (with time-varying jump intensity) would likely be necessary to accurately capture senior spread dynamics. This is an interesting topic for future research.

6 Hedging

In correlation trading investors take views on the future direction of credit correlation and structure trades that are exposed to the level of default correlation but hedged against small movements in credit spreads¹⁵. Hedging changes in market value of for example the equity tranche due to changes in the underlying CDS spreads can be done by entering an offsetting position in 1) another tranche, typically the 3 – 7% tranche, 2) the CDX index, or 3) the underlying CDS contracts. However, the 'delta', the amount invested in the offsetting position, is model-dependent and the quality of the spread hedge depends strongly on the quality of the model used to calculate deltas. Therefore, for such correlation trading as well as for risk management in general, it is important that the model used correctly specifies the relation between changes in the underlying CDS spreads and changes in the market value of CDO tranches.

The literature on CDO modeling has generally focused on the ability of models to fit tranche prices on a single or a few dates while the accuracy of model-implied deltas is not well examined. An exception is Houdain and Guegan (2006) who examine different models' ability to hedge an equity tranche by an offsetting position in the first mezzanine tranche. However, whether one hedges a CDO tranche with another tranche, the CDX index, or the underlying CDS contracts, a necessary condition for a successful hedge is that the model correctly predicts the change in CDO tranche price with respect to any changes in the underlying CDS spreads. I therefore conduct the following test of the hedging ability of a pricing model: For a given day the spread of a CDO tranche is calculated. Keeping the parameters of the model fixed a new spread is calculated based on the CDS spreads of tomorrow. The difference between the two spread is the model-predicted change in tranche spread due to spread changes in the underlying CDS contracts. Finally, the model-predicted spread change is compared with the actual spread

¹⁵See for example Merrill Lynch (2003) and Belsham, Vause, and Wells (2005).

change. The average difference between model-implied and actual spread changes should be close to zero and comparing two models, the better model should have lower RMSEs. This procedure measures how well a model hedges spread risk, but even a perfect model will have positive RMSEs since changes in tranche spreads due to changes in correlation is not accounted for (and should not in a correlation trade)¹⁶. The hedge ratios do not represent an actual hedging strategy since the hedge ratios at time t are constructed using information at time $t + 1$, but they do strongly indicate the accuracy of 'real-time' hedge strategies.

The commonly used one-factor Gaussian model with homogeneous default intensities is chosen as a benchmark model in comparing the hedging ability of the affine model. The model plays a role in the CDO market similar to that of the Black-Scholes model in the options market, and a brief review of the model is given in Appendix D. For every day and tranche, the correlation coefficient in the Gaussian model is fitted to the actual tranche price on that day and the model-predicted spread change is calculated on basis of this correlation.

In addition to the Gaussian model, I also report hedging results for a simple Random Walk where the model-implied price change is zero every day. Table 5 presents the results.

[Table 5 about here.]

In the table we see that the hedging performances of the affine and Gaussian model are comparable for the 3 – 7% tranche, while the affine model is significantly better than the Gaussian model in hedging the equity tranche. Figure 5 shows why the affine model does better in hedging the equity tranche. The deltas of the Gaussian model are too low while the affine model has deltas matching actual deltas better. This is seen in the figure because model-implied spread changes in the affine model in general

¹⁶For simplicity, I look at changes in tranche spreads instead of changes in tranche prices. The change in market value is given as the change in spread times *risky duration* (see Houdain and Guegan (2006) for a definition of risky duration). The equity tranche is an exception because a change in up-front fee is a direct measure of a change in market value. To see this, consider entering an equity tranche as protection buyer at time t and paying an up-front fee of U_t . At time $t + 1$ the position can be closed by entering an equity tranche as protection seller and receiving an up-front fee of U_{t+1} . All future cash flows are matched and at $t + 1$ the gain/loss is $U_{t+1} - U_t$ times the notional (ignoring yield on U_t between t and $t + 1$).

match actual spread changes well, while the price changes in the Gaussian copula are too small in response to CDS spread changes. Therefore, using a Gaussian copula in hedging spread risk in the equity tranche underestimates the price sensitivity to CDS spread changes, while the affine model better matches the sensitivities.

From the table we also see that both in terms of average hedging errors and RMSEs the Gaussian copula hedges senior tranches better than the affine model. In fact, the affine model is only slightly better than a Random Walk. In light of the discussion in the previous section, namely that the prices of senior tranches in the affine model are almost constant, it is not surprising that the affine model does not hedge well.

[Figure 5 about here.]

7 Conclusion

I estimate a dynamic intensity-based model for multi-name default on an extensive data set of 15,600 CDS and CDO tranche spreads on the Dow Jones North American Investment Grade Index. The modeling framework was first proposed by Duffie and Gârleanu (2001) and the estimated model allows heterogeneous default intensity dynamics as in Mortensen (2006).

The empirical results document that the assumptions underlying the model are reasonable and that the model can match average CDO tranche spreads. The variation over time in actual tranche spreads is matched well for the most risky tranches and the model's sensitivities in the equity tranche with respect to underlying CDS premium movements is more in accordance with actual sensitivities than in the Gaussian copula. The last result strongly suggests that the intensity-based model is more accurate in hedging the equity tranche than the Gaussian copula.

The empirical study in this paper points to the importance of testing CDO pricing models along several dimensions such as model assumptions, pricing ability in both the cross section and time series, and hedging performance. For the model examined in this paper, the results suggest that incorporating time-varying jump intensity can in a parsimonious way improve pricing and hedging performance of senior tranches. It is left for future research to modify and test a model with this modification.

A CDS Pricing

This section briefly explains how to price credit default swaps. A more thorough introduction is given in Duffie (1999).

A CDS contract is an insurance agreement between two counterparties written on the default event of a specified underlying reference obligation. The protection buyer pays fixed premium payments periodically until a default occurs or the contracts expires, whichever happens first. If default occurs, the protection buyer delivers the reference obligation to the protection seller in exchange for face value.

A number of assumptions are made in order to price the CDS contract. First, assume that the recovery rate δ is constant. Second, assume that default-free interest rates and default probabilities are independent (under the risk-neutral measure). Third, assume that default occurs halfway between two premium payments.

To be specific, consider a CDS with maturity T and denote ρ the time between two premium payments (the CDS contracts considered in the main text have quarterly payments, i.e. $\rho = \frac{1}{4}$). If τ is the time of default and premium payments occur at $t_1, t_2, \dots, t_{T/\rho}$, the value of the protection leg at time $t_0 = 0$ is

$$\begin{aligned} Prot(0, T) &= E^Q \left[\sum_{j=1}^{T/\rho} \exp \left(- \int_0^{t_{j-1} + \frac{1}{2}\rho} r_s ds \right) (1 - \delta) 1_{\tau \in (t_{j-1}, t_j)} \right] \\ &= (1 - \delta) \sum_{j=1}^{T/\rho} P(0, t_{j-1} + \frac{1}{2}\rho) Q[\tau \in (t_{j-1}, t_j)]. \end{aligned}$$

If S is the CDS price the protection buyer pays ρS until default or contract maturity. At default the protection buyer pays $\frac{1}{2}\rho S$. The value of the premium leg is therefore

$$\begin{aligned} Prem(0, T; S) &= E^Q \left[\sum_{j=1}^{T/\rho} [S\rho \exp \left(- \int_0^{t_j} r_s ds \right) 1_{\tau > t_j} \right. \\ &\quad \left. + \frac{1}{2} S\rho \exp \left(- \int_0^{t_{j-1} + \frac{1}{2}\rho} r_s ds \right) 1_{\tau \in (t_{j-1}, t_j)}] \right] \\ &= S\rho \sum_{j=1}^{T/\rho} \left(P(0, t_j) Q(\tau > t_j) + \frac{1}{2} P(0, t_{j-1} + \frac{1}{2}\rho) Q[\tau \in (t_{j-1}, t_j)] \right) \end{aligned}$$

and the CDS premium S is found by letting the value of the protection leg equal the value of the premium leg, $Prot(0, T) = Prem(0, T; S)$.

B Default Probabilities for Affine Jump Diffusion

For the one-dimensional affine jump-diffusion X_t ,

$$dX_t = (\kappa_0 + \kappa_1 X_t)dt + \sigma \sqrt{X_t} dW_t + dJ_t,$$

whose jump times are those of a Poisson process with intensity l and the jump sizes are independent of the jump times and follow an exponential distribution with mean μ , the following result follows from Duffie and Gârleanu (2001),

$$E[\exp(-\int_t^{t+s} X_u du)] = e^{\alpha(s) + \beta(s)X_t},$$

where

$$\begin{aligned} \beta(s) &= \frac{\lambda_1(s)}{\lambda_2(s)}, \\ \alpha(s) &= -\frac{2\kappa_0}{\sigma^2} \log\left(\frac{-\lambda_2(s)}{\gamma}\right) + \frac{\kappa_0}{c_1} s + \\ &\quad \frac{-2l\mu}{\sigma^2 + 2\mu\kappa_1 - 2\mu^2} \log\left(\frac{\lambda_1(s)\mu - \lambda_2(s)}{\gamma}\right) + \left(\frac{l\mu}{c_1 - \mu}\right)s, \\ \gamma &= \sqrt{\kappa_1^2 + 2\sigma^2} \\ \lambda_1(s) &= 1 - e^{-\gamma s} \\ \lambda_2(s) &= \frac{1}{2}(\kappa_1 + \gamma)\lambda_1(s) - \gamma \\ c_1 &= \frac{1}{2}(\kappa_1 - \gamma) \end{aligned}$$

C Conditional Posteriors in MCMC Estimation

In this Appendix the conditional posteriors stated in the main text and used in MCMC estimation are derived. Bayes' rule

$$p(X|Y) \propto p(Y|X)p(X)$$

is repeatedly used in the calculations.

C.1 Conditionals of $S, \xi, J,$ and Z

The conditional posteriors of $S, \xi, J,$ and Z are used in most of the conditional posteriors for the parameters and are therefore derived in this section.

C.1.1 $p(\xi|\Theta, \Sigma_\epsilon, J, Z)$ and $p(S|\Theta, \Sigma_\epsilon, \xi, J, Z)$

With the discretization in (11) we have that

$$\begin{aligned} p(\xi|\Theta, \Sigma_\epsilon, J, Z) &= \left(\prod_{t=1}^T p(\xi_t|\xi_{t-1}, \Theta, \Sigma_\epsilon, J, Z) \right) p(\xi_0) \\ &= p(\xi_0) \prod_{t=1}^T \frac{1}{\sigma \sqrt{\Delta_t \xi_{t-1}}} \exp \left(-\frac{1}{2} \frac{[\xi_t - (\kappa_0 \Delta_t + (\kappa_1^P \Delta_t + 1)\xi_{t-1} + J_t Z_t)]^2}{\sigma^2 \Delta_t \xi_{t-1}} \right) \\ &\propto p(\xi_0) \sigma^{-T} \xi_x^{-\frac{1}{2}} \exp \left(-\frac{1}{2} \sum_{t=1}^T \frac{[\xi_t - (\kappa_0 \Delta_t + (\kappa_1^P \Delta_t + 1)\xi_{t-1} + J_t Z_t)]^2}{\sigma^2 \Delta_t \xi_{t-1}} \right) \end{aligned} \quad (13)$$

where $\xi_x = \prod_{t=1}^T \xi_{t-1}$. Note that the posterior $p(\xi|\Theta, \Sigma_\epsilon, J, Z)$ differs from $p(\xi|\Theta, \Sigma_\epsilon, J, Z, S)$.

The conditional posterior of S is found as

$$\begin{aligned} p(S|\Theta, \Sigma_\epsilon, \xi, J, Z) &= \prod_{t=1}^T |\Sigma_\epsilon|^{-\frac{1}{2}} \exp \left(-\frac{1}{2} [S_t - f(\Theta^Q, \xi_t)]' \Sigma_\epsilon^{-1} [S_t - f(\Theta^Q, \xi_t)] \right) \\ &= |\Sigma_\epsilon|^{-\frac{T}{2}} \exp \left(-\frac{1}{2} \sum_{t=1}^T \hat{e}_t' \Sigma_\epsilon^{-1} \hat{e}_t \right), \end{aligned} \quad (14)$$

where $\hat{e}_t = S_t - f(\Theta^Q, \xi_t)$. If Σ_ϵ is diagonal this simplifies to

$$p(S|\Theta, \Sigma_\epsilon, \xi, J, Z) \propto \prod_{i=1}^N \Sigma_{\epsilon, ii}^{-\frac{T}{2}} \exp\left(-\frac{1}{2\Sigma_{\epsilon, ii}} \sum_{t=1}^T \hat{e}_{t,i}^2\right).$$

This posterior does not depend on J, Z, κ_0^P , and κ_1^P .

C.1.2 $p(Z|\Theta, \Sigma_\epsilon, \xi, J, S)$ and $p(J|\Theta, \Sigma_\epsilon, \xi, Z, S)$

Since Z_t is exponentially distributed we have that

$$\begin{aligned} p(Z|\Theta, \Sigma_\epsilon, \xi, J, S) &\propto p(S|\Theta, \Sigma_\epsilon, \xi, J, Z)p(Z|\Theta, \Sigma_\epsilon, \xi, J) & (15) \\ &\propto p(\xi|\Theta, \Sigma_\epsilon, J, Z)p(Z|\Theta, \Sigma_\epsilon, J) \\ &\propto p(\xi|\Theta, \Sigma_\epsilon, J, Z) \prod_{t=1}^T \frac{1}{\mu^P} \exp\left(-\frac{Z_t}{\mu^P}\right) \\ &\propto p(\xi|\Theta, \Sigma_\epsilon, J, Z)(\mu^P)^{-T} \exp\left(-\frac{Z_\bullet}{\mu^P}\right) & (16) \end{aligned}$$

where $Z_\bullet = \sum_{t=1}^T Z_t$.

The jump time J_t can only take on two values so the conditional posterior for J_t is Bernoulli. The Bernoulli probabilities are given as

$$\begin{aligned} p(J|\Theta, \Sigma_\epsilon, \xi, Z, S) &\propto p(S|\Theta, \Sigma_\epsilon, \xi, J, Z)p(J|\Theta, \Sigma_\epsilon, \xi, Z) & (17) \\ &\propto p(\xi|\Theta, \Sigma_\epsilon, J, Z)p(J|\Theta, \Sigma_\epsilon, Z) \\ &\propto p(\xi|\Theta, \Sigma_\epsilon, J, Z)p(J|\Theta) \\ &\propto p(\xi|\Theta, \Sigma_\epsilon, J, Z) \prod_{t=1}^T \left((l^P \Delta_t)^{J_t} (1 - l^P \Delta_t)^{1-J_t} \right) \\ &\propto p(\xi|\Theta, \Sigma_\epsilon, J, Z)(l^P \Delta_t)^{J_\bullet} (1 - l^P \Delta_t)^{T-J_\bullet} & (18) \end{aligned}$$

with $J_\bullet = \sum_{t=1}^T J_t$

C.2 Conditional Posteriors

The conditional posteriors are derived and the choice of priors for the posteriors are discussed in this section.

1. The conditional posterior of the error matrix Σ_ϵ is given as

$$\begin{aligned}
p(\Sigma_\epsilon|\Theta, \xi, J, Z, S) &\propto p(S|\Theta, \Sigma_\epsilon, \xi, J, Z)p(\Sigma_\epsilon|\Theta, \xi, J, Z) \\
&\propto p(S|\Theta, \Sigma_\epsilon, \xi, J, Z)p(\Sigma_\epsilon|\Theta) \\
&\propto |\Sigma_\epsilon|^{-\frac{T}{2}} \exp\left(-\frac{1}{2}\sum_{t=1}^T \hat{e}'_t \Sigma_\epsilon^{-1} \hat{e}_t\right) p(\Sigma_\epsilon|\Theta) \\
&= |\Sigma_\epsilon|^{-\frac{T}{2}} \exp\left(-\frac{1}{2}\text{tr}(\Sigma_\epsilon^{-1} \sum_{t=1}^T \hat{e}_t \hat{e}'_t)\right) p(\Sigma_\epsilon|\Theta).
\end{aligned}$$

The last line follows because $-\frac{1}{2}\sum_{t=1}^T \hat{e}'_t \Sigma_\epsilon^{-1} \hat{e}_t = -\frac{1}{2}\sum_{t=1}^T \text{tr}(\hat{e}'_t \Sigma_\epsilon^{-1} \hat{e}_t) = -\frac{1}{2}\sum_{t=1}^T \text{tr}(\Sigma_\epsilon^{-1} \hat{e}_t \hat{e}'_t) = -\frac{1}{2}\text{tr}(\Sigma_\epsilon^{-1} \sum_{t=1}^T \hat{e}_t \hat{e}'_t) = -\frac{1}{2}\text{tr}(\Sigma_\epsilon^{-1} \sum_{t=1}^T \hat{e}_t \hat{e}'_t)$. If the prior on Σ_ϵ is independent of the other parameters and has an inverse Wishart distribution with parameters V and m then $p(\Sigma_\epsilon|\dots)$ is inverse Wishart distributed with parameters $V + \sum_{t=1}^T \hat{e}_t \hat{e}'_t$ and $T + m$. The special case of V equal to the zero matrix and $m = 0$ corresponds to a flat prior.

2. The conditional posterior of κ_1^P is found as

$$\begin{aligned}
p(\kappa_1^P|\Theta_{\setminus \kappa_1^P}, \Sigma_\epsilon, \xi, J, Z, S) &\propto p(S|\Theta, \Sigma_\epsilon, \xi, J, Z)p(\kappa_1^P|\Theta_{\setminus \kappa_1^P}, \Sigma_\epsilon, \xi, J, Z) \\
&\propto p(\kappa_1^P|\Theta_{\setminus \kappa_1^P}, \Sigma_\epsilon, \xi, J, Z) \\
&\propto p(\xi|\Theta, \Sigma_\epsilon, J, Z)p(\kappa_1^P|\Theta_{\setminus \kappa_1^P}, \Sigma_\epsilon).
\end{aligned}$$

According to equation (13) we have

$$p(\kappa_1^P|\dots) \propto \exp\left(-\frac{1}{2}\sum_{t=1}^T \frac{[\xi_t - (\kappa_0 \Delta_t + (\kappa_1^P \Delta_t + 1)\xi_{t-1} + J_t Z_t)]^2}{\sigma^2 \Delta_t \xi_{t-1}}\right) p(\kappa_1^P|\Theta_{\setminus \kappa_1^P}, \Sigma_\epsilon)$$

so

$$p(\kappa_1^P|\dots) \propto \exp\left(-\frac{1}{2}\sum_{t=1}^T \frac{[a_t \kappa_1^P - b_t]^2}{\sigma^2 \Delta_t \xi_{t-1}}\right) p(\kappa_1^P|\Theta_{\setminus \kappa_1^P}, \Sigma_\epsilon)$$

where

$$\begin{aligned}
a_t &= -\Delta_t \xi_{t-1} \\
b_t &= \kappa_0 \Delta_t + \xi_{t-1} + J_t Z_t - \xi_t.
\end{aligned}$$

Using the result in Frühwirth-Schnatter and Geyer (1998) p.10 and assuming flat priors we have that $\kappa_1^P \sim N(Qm, Q)$ where

$$m = \sum_{t=1}^T \frac{a_t b_t}{\sigma^2 \Delta_t \xi_{t-1}}$$

$$Q^{-1} = \sum_{t=1}^T \frac{a_t^2}{\sigma^2 \Delta_t \xi_{t-1}}.$$

3. For the jump size parameter μ^P the conditional posterior is found as

$$\begin{aligned} p(\mu^P | \Theta_{\setminus \mu^P}, \Sigma_\epsilon, \xi, J, Z, S) &\propto p(S | \Theta, \Sigma_\epsilon, \xi, J, Z) p(\mu^P | \Theta_{\setminus \mu^P}, \Sigma_\epsilon, \xi, J, Z) \\ &\propto p(\xi | \Theta, \Sigma_\epsilon, J, Z) p(\mu^P | \Theta_{\setminus \mu^P}, \Sigma_\epsilon, J, Z) \\ &\propto p(Z | \Theta, \Sigma_\epsilon, J) p(\mu^P | \Theta_{\setminus \mu^P}, \Sigma_\epsilon, J) \\ &\propto p(Z | \Theta) p(\mu^P | \Theta_{\setminus \mu^P}, \Sigma_\epsilon) \\ &\propto (\mu^P)^{-T} \exp\left(-\frac{Z_\bullet}{\mu^P}\right) p(\mu^P | \Theta_{\setminus \mu^P}, \Sigma_\epsilon). \end{aligned}$$

If the prior on μ^P is flat then the conditional posterior inverse gamma distributed with parameters Z_\bullet and $T - 1$.

4. The same calculations as for the jump-size parameter μ^P yields the conditional posterior of the jump-time parameter l^P as

$$\begin{aligned} p(l^P | \Theta_{\setminus l^P}, \Sigma_\epsilon, \xi, J, Z, S) &\propto p(J | \Theta) p(l^P | \Theta_{\setminus l^P}, \Sigma_\epsilon) \\ &\propto \left((l^P \Delta_t)^{J_\bullet} (1 - l^P \Delta_t)^{T - J_\bullet} \right) p(l^P | \Theta_{\setminus l^P}, \Sigma_\epsilon). \end{aligned}$$

Assuming a flat prior on l^P the conditional posterior of $l^P \Delta_t$ is beta distributed, $l^P \Delta_t \sim B(J_\bullet + 1, T - J_\bullet + 1)$.

5. The parameters σ and κ_0 are sampled by Metropolis-Hastings since the conditional distributions are not known. Denoting any of the two parameters θ_i , the conditional distribution is found as

$$\begin{aligned} p(\theta_i | \Theta_{\setminus \theta_i}, \Sigma_\epsilon, \xi, J, Z, S) &\propto p(S | \Theta, \Sigma_\epsilon, \xi, J, Z) p(\theta_i | \Theta_{\setminus \theta_i}, \Sigma_\epsilon, \xi, J, Z) \\ &\propto p(S | \Theta, \Sigma_\epsilon, \xi, J, Z) p(\xi | \Theta, \Sigma_\epsilon, J, Z) p(\theta_i | \Theta_{\setminus \theta_i}, \Sigma_\epsilon, J, Z) \\ &\propto p(S | \Theta, \Sigma_\epsilon, \xi, J, Z) p(\xi | \Theta, \Sigma_\epsilon, J, Z) p(\theta_i | \Theta_{\setminus \theta_i}, \Sigma_\epsilon). \end{aligned}$$

Flat priors on both parameters are assumed.

6. The parameters κ_1^Q , l^Q , and μ^Q are sampled by Metropolis-Hastings. The only difference in the derivation of their conditional distributions compared to derivation of the distributions of σ and κ_0 is that the distribution of ξ does not depend on these three parameters. Letting θ_i represent any of the three parameters, the conditional distribution is found as

$$\begin{aligned} p(\theta_i|\Theta_{\setminus\theta_i}, \Sigma_\epsilon, \xi, J, Z, S) &\propto p(S|\Theta, \Sigma_\epsilon, \xi, J, Z)p(\theta_i|\Theta_{\setminus\theta_i}, \Sigma_\epsilon, \xi, J, Z) \\ &\propto p(S|\Theta, \Sigma_\epsilon, \xi, J, Z)p(\xi|\Theta, \Sigma_\epsilon, J, Z)p(\theta_i|\Theta_{\setminus\theta_i}, \Sigma_\epsilon, J, Z) \\ &\propto p(S|\Theta, \Sigma_\epsilon, \xi, J, Z)p(\theta_i|\Theta_{\setminus\theta_i}, \Sigma_\epsilon). \end{aligned}$$

Flat priors on all three parameters are assumed.

7. The latent jump indicators J_t 's are sampled individually from Bernoulli distributions. To see this, note that equation (18) implies that

$$\begin{aligned} p(J|\Theta, \Sigma_\epsilon, \xi, Z, S) \\ \propto \prod_{t=1}^T \exp\left(-\frac{1}{2} \frac{[\xi_t - (\kappa_0\Delta_t + (\kappa_1^P\Delta_t + 1)\xi_{t-1} + J_t Z_t)]^2}{\sigma^2\Delta_t\xi_{t-1}}\right) \left(\frac{l^P\Delta_t}{1 - l^P\Delta_t}\right)^{J_t}. \end{aligned}$$

In the actual implementation I use

$$\begin{aligned} p(J|\Theta, \Sigma_\epsilon, \xi, Z, S) \\ \propto \prod_{t=1}^T \exp\left(-\frac{1}{2} \frac{(-2[\xi_t - (\kappa_0\Delta_t + (\kappa_1^P\Delta_t + 1)\xi_{t-1})] + J_t Z_t)J_t Z_t}{\sigma^2\Delta_t\xi_{t-1}}\right) \left(\frac{l^P\Delta_t}{1 - l^P\Delta_t}\right)^{J_t} \end{aligned}$$

since this is numerically more robust.

8. For the latent jump sizes Z_t we have according to equation (16) that

$$p(Z|\Theta, \Sigma_\epsilon, \xi, J, S) \propto \prod_{t=1}^T \exp\left(-\frac{1}{2} \frac{[\xi_t - (\kappa_0\Delta_t + (\kappa_1^P\Delta_t + 1)\xi_{t-1} + J_t Z_t)]^2}{\sigma^2\Delta_t\xi_{t-1}} - \frac{Z_t}{\mu^P}\right)$$

so the Z_t s are conditionally independent and are sampled individually.

If $J_t = 0$ then Z_t is sampled from an exponential distribution with mean μ^P . If $J_t = 1$ tedious calculations show that

$$p(Z_t|\Theta, \Sigma_\epsilon, \xi, J, Z_{\setminus Z_t}, S) \propto \frac{[(\kappa_1^P + \mu^P\sigma^2)\Delta_t + 1]\xi_{t-1} - (\xi_t - \kappa_0\Delta_t) + Z_t^2}{\sigma^2\Delta_t\xi_{t-1}},$$

where $Z_t \geq 0$. Therefore, Z_t is drawn from a $N((\xi_t - \kappa_0 \Delta_t) - ((\kappa_1^P + \mu^P \sigma^2) \Delta_t + 1) \xi_{t-1}, \sigma^2 \Delta_t \xi_{t-1})$ distribution and the draw is rejected if $Z_t < 0$. In practice the number of rejections are small¹⁷.

9. The latent ξ_t s are sampled individually by Metropolis-Hastings and for $t = 1, \dots, T - 1$ the conditional posterior is

$$\begin{aligned} p(\xi_t | \Theta, \Sigma_\epsilon, \xi_{\setminus \xi_t}, J, Z, S) &\propto p(S | \Theta, \Sigma_\epsilon, \xi, J, Z, S) p(\xi_t | \Theta, \Sigma_\epsilon, \xi_{\setminus \xi_t}, J, Z) \\ &\propto p(S_t | \Theta, \Sigma_\epsilon, \xi_t, J, Z, S) p(\xi_t | \Theta, \Sigma_\epsilon, \xi_{t-1}, \xi_{t+1}, J, Z) \\ &\propto p(S_t | \Theta, \Sigma_\epsilon, \xi_t, J, Z, S) \\ &\quad \times p(\xi_t | \Theta, \Sigma_\epsilon, \xi_{t-1}, J, Z) p(\xi_{t+1} | \Theta, \Sigma_\epsilon, \xi_t, J, Z) \end{aligned}$$

For ξ_T the conditional posterior is

$$\begin{aligned} p(\xi_T | \Theta, \Sigma_\epsilon, \xi_{\setminus \xi_T}, J, Z, S) &\propto p(\xi_T | \Theta, \Sigma_\epsilon, \xi_{T-1}, J, Z, S) \\ &\propto p(S_T | \Theta, \Sigma_\epsilon, \xi_T, J, Z, S) p(\xi_T | \Theta, \Sigma_\epsilon, \xi_{T-1}, J, Z) \end{aligned}$$

while for ξ_0 it is

$$\begin{aligned} p(\xi_0 | \Theta, \Sigma_\epsilon, \xi_{\setminus \xi_0}, J, Z, S) &\propto p(\xi_0 | \Theta, \Sigma_\epsilon, \xi_1, J, Z) \\ &\propto p(\xi_1 | \Theta, \Sigma_\epsilon, \xi_0, J, Z) p(\xi_0). \end{aligned}$$

C.3 Implementation Details

In the RW-MH steps of the MCMC sample, the proposal density is chosen to be Gaussian, and the efficiency of the RW-MH algorithm depends crucially on the variance of the proposal normal distribution. If the variance is too low, the Markov chain will accept nearly every draw and converge very slowly while it will reject a too high portion of the draws if the variance is too high. I therefore do an algorithm calibration and adjust the variance in the first half of the burn-in period in the MCMC algorithm. Roberts et al. (1997) recommend acceptance rates close to $\frac{1}{4}$ and therefore the standard deviation during the algorithm calibration is chosen as follows: Every 100'th draw the acceptance ratio of each parameter is evaluated. If it is less than 10 % the standard deviation is doubled while if it is more than 50 % it is cut in half. This step is prior to the second half of the burn-in period since

¹⁷If the draws were frequently rejected the method in Gelfand, Smith, and Lee (1992) could be used.

the convergence results of RW-MH only applies if the variance is constant (otherwise the Markov property of the chain is lost).

In section 2.5 it is stated that for each date $X_i(t)$ is chosen such that the 5-year CDS spread is fitted perfectly. If a price is quoted very low on a given day compared to the average in the sample period (remember the average determines the a_i s) this would restrict the common factor to be very low on that date. In order to mitigate the effect of possible misquotings, I allow five CDS spreads not to be fitted perfectly on a given date. Therefore, if a quoted CDS price is very low and the model-implied CDS price is more than actual CDS price even with $X_i(t) = 0$ this is acceptable for a maximum of five issuers each day and $X_i(t) = 0$ for these issuers. In the actual estimation this occurs very few times.

The Fourier inversion in equation (5) is calculated by using Fast Fourier Transform and the number of points used in FFT is 2^{18} . The characteristic function is not evaluated in every Fourier transform point. Instead, since the characteristic function is exponential-affine with function A and B , the functions A and B are splined from a lower number of points. The spline uses a total number of 40 points. 25 points are spread evenly out in the interval while 15 points are placed near 0. Also, the integration in (4) is done using Gauss-Legendre integration and the number of integration points is 50.

D Standard 1-Factor Gaussian Copula Model

In the Gaussian copula model, for each issuer it is assumed that the value of issuer i 's assets at time t is given as

$$X_i = aZ + \sqrt{1 - a^2}\epsilon_i,$$

where $Z, \epsilon_1, \dots, \epsilon_N$ are independent standard normal random variables. The correlation between any pairs of X_i and X_j is a^2 . Defaults are modelled in an intensity-based setting with constant default intensities through time and across issuers, i.e. the probability of default for issuer i at time t is given as

$$p_i(t) = 1 - e^{-\lambda t}.$$

To determine λ the approximation $CDS \approx (1 - \delta)\lambda$ is applied where δ is the recovery rate and CDS is the average 5-year CDS spread. The recovery rate δ is set to 40%.

There is a default barrier x_i such that issuer i defaults if $X_i < x_i$ and for time t we have that

$$p_i(t) = P(X_i < x_i).$$

Since X_i is standard normal we have that

$$x_i = \Phi^{-1}(p_i(t)).$$

Conditional on Z we have that

$$p_i(t|Z) = P(\epsilon_i < \frac{x_i - aZ}{\sqrt{1 - a^2}}) = \Phi(\frac{\Phi^{-1}(p_i(t)) - aZ}{\sqrt{1 - a^2}}),$$

and since we assume that the default probabilities are the same across issuers we have that the probability of k defaults among N issuers is

$$P_t(X = k|Z) = \binom{N}{k} p(t|Z)^k (1 - p(t|Z))^{N-k}.$$

The unconditional probability of k defaults at time t is

$$P_t(X = k) = \int_{-\infty}^{\infty} P_t(k|x) \phi(x) dx$$

and the cumulative probability is

$$P_t(X \leq m) = \sum_{k=0}^m \int_{-\infty}^{\infty} P_t(k|x) \phi(x) dx$$

which can be approximated as shown by Vasiček (1987) as

$$P_t(X \leq m) = \Phi\left(\frac{\sqrt{1 - a^2} \Phi^{-1}(m) - \Phi^{-1}(p_i(t))}{a}\right).$$

The last formula yields the loss distribution on each time t and the formulas in section 2.4 can be applied to find CDO tranche spreads for a given level of correlation.

References

- Andersen, L., J. Sidenius, and S. Basu (2003). All Your Hedges in one Basket. *Risk*, 67–72.
- Belsham, T., N. Vause, and S. Wells (2005). Credit Correlation: Interpretation and Risks. *Financial Stability Review, Bank of England* (19), 103–115.
- Besag, J. (1974). Spatial Interaction and the Statistical Analysis of Lattice Systems. *Journal of the Royal Statistical Association Series B* 36, 192–236.
- Bhansali, V., R. Gingrich, and F. Longstaff (2008). Systemic credit risk: what is the market telling us? *Working Paper, UCLA Anderson School*.
- Cheridito, P., D. Filipovic, and R. L. Kimmel (2007). Market Price of Risk Specifications for Affine Models: Theory and Evidence. *Journal of Financial Economics* 83, 123–170.
- 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, 73–85.
- Duffie, D. and N. Gârleanu (2001). Risk and Valuation of Collateral Debt Obligations. *Financial Analysts Journal* 57(1), 41–59.
- Duffie, D. and K. Singleton (1997). An econometric model of the term structure of interest rate swap yields. *Journal of Finance* 52, 1287–1321.
- Duffie, D. and K. J. Singleton (1999). Modeling Term Structures of Defaultable Bonds. *Review of Financial Studies* 12, 687–720.
- Eckner, A. (2007a). Computational Techniques for basic Affine Models of Portfolio Credit Risk. *Working Paper, Stanford*.
- Eckner, A. (2007b). Risk Premia in Structured Credit Derivatives. *Working Paper, Stanford*.
- Feldhütter, P. (2006). Can Affine Models Match the Moments in Bond Yields? *Working Paper, Copenhagen Business School*.

- Feldhütter, P. and D. Lando (2007). Decomposing Swap Spreads. *Forthcoming, Journal of Financial Economics*.
- Frühwirth-Schnatter, S. and A. L. Geyer (1998). Bayesian Estimation of Econometric Multifactor Cox Ingersoll Ross Models of the Term Structure of Interest Rates via MCMC Methods. *Working Paper, Department of Statistics, Vienna University of Economics and Business Administration*.
- Gelfand, A. E., A. F. Smith, and T.-M. Lee (1992). Bayesian Analysis of Constrained Parameter and Truncated Data problems Using Gibbs Sampling. *Journal of the American Statistical Association* 87(418), 523–532.
- Geman, S. and D. Geman (1984). Stochastic Relaxation, Gibbs Distributions and the Bayesian Restoration of Images. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 6, 721–741.
- Hammersley, J. and P. Clifford (1970). Markov Fields on Finite Graphs and Lattices. *Unpublished Manuscript*.
- Houdain, J. and D. Guegan (2006). Hedging Tranched Index Products: Illustration of the Model Dependency. *Working Paper*.
- Hull, J. and A. White (2004). Valuation of a CDO and nth to Default CDS Without Monte Carlo Simulation. *Journal of Derivatives* 12(2), 8–23.
- Johannes, M. and N. Polson (2003). MCMC Methods for Financial Econometrics. *forthcoming, Handbook of Financial Econometrics*.
- Jorion, P. and G. Zhang (2006). Good and Bad Credit Contagion: Evidence from Credit Default Swaps. *forthcoming, Journal of Financial Economics*.
- Lando, D. (1994). *Three Essays on Contingent Claims Pricing*. Ph. D. thesis, Cornell University.
- Longstaff, F., S. Mithal, and E. Neis (2005). Corporate yield spreads: default risk or liquidity? New evidence from the credit-default swap market. *Journal of Finance* 60(5), 2213–2253.
- Longstaff, F. and A. Rajan (2006). An Empirical Analysis of the Pricing of Collateralized Debt Obligations. *forthcoming, Journal of Finance*.
- Marsaglia, G. and A. Zaman (1991). A New Class of Random Number Generators. *Annals of Applied Probability* 3, 462–480.

- Mashal, R. and M. Naldi (2002). Pricing Multiname Credit Derivatives: Heavy Tailed Hybrid Approach. *Working Paper, Columbia Business School*.
- Merrill Lynch (2003). Correlation Trading. *Technical Report, Merrill Lynch Global Securities Research and Economics Group, Fixed Income Strategy, Merrill Lynch*.
- Metropolis, N., A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller (1953). Equations of State Calculations by fast Computing Machines. *Journal of Chemical Physics* 21, 1087–1091.
- Micu, M., E. Remolona, and P. Wooldridge (2004). The Pricing Impact of Rating Announcements. *Bis Quarterly Review(June)*, 55–65.
- Mortensen, A. (2006). Semi-Analytic Valuation of Basket Credit Derivatives in Intensity-Based Models. *Journal of Derivatives*.
- Robert, C. P. and G. Casella (2004). *Monte Carlo Statistical Methods* (2 ed.). Springer-Verlag.
- Roberts, G., A. Gelman, and W. Gilks (1997). Weak Convergence and Optimal Scaling of Random Walk Metropolis Algorithms. *Annals of Applied Probability* 7(1), 110–120.
- Tanner, M. and W. Wong (1987). The Calculation of Posterior Distributions by Data Augmentation. *Journal of the American Statistical Association* 82, 528–549.
- Vasiček, O. (1987). Limiting Loan Loss Probability Distribution. *Technical Report, KMV Corporation*.
- Willemann, S. (2005). An Evaluation of the Base Correlation Framework for Synthetic CDOs. *Journal of Credit Risk* 1(4).
- Zhu, H. (2006). An Empirical Comparison of Credit Spreads Between the Bond Market and Credit Default Swap Market. *Journal of Financial Services Research* 29(3).

Figures

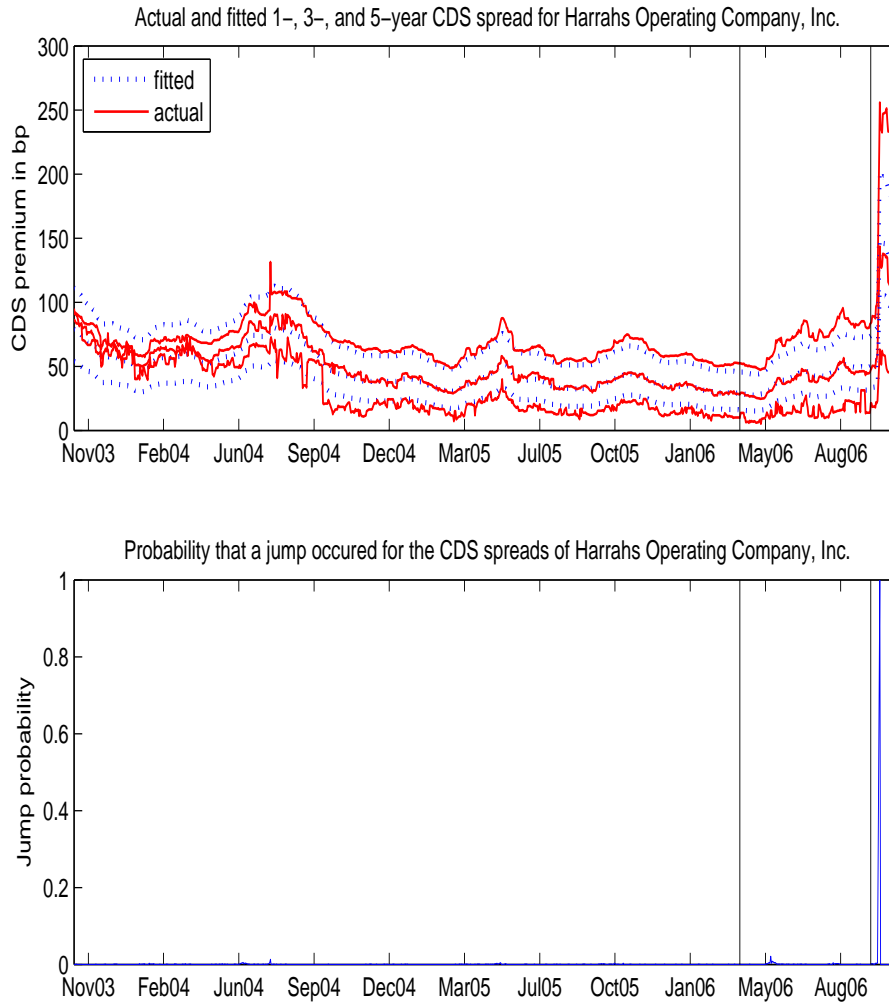


Figure 1: Actual and fitted CDS spreads and jump probabilities for Harrah’s Operating Company, Inc. The top figure shows for Harrah’s Operating Company, Inc the actual and mean fitted 1-, 3-, and 5-year CDS spreads. The 5-year spread is the highest, 3-year spread in the middle, and 1-year spread lowest. The bottom figure shows the jump probabilities. At time t the jump probability is calculated as the mean of the jump indicator across all simulations. In both figures the period for which the CDO model is estimated is marked by two vertical lines.




Figure 2: Actual and fitted CDS spreads and jump probabilities for American International Group, Inc. The top figure shows for American International Group, Inc the actual and mean fitted 1-, 3-, and 5-year CDS spreads. The 5-year spread is the highest, 3-year spread in the middle, and 1-year spread lowest. The bottom figure shows the jump probabilities. At time t the jump probability is calculated as the mean of the jump indicator across all simulations. In both figures the period for which the CDO model is estimated is marked by two vertical lines.

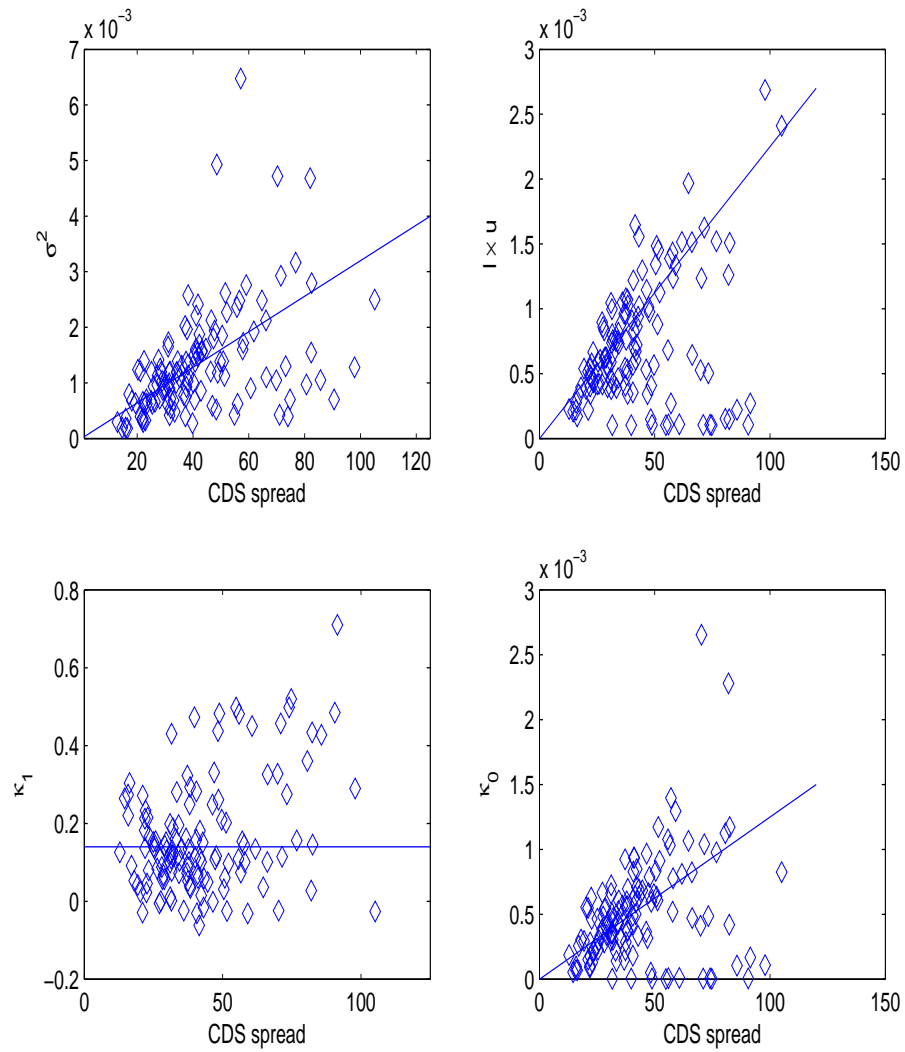


Figure 3: Assumptions of the parsimonious model. Assumption (6) in the text imply that the mean-reversion level κ_1 for all CDS issuers is the same while assumptions (7)-(10) imply that σ^2 , κ_0 , and $l \times \mu$ are linear in a_i . Since a_i is chosen to be the average 5-year CDS spread for issuer i divided by the total average 5-year spread for all issuers, this implies that σ^2 , κ_0 , and $l \times \mu$ are linear in the average CDS spread. The four figures show the relation between σ^2 , κ_0 , $l \times \mu$, and κ_1 and the average 5-year CDS spread across issuers. Each graph has a fitted line illustrating the relevant model assumption.

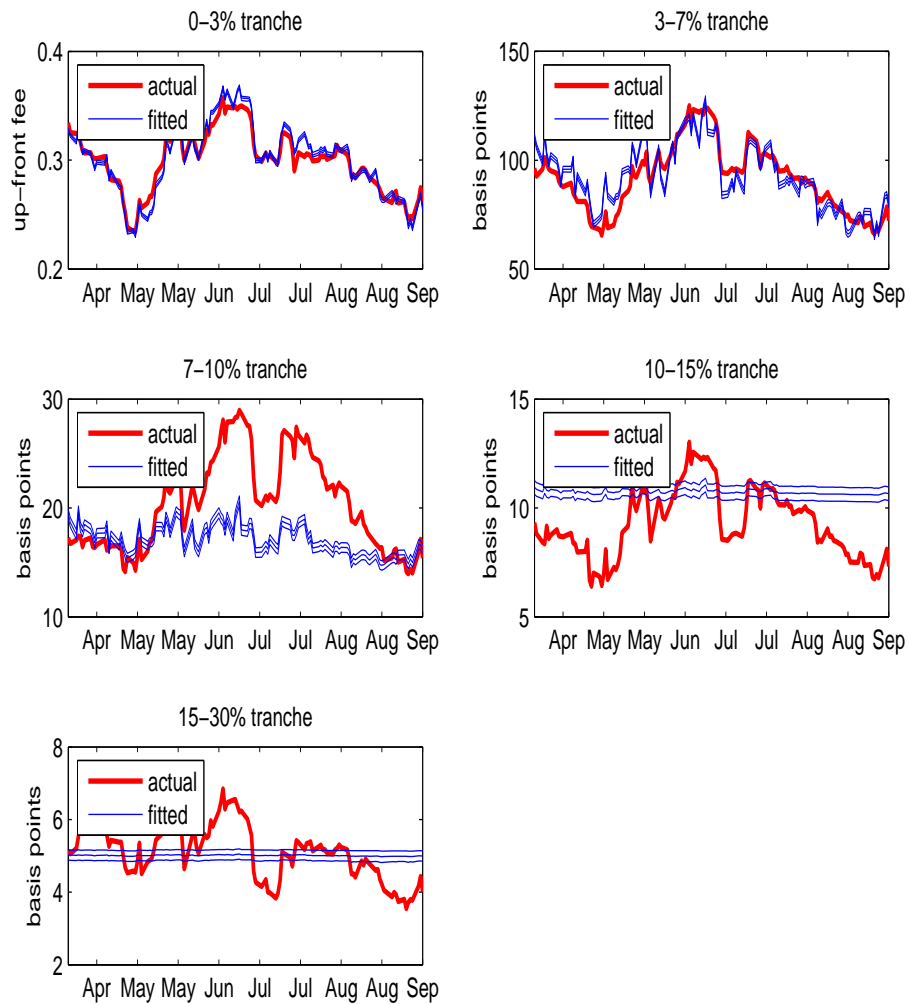


Figure 4: Actual and model-implied tranche spreads. This graph shows for the five CDO tranches used in estimation the actual and model-fitted CDO spreads along with a 95% confidence interval around the model-fitted spreads. The period is March 30, 2006 to September 20, 2006.




Figure 5: Hedging the equity tranche. This graph shows the actual spread changes on the equity tranche (0 – 3%) along with model-predicted spread changes due to underlying CDS spread changes for the affine model along with the standard Gaussian copula.

Tables

CDO spreads, CDX NA IG, March 30, 2006 - September 20, 2006

	0-3%	3-7%	7-10%	10-15%	15-30%
mean	29.92%	91.69bp	20.41bp	9.32bp	5.12bp
standard deviation	2.96%	15.54bp	4.30bp	1.59bp	0.75bp
median	30.29%	92.48bp	20.31bp	9.06bp	5.17bp
min	21.97%	65.52bp	13.96bp	6.40bp	3.54bp
max	35.75%	125.02bp	28.97bp	13.02bp	6.84bp
N	120	120	120	120	120

5-year CDS spreads for constituents, March 30, 2006 - September 20, 2006

	5y
mean	37.67bp
standard deviation	28.62bp
median	27.97bp
minimum	4.87bp
maximum	208.01bp
N	15,000

CDS spreads for constituents, October 27, 2003 - October 26, 2006

	1y	3y	5y
mean	15.58bp	27.80bp	42.19bp
standard deviation	14.39bp	18.80bp	25.39bp
median	11.07bp	22.94bp	36.12bp
minimum	1.34bp	2.93bp	5.35bp
maximum	215.75bp	228.02bp	259.30bp
N	96,631	96,631	96,631

Table 1: *Summary Statistics.* This table shows summary statistics for the data used in estimation. Panel data for the 5-year CDS spreads and all CDO tranche prices for March 30, 2006-September 20, 2006 is used in estimation of an affine correlation model. Panel data for the 1-, 3-, and 5-year CDS spread for each underlying issuer for October 27, 2003-October 26, 2006 is used in estimation of marginal intensity models for each issuer.

	$\kappa_0(\times 10^4)$	κ_1	$l\mu(\times 10^4)$	$\sigma(\times 10^2)$	κ_1^P	$l^P\mu^P(\times 10^4)$	s
Mean	4.94	0.164	7.29	3.56	-1.66	6.98e+024	3.72
Median	4.36	0.13	6.2	3.35	-1.61	14.6	3.07
Min	0.0117	-0.0618	1.01	1.34	-4.6	0.111	0.766
Max	27.7	0.71	27.1	18.1	0.121	8.72e+026	11.1
2.5% quantile	0.0343	-0.0257	1.04	1.72	-3.12	0.5	1.18
97.5% quantile	11.3	0.481	15.6	5.4	-0.331	361	8.07

Table 2: *Estimates of a one-factor affine jump-diffusion model for CDS issuers.* The Duffie and Gârleanu (2001) model for correlated defaults implies that the marginal default intensity for all underlying issuers is given as $d\xi_t = (\kappa_0^i + \kappa_1^i \xi_t)dt + \sigma^i \sqrt{\xi_t} dW_t^{Q,i} + dJ_t^{Q,i}$ under the pricing measure where the jump times arrive with intensity l^i and the jump sizes are exponentially distributed with mean μ^i while the dynamics under the historical measure is $d\xi_t = (\kappa_0^{P,i} + \kappa_1^{P,i} \xi_t)dt + \sigma^i \sqrt{\xi_t} dW_t^P + dJ_t^{P,i}$ with jumps arriving with intensity $l^{P,i}$ and the jump sizes are exponentially distributed with mean $\mu^{P,i}$. This table shows summary statistics for univariate estimations for all 125 issuers underlying the Dow Jones CDX North American Investment Grade Index for the period March 21, 2006, to September 20, 2006 (CDX NA IG series 6). For each CDS estimation all parameters are estimated by taking the median across MCMC draws. This table shows summary statistics for the parameter estimates across the 125 CDS issuers. The univariate panel data estimations are based on daily 1-, 3-, and 5-year CDS premia for the period from October 27, 2003 to October 26, 2006.

$\kappa_0(\times 10^5)$	κ_1	$\sigma(\times 10^2)$	$l(\times 10^3)$	$l\mu(\times 10^3)$
1.59	0.4648	3.668	3.186	3.92
(1.301;1.828)	(0.4539;0.4712)	(3.607;3.76)	(3.023;3.304)	(2.427;5.611)
κ_1^P	$l^P(\times 10^3)$	$l^P u^P(\times 10^4)$	ω	
0.4402	3.377	0.007893	0.9742	
(-4.725;5.545)	(4.305e-012;1.186e+004)	(3.243e-015;44.89)	(0.8565;0.9989)	
$\sqrt{\Sigma_{11}}(\times 10^3)$	$\sqrt{\Sigma_{22}}(\times 10^4)$	$\sqrt{\Sigma_{33}}(\times 10^4)$	$\sqrt{\Sigma_{44}}(\times 10^4)$	$\sqrt{\Sigma_{55}}(\times 10^4)$
7.998	6.806	5.296	2.99	0.9974
(7.89;8.113)	(6.711;6.897)	(5.219;5.371)	(2.95;3.033)	(0.9832;1.011)

Table 3: *Parameter estimates of CDO pricing model.* This table shows the parameter estimates for the multi-name default model outlined in section 2.5. Estimates parameters are median values. The estimation is done on a panel data set of 5 CDO tranche prices and 125 CDS 5-year spreads on 120 days from March 30, 2006, to September 30, 2006. The CDO tranches are the 0 – 3%, 3 – 7%, 7 – 10%, 10 – 15%, and 15 – 30% tranche prices for the CDX North American Investment Grade Index (Series 6).

	Mean error	RMSE
0-3%	0.00142% (-6.52e-006%;0.00273%)	0.00787% (0.00772%;0.0082%)
3-7%	1.11 bp (-0.139 bp;2.44 bp)	7.05 bp (6.76 bp;7.35 bp)
7-10%	-3.28 bp (-3.72 bp;-2.85 bp)	4.91 bp (4.63 bp;5.22 bp)
10-15%	1.42 bp (1.11 bp;1.73 bp)	2.09 bp (1.88 bp;2.31 bp)
15-30%	-0.113 bp (-0.255 bp;0.0314 bp)	0.742 bp (0.731 bp;0.775 bp)

	Relative mean error	Relative RMSE
0-3%	0.36% (-0.125%;0.805%)	2.61% (2.56%;2.68%)
3-7%	1.7% (0.276%;3.19%)	8.23% (7.83%;8.7%)
7-10%	-13.2% (-15.4%;-11%)	20.1% (19.1%;21.4%)
10-15%	18.6% (15.1%;22%)	27.2% (24.4%;30.1%)
15-30%	-0.0364% (-2.88%;2.85%)	15.1% (14.8%;15.9%)

Table 4: *CDO pricing errors.* The first part of this table shows the average pricing error and the root-mean-squared-error for each CDO tranche in the CDO estimation. The average pricing error is the average model-implied minus actual price, while RMSE is the square root of average squared pricing errors. The second part of this table shows the relative pricing error, defined as the pricing error divided by the spread, and the relative RMSE, defined as the RMSE for the relative error. For each simulation in MCMC estimation, mean pricing errors and RMSEs are calculated, thereby yielding time series of pricing errors and RMSEs. For each time series a 95% confidence band is defined as the interval between the 2.5% and 97.5% quantile.

Affine model		
	Mean hedging error	RMSE
0-3%	-8.8e-005%	0.00608%
	(-0.000102%;-7.16e-005%)	(0.00594%;0.00625%)
3-7%	-0.0839 bp	3.2 bp
	(-0.0887 bp;-0.08 bp)	(3.16 bp;3.27 bp)
7-10%	-0.0124 bp	1.12 bp
	(-0.013 bp;-0.012 bp)	(1.11 bp;1.12 bp)
10-15%	-0.0168 bp	0.578 bp
	(-0.0169 bp;-0.0167 bp)	(0.578 bp;0.579 bp)
15-30%	-0.00939 bp	0.299 bp
	(-0.00942 bp;-0.00936 bp)	(0.299 bp;0.299 bp)

Gaussian copula		
	Mean hedging error	RMSE
0-3%	-0.000392%	0.00663%
3-7%	0.0308 bp	3.27 bp
7-10%	0.0217 bp	0.981 bp
10-15%	-0.000139 bp	0.471 bp
15-30%	-0.00177 bp	0.257 bp

Random Walk		
	Mean hedging error	RMSE
0-3%	-0.000609%	0.00855%
3-7%	-0.202 bp	4.12 bp
7-10%	-0.0156 bp	1.18 bp
10-15%	-0.0168 bp	0.583 bp
15-30%	-0.00948 bp	0.3 bp

Table 5: *CDO hedging errors.* This table shows the average hedging error and RMSE for each CDO tranche in the CDO estimation. Denoting P_t^a the actual and $P_t(S_t)$ a model-implied tranche price at time t with CDS spreads at time t as input, then the hedging error at time $t + 1$ is defined as $(P_{t+1}^a - P_t^a) - (P_t(S_{t+1}) - P_t(S_t))$. For each hedging error and RMSE a 95% confidence band is given in parenthesis.