

Systematic and idiosyncratic default risk in synthetic credit markets *

Peter Feldhütter

Mads Stenbo Nielsen[†]

This draft: July 8, 2010

Abstract

We present a new estimation approach that allows us to extract from synthetic credit markets the contribution of systematic and idiosyncratic default risk to total default risk. Using an extensive data set of 90,600 CDS and CDO tranche spreads on the North American Investment Grade CDX index we conduct an empirical analysis of an intensity-based model for correlated defaults. Our results show that systematic default risk is an explosive process with low volatility, while idiosyncratic default risk is more volatile but less explosive. Also, we find that the model is able to capture both the level and time series dynamics of CDO tranche spreads.

*We 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, msn.fi@cbs.dk).

1 Introduction

Campbell and Taksler (2003) show that idiosyncratic firm-level volatility is a major driver of corporate bond yield spreads and that there has been an upward trend over time in idiosyncratic equity volatility in contrast to market-wide volatility. This suggests that in order to understand changing asset prices over time, it is important to separate out and understand the dynamics of both idiosyncratic and systematic volatility. In this paper, we present a new approach to separate out the size and time series behavior of idiosyncratic and systematic (default intensity) volatility by using information in synthetic credit markets.

Markets for credit derivatives have experienced massive growth in recent years (see Duffie (2008)) and numerous models specifying default and correlation dynamics have been proposed. A good model of multi-name default should ideally have the following properties (see Collin-Dufresne (2009) for a discussion). First, the model should be able to match 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 products in a market where prices are available for standard products. Second, the model should have parameters that are economically interpretable, such that parameter values can be discussed and critically evaluated. If a non-standard product needs to be priced and parameters cannot be inferred from existing market prices, economic interpretability provides guidance in choosing parameters. Third, credit spreads and their correlation should be modelled dynamically such that options on multi-name products can be priced. And fourth, since market makers quote spreads at any given time, pricing formulas should not be too time-consuming to evaluate.

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 used.¹ 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. We follow Duffie and Gârleanu (2001) and model the default intensity of each issuer as 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 credit spreads are matched. While the model has many attractive properties it has not been much used because empirically estimating the model has proven to be challenging.

We present a new approach to estimate intensity-based models based on quotes from synthetic credit markets. The main challenge so far has been that the estimation of a model based on an index with 125 names required the simultaneous estimation of a common factor and 125 idiosyncratic factors. This challenge has been dealt with by imposing strong parameter restrictions on the idiosyncratic factors (see among others Mortensen (2006), Eckner (2007a), and Eckner (2007b)). Our approach circumvents this problem by using the insight that total firm-specific default risk can be measured model-independently from CDS spreads and then only systematic default risk needs to be specified while idiosyncratic risk can be left unmodelled. This reduces the problem of estimating 126 factors to estimating one factor. Subsequently, idiosyncratic default risk can be parameterized and idiosyncratic factors can be estimated one at a time. Thus, the approach reduces the problem of estimating 126 factors simultaneously to 126 single-factor estimations and in the estimation restrictions on the factors are not

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

necessary.

We apply our approach to the North American Investment Grade CDX index, and estimate both systematic and idiosyncratic default risk as affine jump-diffusion processes using CDS and CDO spreads. Papers imposing strong parameter restrictions have found that intensity-based jump-diffusion models can match the levels but not the time series behavior of CDO tranche spreads (see Eckner (2007b) for a discussion). We find that the models can in fact match not only the levels but also the time series behavior in tranche spreads. That is, once parameter restrictions are not imposed, the model gains the ability to match time series dynamics of systematic and unsystematic default risk. We also find that idiosyncratic default risk is a major driver of total default risk consistent with the findings in Campbell and Taksler (2003). Furthermore, we confirm the finding in Zhang, Zhou, and Zhu (2009) that both diffusion volatility and jumps are important for default risk. More importantly, our analysis allows us to separate both idiosyncratic and systematic default risk into a diffusion part and jump part yielding new insights: compared to systematic default risk, idiosyncratic default risk has a higher diffusion volatility, a higher contribution from jumps, and is less explosive.

An alternative modelling approach to that of ours is to model aggregate portfolio loss and fit the model to CDO tranche spreads. This is the approach taken in for example Longstaff and Rajan (2008), and Giesecke and Goldberg (2006), Errais, Giesecke, and Goldberg (2007). Since the default intensity of individual firms is not modelled, this approach is not useful for examining individual default risk, whether it is systematic or unsystematic.

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 ability of the model to match CDO tranche spreads and examines the properties of systematic default risk, while idiosyncratic default risks are examined in section 6 . Section 7 concludes.

2 Intensity-Based Default Risk Model

This section explains the model framework that we employ for pricing single- and multi-name credit securities. For single-name Credit Default Swaps (CDSs) we use the intensity-based framework introduced in Lando (1994) and Duffie and Singleton (1999). For multi-name Collateralized Debt Obligation (CDO) valuation we follow Duffie and Gârleanu (2001) and model the default intensity of each underlying issuer as the sum of an idiosyncratic and a common process. Default correlation among issuers thus arises through the joint dependence of individual default intensities on the common factor. Furthermore, we generalize the model in Duffie and Gârleanu (2001) by allowing for a flexible specification of the idiosyncratic processes, while maintaining semi-analytical calculation of the loss distribution as in Mortensen (2006). This extension allows us to avoid the ad hoc parameter restrictions that are common in the existing literature.

2.1 Default Modeling

We assume that the time of default of a single issuer, τ , is modelled through an intensity $(\lambda_t)_{t \geq 0}$, which implies that the risk-neutral probability at time t of defaulting within a short period of time Δt is approximately

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

Unconditional default probabilities are given by

$$Q_t(\tau \leq s) = 1 - E_t^Q \left[\exp \left(- \int_t^s \lambda_u du \right) \right] \quad (1)$$

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

In our model we consider a total of N different issuers. To model correlation between individual issuers we follow Mortensen (2006) and assume that the intensity of each issuer is given as the sum of an idiosyncratic component and a scaled common component

$$\lambda_{i,t} = a_i Y_t + X_{i,t} \quad (2)$$

where a_1, \dots, a_N are non-negative constants and Y, X_1, X_2, \dots, X_N are independent stochastic processes. The common factor Y creates dependence in default occurrences among the N issuers and may be viewed as reflecting the overall state of the economy, while X_i similarly represents the idiosyncratic default risk for firm i . Thus, a_i indicates the sensitivity of firm i to the performance of the macro economy, and we allow this parameter to vary across firms, contrary to Duffie and Gârleanu (2001) that assume $a_i = 1$ for all i and thereby enforce a homogeneous impact of the macro economy on all issuers.

We assume that the common factor follows an affine jump diffusion under the risk-neutral measure

$$dY_t = (\kappa_0 + \kappa_1 Y_t)dt + \sigma \sqrt{Y_t} dW_t^Q + dJ_t^Q \quad (3)$$

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 $\mu > 0$. This process is well-defined for $\kappa_0 > 0$. As a special case, if the jump intensity is equal to zero the default intensity then follows a CIR process.

We do *not* impose any distributional assumptions on the evolution of the idiosyncratic factors X_1, \dots, X_N . In particular, they are not required to be affine jump diffusions. This generalizes the setup in Duffie and Gârleanu (2001), Mortensen (2006), and Eckner

(2009), where the idiosyncratic factors are required to be affine jump diffusions with very restrictive assumptions on their parameters.

2.2 Risk Premium

For the basic affine process in equation (3) we 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 and since this restriction is likely to be violated as discussed in Feldhütter (2006) we choose the more parsimonious essentially affine risk premium.²

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

$$dY_t = (\kappa_0 + \kappa_1^P Y_t)dt + \sigma\sqrt{Y_t}dW_t^P + dJ_t^P \quad (4)$$

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

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

2.3 Aggregate Default Distribution

Our model allows for semi-analytic calculation of the distribution of the aggregate number of defaults among the N issuers. More specifically, we can at time t calculate in semi-closed form the distribution of the aggregate number of defaults at time $s \geq t$ by conditioning on the common factor. If we let

$$Z_{t,s} = \int_t^s Y_u du$$

denote the integrated common factor, then it follows from (1) and (2) that conditional on $Z_{t,s}$, defaults are independent and given as

$$p_{i,t}(s|z) = Q_t(\tau_i \leq s | Z_{t,s} = z) = 1 - \exp(-a_i z) E_t^Q \left[\exp \left(- \int_t^s X_{i,u} du \right) \right]. \quad (5)$$

The total number of defaults at time s among the N issuers, D_s^N , is then found by the recursive algorithm³

$$Q_t(D_s^N = j|z) = Q_t(D_s^{N-1} = j|z)(1 - p_{N,t}(s|z)) + Q_t(D_s^{N-1} = j-1|z)p_{N,t}(s|z)$$

due to Andersen, Sidenius, and Basu (2003). The unconditional default distribution is therefore given as

$$Q_t(D_s^N = j) = \int_0^\infty Q_t(D_s^N = j|z) f_{t,s}(z) dz \quad (6)$$

where $f_{t,s}$ is the density function for $Z_{t,s}$. Finally, $f_{t,s}$ can be determined by Fourier inversion of the characteristic function $\phi_{Z_{t,s}}$ for $Z_{t,s}$ as

$$f_{t,s}(z) = \frac{1}{2\pi} \int_{-\infty}^\infty \exp(-iuz) \phi_{Z_{t,s}}(u) du \quad (7)$$

where we apply the closed-form expression for $\phi_{Z_{t,s}}$ derived in Duffie and Gârleanu (2001).⁴

³The last term disappears if $j = 0$.

⁴Duffie and Gârleanu (2001) derive an explicit solution for $E_t^Q[\exp(q \int_t^s Y_u du)]$ when q is a real number, but as noted by Eckner (2009) the formula works equally well for q complex.

2.4 Synthetic CDO Pricing

CDOs began to trade frequently in the mid-nineties and in the last decade issuance of CDOs has experienced massive growth, see BIS (2007). In a CDO the credit risk of a portfolio of debt securities is passed on to investors by issuing CDO 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 credit 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 agent *A* receives quarterly a CDS premium until the CDS contract expires or the bond defaults. In case of default, agent *A* receives the defaulted bond in exchange for face value. The loss is therefore the difference between face value and market value of the bond.⁵

Agent *A* at the same time issues a CDO tranche on the first 3% of losses in his CDS portfolio and agent *B* "buys" this tranche, which has a principal of \$3 million. No money is exchanged at time 0, when the tranche is sold. If the premium on the tranche is, say, 2,000 basis points, agent *A* pays a quarterly premium of 500 basis points to agent *B* on the remaining principal. If a default occurs on any of the underlying CDS contracts, the loss is covered by agent *B* and his principal is reduced accordingly. Agent *B* continues to receive the premium on the remaining principal until either the CDO contract matures or the remaining principal is exhausted. Since the first 3% of portfolio losses are covered by this tranche it is called the 0% – 3% tranche. Similarly, agent *A*

⁵Pricing CDS contracts is explained in Appendix 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, we find the fair spread at time t on a specific CDO tranche. Consider a tranche that covers portfolio losses between K_1 and K_2 from time $t_0 = t$ to $t_M = T$, and assume that the tranche has quarterly payments at time t_1, \dots, t_M . The tranche premium is found by equating the value of the protection and premium payments. We denote the total portfolio loss in percent at time s as L_s , i.e. the percentage number of defaults D_s^N / N times $1 - \delta$, where δ is the recovery rate, which we assume to be constant at 40%. The tranche loss is then given as

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

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

$$Prot(t, T) = E_t^Q \left[\int_t^T \exp\left(-\int_t^s r_u du\right) dT_{K_1, K_2}(L_s) \right]$$

while the value of the premium payments is the annual tranche premium $S(t, T)$ times

$$\begin{aligned} & Prem(t, T) \\ &= E_t^Q \left[\sum_{j=1}^M \exp\left(-\int_t^{t_j} r_u du\right) (t_j - t_{j-1}) \int_{t_{j-1}}^{t_j} \frac{K_2 - K_1 - T_{K_1, K_2}(L_s)}{t_j - t_{j-1}} ds \right] \end{aligned}$$

where r_u is the riskfree interest 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 . The CDO tranche premium at time t is thus given as

$$S(t, T) = \frac{Prot(t, T)}{Prem(t, T)}.$$

We follow Mortensen (2006) and discretize the integrals appearing in $Prot(t, T)$ and $Prem(t, T)$ at premium payment dates, we assume that the riskfree rate is uncorrelated

with portfolio losses, and that defaults occur halfway between premium payments. Under these assumptions the value of the protection payment is

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

while the expression for the premium payments reduces to

$$\begin{aligned} & Prem(t, T) \\ &= \sum_{j=1}^M (t_j - t_{j-1}) P(t, t_j) \left(K_2 - K_1 - \frac{E_t^Q[T_{K_1, K_2}(L_{t_{j-1}})] + E_t^Q[T_{K_1, K_2}(L_{t_j})]}{2} \right) \end{aligned}$$

where $P(t, s) = E_t^Q[\exp(-\int_t^s r_u du)]$ is the price at time t of a riskless zero coupon bond maturing at time s .

3 Estimation

The parameters in our intensity model are estimated in three separate steps. First we imply out firm-specific term structures of risk-neutral survival probabilities from daily observations of CDS spreads, second, we use the inferred survival probabilities to estimate each issuer's sensitivity a_i to the economy-wide common factor Y , and finally we estimate the parameters and the path of the common factor using a Bayesian MCMC approach.⁶ An important ingredient in the third step is our explicit use of the calibrated survival probabilities, which implies that we do not need to impose any structure on the idiosyncratic factors X_i .

In other words, we can estimate the model without putting specific structure on the idiosyncratic factors and this has several advantages, which we discuss in section 3.1.

⁶For a general introduction to MCMC see Robert and Casella (2004) and for a survey of MCMC methods in financial econometrics see Johannes and Polson (2006).

Note also that our estimation approach is consistent with the common view that CDS contracts may be used to read off market views of marginal default probabilities, whereas basket credit derivatives instead reflect the correlation patterns among the underlying entities, see e.g. Mortensen (2006).

In an additional fourth step of the estimation procedure, we take in section 6 a closer look at the cross-section of the idiosyncratic factors implicitly given by the inferred survival probabilities and the estimated common factor. Here, we impose a dynamic structure on each X_i and then estimate the parameters for each idiosyncratic factor separately, again using MCMC methods.

3.1 A General Estimation Approach

For each day in our data sample we observe 5 CDO tranche spreads as well as CDS spreads for a range of maturities for each of the 125 firms underlying the CDO tranches. Previous literature on CDO pricing has also studied models of the form (2), but only by imposing strong assumptions on the parameters of the idiosyncratic factors X_i , as well as by disregarding the information in the term structure of CDS spreads (Mortensen (2006), Eckner (2007b), Eckner (2009)). In this paper, we remove both of these shortcomings by allowing the idiosyncratic factors to be of a very general form, while we at the same time use all the available information from each issuer's term structure of CDS spreads.

Theoretically, if we had CDS contracts for any maturity we could extract survival probabilities for any future time-horizon, but in practice CDS contracts are only traded for a limited range of maturities. To circumvent this problem we assume a flexible parametric form for the term structure of risk-neutral survival probabilities, and use

that to infer survival probabilities from the observed CDS spreads.⁷ That is, on any given day t and for any given firm i we extract from the observed term structure of CDS spreads the term structure of marginal survival probabilities $s \mapsto q_{i,t}(s)$, where

$$q_{i,t}(s) = Q_t(\tau_i > s) = E_t^Q \left[\exp \left(- \int_t^s (a_i Y_u + X_{i,u}) du \right) \right],$$

see appendix B for details. Once we condition on the value of the common factor, this directly gives us the idiosyncratic component

$$E_t^Q \left[\exp \left(- \int_t^s X_{i,u} du \right) \right]$$

of the risk-neutral survival probability. Thus, we can use observed CDS spreads to derive values of the function $s \mapsto E_t^Q[\exp(-\int_t^s X_{i,u} du)]$, which is all we need to calculate the aggregate default distribution (and hence compute CDO tranche spreads) using equation (5). Therefore, we do not need to explicitly model the stochastic behavior of each X_i in order to price CDO tranches.

For each firm i in our sample, the parameter a_i measures that firm's sensitivity to the overall state of the economy, and this parameter can be estimated directly from the inferred term structures of survival probabilities $s \mapsto q_{i,t}(s)$. Intuitively, a_i measures to what extent the default probability of firm i is correlated with the average default probability (since this average mainly reflects exposure to the systematic risk factor Y), and therefore a consistent estimate of a_i is given by the slope coefficient in the regression of firm i 's short-term default probability on the average short-term default probability of all 125 issuers.⁸ Appendix C provides the technical details.

⁷This procedure is essentially similar to the well-known technique for inferring a term structure of interest rates from observed prices of coupon bonds, see Nelson and Siegel (1987).

⁸The average of the a_i 's are without loss of generality normalized to 1.

Once we have inferred marginal default probabilities from CDS spreads and estimated common factor loadings a_i , we can then, given the parameters and current value of the common factor Y , price CDO tranches.

3.2 MCMC Methodology

In order to write the CDO pricing model on state space form, the continuous-time specification in equation (4) is approximated using an Euler scheme

$$Y_{t+1} - Y_t = (\kappa_0 + \kappa_1^P Y_t) \Delta_t + \sigma \sqrt{\Delta_t} Y_t \epsilon_{t+1}^Y + J_{t+1} Z_{t+1} \quad (8)$$

where Δ_t is the time between two observations and

$$\epsilon_{t+1}^Y \sim N(0, 1)$$

$$Z_{t+1} \sim \exp(\mu^P)$$

$$P(J_{t+1} = 1) = l^P \Delta_t.$$

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

On each day $t = 1, \dots, T$, 5 CDO tranche spreads are recorded and stacked in the 5×1 vector S_t , and we let S denote the $5 \times T$ matrix with S_t in the t 'th column. The logarithm of the observed CDO spreads are assumed to be observed with measurement error, so the observation equation is

$$\log(S_t) = \log(f(\Theta, Y_t)) + \epsilon_t, \quad \epsilon_t \sim N(0, \Sigma_\epsilon) \quad (9)$$

where f is the CDO pricing formula. Appendix D gives details on how to calculate f in the estimation of the common factor Y . For the estimation of each of the idiosyncratic

factors X_i in section 6, Y is replaced by X_i and f is instead the model-implied idiosyncratic part of the survival probability, i.e. $E_t^Q[\exp(-\int_t^{t+s} X_{i,u} du)]$, calculated for each of the time horizons $s = 0.5, 1, 2, 3, 4$, and 5 years.

The interest lies in samples from the target distribution $p(\Phi, \Sigma_\epsilon, Y, J, Z|S)$. 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)). We 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)⁹

⁹All 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.

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

Details of the derivations of the conditional and proposal distributions in the Metropolis-Hastings steps are given in Appendix E. 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

In our estimation we use daily CDS and CDO 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 a minimum of three contributors provide data.

We focus in this paper on CDS and CDO prices (i.e. spreads) for defaultable entities in the Dow Jones CDX North America Investment Grade (NA IG) index. The index

contains 125 North American investment grade entities and is updated semi-annually. For our sample period March 21, 2006 to September 20, 2006, the latest version of the index is CDX NA IG Series 6. We specifically select the most liquid CDO tranches, the 5-year tranches, with CDX NA IG 6 as the underlying pool of reference CDSs. These tranches mature on June 20, 2011. Daily spreads of the five CDO tranches we consider: 0% – 3%, 3% – 7%, 7% – 10%, 10% – 15%, and 15% – 30%, are not available for the first 7 days of the period, so the data we use in the estimation covers the period from March 30, 2006 to September 20, 2006. There are holidays on April 14, April 21, June 3, July 4, and September 4, thus leaving a total of 120 days with spreads available.

The quoting convention for the equity tranche (i.e. the 0% – 3% tranche) differs from that of the other tranches. Instead of quoting a running premium, the equity tranche is quoted in terms of an upfront fee. Specifically, an upfront fee of 30% means that the investor receives 30% of the tranche notional at time 0 plus a fixed running premium of 500 basis points per year, paid quarterly.¹⁰

In addition to the CDO tranche spreads, we also use 0.5-, 1-, 2-, 3-, 4-, and 5-year CDS spreads for each of the 125 index constituents.¹¹ The total number of observations

¹⁰Upfront payments may be converted to running spreads using so-called “risky duration”, see e.g. Amato and Gyntelberg (2005). This calculation requires a fully parametric model, and hence is not possible within our modeling framework. Instead we use the original upfront payment quotes available from MarkIt for the equity tranche.

¹¹The 5-year CDS contracts for the period March 21, 2006 to June 19, 2006 mature on June 20, 2011, consistent with the maturity of the 5-year CDO tranches, but for the period June 20 to September 19, 2006, the maturity of the 5-year CDS contracts is September 20, 2011 (and the maturity of the other CDS contracts are similarly shifted forward by 3 months from June 20 and onwards). However, this maturity mismatch between the CDS and CDO contracts in the latter part of our sample period is automatically corrected for, when we imply out the term structures of firm-specific survival probabilities

in the estimation of the multi-name default model is therefore 90,600: 125×6 CDS spreads and 5 CDO tranche spreads observed on 120 days. Table 1 shows summary statistics of the CDS and CDO data.

[Table 1 about here.]

As a proxy for riskless rates we use LIBOR and swap rates since Feldhütter and Lando (2008) show that swap rates are a more accurate proxy for riskless rates than Treasury yields. Thus, prices of riskless zero coupon bonds with maturities up to 1 year are calculated from 1-12 month LIBOR rates (taking into account money market quoting conventions), and for longer maturities are bootstrapped from 1-, 2-, 3-, 4-, and 5-year swap rates (using cubic spline to infer swap rates for semi-annual maturities). This gives a total of 20 zero coupon bond prices on any given day (maturities of 1-12 months, 1.5, 2, 2.5, . . . , 5 years) from which zero coupon bond prices at any maturity up to 5 years can be found by interpolation (again using cubic spline).

5 Results

5.1 Marginal Default Probabilities

As the first step in the estimation of the multi-name default model, we calibrate for each firm daily term structures of risk-neutral default probabilities using all the available information from CDS contracts with maturities up to 5 years. With 125 firms and a sample period of 120 days, we calibrate a total of $125 \times 120 = 15,000$ term structures of default probabilities, with each term structure based on 6 CDS contracts.

from observed CDS spreads (see appendix B), and hence poses no problem to the estimation of the model.

Figure 1 plots for each day in the sample the average term structure of default probabilities across the 125 firms. By definition, the term structures are upward sloping since the probability of defaulting increases as maturity increases. Also, the graph shows that on average the first derivative with respect to maturity is increasing.¹² Thus, forward default probabilities $\frac{\partial Q_t(\tau \leq s)}{\partial s}$, which measure the probability of defaulting at time s given that the firm has not yet defaulted, are upward-sloping. Hence, the market expects the marginal probability of default to increase over time for the average firm. This is likely caused by the fact that the CDX NA IG index consists of solid investment grade firms with low short-term default probabilities, and it is therefore more probable that credit conditions worsen for a given firm than improve.

[Figure 1 about here.]

The sensitivity of each firm's default probability to the economy-wide factor Y is captured in the parameter a_i , which is estimated model-independently through the covariance between firm-specific instantaneous default probabilities and market-wide instantaneous default probabilities. Figure 2 shows the distribution of a_i 's across firms (remember that the a_i 's are normalized such that the average across firms is 1). There is a significant amount of variation in the a_i 's, and for a large fraction of the firms the default probabilities are quite insensitive to market-wide fluctuations in credit risk. This suggests that the assumption in Duffie and Gârleanu (2001) to let all firms have the same sensitivity through identical a_i 's is not supported by the data.

[Figure 2 about here.]

¹²This observation is apparent from a visual inspection of the graph, and quantitative estimates are available upon request.

To examine whether the subset of firms with large a_i 's have common characteristics, we split the index into its five subindices (fraction of total index in parenthesis): Energy (11%), Financials (19%), Basic Industrials (23%), Telecommunications, Media and Technology (18%), Consumer Products and Retail (29%), and we find that firms with large a_i 's are fairly evenly distributed across these five sectors.¹³

The correlation between a_i 's and the average 5-year CDS spread for firm i (averaging across the 120 days) is 0.78 across the 125 firms. This strong positive correlation indicates that the ad hoc assumption in Mortensen (2006), Eckner (2007b) and Eckner (2009), where a_i is exogenously set based on the firm-specific 5-year CDS spread, is reasonable.¹⁴

5.2 CDO Parameter Estimates and Pricing Results

The multi-name default model is estimated on the basis of a panel data set of daily CDS and CDO tranche spreads as described in section 3, and we assume that the measurement error matrix Σ_ϵ in (9) is diagonal and use diffuse priors. We run the MCMC estimation routine using a burn-in period of 20,000 simulations and a subsequent estimation period of another 10,000 simulations, where we use every 10th simulation to calculate parameter estimates.

The parameter estimates are given in Table 2, and the first thing we note is that the volatility of the common factor is $\sigma = 0.0166$, which is low compared to estimates in

¹³The distribution on sectors of the firms with the 20% largest market sensitivities a_i is: Energy (8%), Financials (8%), Basic Industrials (16%), Telecommunications, Media and Technology (28%), Consumer Products and Retail (40%).

¹⁴Mortensen (2006) fixes a_i implicitly through a parameter restriction but notes that it effectively corresponds to setting a_i equal to the fraction of firm-specific to average (across all firms) 5-year CDS spread.

the previous literature: Duffee (1999) fits CIR processes to firm default intensities using corporate bond data and finds an average σ of 0.074 and Eckner (2009) uses a panel data set of CDS and CDO spreads similar to the data set used here and estimates σ to be 0.103. An important factor in explaining this difference in the estimated size of σ is the extent to which systematic and idiosyncratic default risk is separated. Duffee (1999) is not concerned with such a subdivision of the default risk and therefore estimates a factor that includes both systematic and unsystematic risk. Eckner (2009) has a model that is similar to ours, but when estimating the model he imposes strong restrictions on the parameters of the systematic and idiosyncratic factors. For example, he requires σ^2 of the common factor to be equal to the average σ_i^2 of the idiosyncratic factors.

Our results suggest that separating default risk into an idiosyncratic and a common component, and letting these factors be fully flexible during the estimation, reveals that the common factor is "slow-moving" in the sense that the volatility is low. In addition, we estimate the total contribution of jumps $l \times \mu$ to be $6 \cdot 10^{-5}$ which is lower than the estimate of $3 \cdot 10^{-3}$ in Eckner (2009), further underlining that the total volatility of the common factor is low when properly estimated.¹⁵ Finally, we note that although the common factor is not very volatile, it is explosive with a mean reversion coefficient of 0.94 under the risk-neutral measure. Under the actual measure, the factor is estimated to be mean-reverting, although the mean-reversion coefficient is hard to pin down with any precision due to the relatively short time span of our data sample.

[Table 2 about here.]

We now examine the pricing ability of our model by considering the average price

¹⁵In a previous version of this paper, we imposed parameter restrictions similar to Eckner (2009), which resulted in parameter estimates consistent with those that he reports.

ing errors and RMSEs (Root-Mean-Squared-Errors) given in Table 3. We see that on average the model underestimates spreads for the 3% – 7% tranche by 7 basis points and overestimates the 10% – 15% tranche by 4 basis points. For comparison, Mortensen (2006) reports average bid-ask spreads for the 3% – 7% tranche to be 10.9 basis points and for the 10% – 15% to be 5 basis points. In both cases, average pricing errors are smaller than the bid-ask spread. The RMSEs of the model are larger than the average pricing errors, so the model errors are not consistently within the bid-ask spread, but RMSEs and pricing errors do suggest a good overall fit.

[Table 3 about here.]

Figure 3 shows the observed and fitted CDO tranche spreads over time, and the graphs confirm a reasonable fit to all tranches apart from a slight underestimation of the 3% – 7% tranche and overestimation of the 15% – 30% tranche. It is particularly noteworthy that the time series variation in the most senior tranches – especially the 15% – 30% tranche – is well matched. This is surprising because both the level and the time series variation of the 15% – 30% tranche have been difficult to capture by models in the previous literature. Mortensen (2006) finds that jumps in the common factor are necessary to generate sufficiently high senior tranche spreads, but even with jumps it has been difficult to reproduce the observed time series variation in senior tranche spreads, as argued by Eckner (2009) and in a previous version of this paper.¹⁶ What enables our model to fit the time series variation of senior tranche spreads well is that we not imposed the usual set of strong assumptions on the parameters of the common and idiosyncratic factors as done in Mortensen (2006), Eckner (2009), and in a previous

¹⁶The previous version of the paper entitled "An empirical investigation of an intensity-based model for pricing CDO tranches" is available upon request.

version of this paper. Thus, a careful implementation of the multi-name default model frees up the model's ability to fit tranche spreads in important dimensions.

[Figure 3 about here.]

To examine the contribution of systematic default risk to the total default risk across different maturities, we calculate the following: for each maturity, date, and firm we use the estimated sensitivities a_i and the path and parameters of the common factor Y to calculate the systematic part of the risk-neutral default probability according to equation (1) and (2). We then find an average term-structure of systematic default risk by averaging across firms and dates and plot the result in Figure 4 together with the average total default risk inferred from observed CDS spreads. The figure shows that the systematic contribution to the overall default risk is small for short maturities but increases with maturity. As shown in Table 4 the average exposure to systematic default risk on a 6-month horizon is merely 0.003% and constitutes only 6% of the overall default risk, but increases to 0.874% and a fraction of 26% of the total default risk for a 5-year horizon. Hence, out of the total average 5-year default probability of 3.309%, 0.874% is systematic and non-diversifiable.

[Figure 4 about here.]

[Table 4 about here.]

6 Idiosyncratic Default Risk

So far in the estimation we have put structure on the systematic part of default risk through the specification of the common factor, while total default risk has been estimated model-independently. Thus, for each firm and each date, this gives us a term

structure of idiosyncratic default risk which essentially is the difference between total default risk and systematic default risk.¹⁷ Given this panel data set of default risk term structures for each firm, we can now put structure on the idiosyncratic default risk and estimate the parameters of this structural form. Thus, the data set we observe for each firm are idiosyncratic survival probabilities $t \rightarrow E_t^Q[\exp(-\int_t^{t+s} X_{i,u} du)]$ for maturities $s = 0.5, 1, 2, 3, 4, 5$ years on the 120 days in the sample.

We can allow idiosyncratic default risk to be the sum of several factors and the factors can be of any distributional form subject only to the requirements of non-negativity and that we can calculate the expectation

$$E_t^Q \left[\exp \left(- \int_t^s X_{i,u} du \right) \right].$$

We choose to let the idiosyncratic factors have the same functional form as the common factor, namely be a one-factor affine jump-diffusion

$$dX_{i,t} = (\kappa_{i,0} + \kappa_{i,1} X_{i,t}) dt + \sigma_i \sqrt{X_{i,t}} dW_{i,t}^Q + dJ_{i,t}^Q$$

with an essentially affine risk premium for diffusive risk and constant risk premium for the jump risk. This allows us to compare the results of our general estimation approach with those in previous literature, where a number of restrictions are placed jointly on the common and idiosyncratic factors. Thus, for each of the 125 firms in the sample, we estimate by MCMC the parameters of the idiosyncratic factor in the same way as the parameters of the common factor, but in this estimation we observe a panel data set of

¹⁷More specifically, the relation

$$Q_t(\tau_i > s) = E_t^Q \left[\left(-a_i \int_t^s Y_u du \right) \right] \cdot E_t^Q \left[\left(- \int_t^s X_{i,u} du \right) \right]$$

allows us to infer the idiosyncratic survival probabilities directly from the estimated common and total survival probabilities $E_t^Q \left[\left(-a_i \int_t^s Y_u du \right) \right]$, respectively $Q_t(\tau_i > s)$.

idiosyncratic default probabilities. Note that structural assumptions on the idiosyncratic risk were not necessary in order to price CDO tranches in the previous section, but adding structure here enables us to gain further understanding of the nature of the idiosyncratic default risk.

The results from the estimation of the idiosyncratic default factors are given in Table 5. We see that the average volatility across all firms is $\sigma = 0.14$, almost 10 times higher than the volatility estimate of 0.017 for the common factor. Combined with the parameter estimates discussed in the previous section, this shows that the idiosyncratic factors are more volatile than the systematic factor. The fact that the volatility of our systematic factor is lower than that reported in previous papers reflects that our estimation procedure allows us to fully separate the dynamics of the systematic factor from the dynamics of the idiosyncratic factors. This leads to a low-volatility systematic factor and high-volatility idiosyncratic factors, while previous research finds something in-between. In addition, we see that the average total (risk-neutral) contribution from jumps is $l \times \mu = 4 \cdot 10^{-2}$, which is higher than the total jump contribution in the systematic factor of $6 \cdot 10^{-5}$, reinforcing the conclusion that volatilities of the idiosyncratic factors are higher than that of the systematic factor.

[Table 5 about here.]

We see that κ_1 is positive on average, so the idiosyncratic factors are on average explosive under the risk-neutral measure. However, they are less explosive than the systematic factor, implying that when pricing securities sensitive to default risk, the relative importance of systematic risk increases as maturity increases in accordance with our observations in Figure 4.

7 Conclusion

We present a new estimation procedure that allows us to estimate the contribution and dynamics of systematic default risk in an intensity-based model while allowing idiosyncratic default risk to be estimated using a model-free approach. Based on a large data set of CDS and CDO tranche spreads on the North American Investment Grade CDX index, we find that our estimated model is able to capture both the level and time series dynamics of CDO tranche spreads. We then go on and split the total default risk of a given entity into its idiosyncratic and systematic part. We find that the systematic default risk is explosive but has low volatility and that the relative contribution of systematic default risk is small for short maturities, but of growing importance as maturity increases. Our subsequent parametric estimation of the idiosyncratic default risks shows that idiosyncratic risk is more volatile and less explosive than systematic risk.

A CDS Pricing

This section briefly explains how to price credit default swaps (CDSs). More thorough introductions are given in Duffie (1999) and O’Kane (2008).

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

For a CDS contract covering default risk between time $t_0 = t$ and $t_M = T$ and with premium payment dates t_1, \dots, t_M , the value of the protection payment is given as

$$Prot(t, T) = E_t^Q \left[(1 - \delta) \exp \left(- \int_t^\tau r_u du \right) 1_{(\tau \leq T)} \right]$$

where δ is the recovery rate, while the value of the premium payment stream is $S \cdot Prem(t, T)$, where S is the annual CDS premium and

$$Prem(t, T) = E_t^Q \left[\sum_{j=1}^M \exp \left(- \int_t^{\min\{t_j, \tau\}} r_u du \right) \int_{t_{j-1}}^{t_j} 1_{(\tau > s)} ds \right].$$

The CDS premium at time t is settled such that it equates the two payment streams, i.e. $S(t, T) = \frac{Prot(t, T)}{Prem(t, T)}$.

In order to calculate the CDS premium $S(t, T)$ we make the simplifying assumptions that the recovery rate δ is constant at 40%, that the riskfree interest rate is independent of the default time τ , and finally that default, if it occurs, will occur halfway between two premium payment dates. With these assumptions we can rewrite the two expressions

above as

$$Prot(t, T) = (1 - \delta) \sum_{j=1}^M P\left(t, \frac{t_{j-1} + t_j}{2}\right) \cdot (Q_t(\tau > t_{j-1}) - Q_t(\tau > t_j)) \quad (10)$$

$$\begin{aligned} Prem(t, T) &= \sum_{j=1}^M P\left(t, \frac{t_{j-1} + t_j}{2}\right) \cdot \frac{t_j - t_{j-1}}{2} \cdot (Q_t(\tau > t_{j-1}) - Q_t(\tau > t_j)) \\ &+ \sum_{j=1}^M P(t, t_j) \cdot (t_j - t_{j-1}) \cdot Q_t(\tau > t_j). \end{aligned} \quad (11)$$

B Calibration of Survival Probabilities

For the calibration of firm-specific survival probabilities from observed CDS spreads we assume that risk-neutral probabilities take the flexible form

$$Q_t(\tau > s) = \frac{1}{1 + \alpha_2 + \alpha_4} \left(e^{-\alpha_1(s-t)} + \alpha_2 e^{-\alpha_3(s-t)^2} + \alpha_4 e^{-\alpha_5(s-t)^3} \right) \quad s \geq t \quad (12)$$

with all $\alpha_j \geq 0$. The calibrated survival probabilities $s \mapsto Q_t(\tau > s)$ for a given firm at time t are then calculated by minimizing relative pricing errors using (10)–(12)

$$\sum_T \left(\frac{Prot(t, T)/Prem(t, T) - S_{obs}(t, T)}{S_{obs}(t, T)} \right)^2$$

where $S_{obs}(t, T)$ is the empirically observed CDS spread at time t on a contract with maturity T . The calibration is based on observed CDS spreads for maturities of $T = 0.5, 1, 2, 3, 4, 5$ years and is carried out separately for each firm, at each time t , and results in a very accurate fit to the observed CDS term structure.¹⁸

¹⁸This calibration approach is close to the industry benchmark of fitting the observed CDS term structure perfectly using piecewise constant intensities, see O’Kane (2008).

C Estimation of Common Factor Sensitivities

The common factor sensitivities a_i appearing in the specification (2) of individual default intensities can be estimated by ordinary linear regression, and without exploiting specific assumptions on the dynamic evolution of the processes Y, X_1, \dots, X_N except for a mild stationarity condition. As we argue in the following, this model-independent technique only relies on the availability of term structures of risk-neutral survival probabilities for each of the N issuers in the portfolio.

The simple idea that we build upon is the fact that (1) and (2) imply

$$-\lim_{s \searrow 0} \frac{\partial}{\partial s} Q_t(\tau_i > t + s) = \lambda_{i,t} = a_i Y_t + X_{i,t}$$

and that we can calculate this quantity simply by inserting the calibrated survival probabilities on the left-hand-side of this expression.

If we now for fixed i consider the regression

$$W_{i,t} = \beta_{0,i} + \beta_{1,i}(V_t - \bar{V}) + \varepsilon_t \quad t = 1, \dots, T$$

where

$$\begin{aligned} W_{i,t} &= a_i Y_t + X_{i,t} \\ \bar{W}_i &= \frac{1}{T} \sum_{t=1}^T W_{i,t} \\ V_t &= \frac{1}{N} \sum_{j=1}^N W_{j,t} \\ \bar{V} &= \frac{1}{T} \sum_{t=1}^T V_t \end{aligned}$$

and ε_t is a Gaussian noise term, then it follows by standard estimation theory that

$$\hat{\beta}_{1,i} = \frac{\sum_t (W_{i,t} - \bar{W}_i)(V_t - \bar{V})}{\sum_t (V_t - \bar{V})^2}.$$

Under the assumption of stationarity of each of the processes X_1, \dots, X_N, Y (and hence also of W_i and V), we can rewrite the estimated regression coefficient as

$$\hat{\beta}_{1,i} = \frac{\widehat{Cov}(W_i, V)}{\widehat{Var}(V)}. \quad (13)$$

Since X_1, \dots, X_N, Y are mutually independent then for sufficiently large N

$$Cov(W_i, V) = \frac{1}{N}Var(X_i) + a_iVar(Y) \approx a_iVar(Y) \quad (14)$$

and similarly

$$Var(V) = \frac{1}{N^2} \sum_{j=1}^N Var(X_j) + Var(Y) \approx Var(Y) \quad (15)$$

where we have applied the normalization $\frac{1}{N} \sum_i a_i = 1$. By combining (13), (14) and (15) it is now straightforward to see that $\hat{\beta}_{1,i}$ is an approximate estimator of the unknown sensitivity a_i .

To increase numerical robustness in the calculations, we make a small approximation and everywhere replace the instantaneous derivative

$$-\lim_{s \searrow 0} \frac{\partial}{\partial s} Q_t(\tau_i > t + s)$$

with the one-year default probability

$$1 - Q_t(\tau_i > t + 1) = -\frac{Q_t(\tau_i > t + 1) - Q_t(\tau_i > t)}{1 - 0} \approx -\lim_{s \searrow 0} \frac{\partial}{\partial s} Q_t(\tau_i > t + s)$$

since our calibration of the term structure of survival probabilities uses CDS contracts with maturities from 0.5 to 5 years, which results in minor numerical instabilities (across calendar time) in the very short end of the term structure.

D Estimation of Common Factor

Once we have inferred marginal risk-neutral survival probabilities $s \mapsto q_{i,t}(s)$ from CDS spreads and estimated the common factors sensitivities a_i , we are ready to estimate the parameters and the path of the common factor process Y . Throughout the estimation of the common factor process, all the $q_{i,t}(s)$ and all a_i are taken as given (and thus held fixed).

Given an initial path of Y and initial values of the common factor parameters, the estimation procedure runs as follows:

- (i) Calculate the common factor component of survival probabilities

$$E_t^Q \left[\left(-a_i \int_t^s Y_u du \right) \right]$$

for all firms i , all dates t and all maturities s .

- (ii) Use the common factor components $E_t^Q[(-a_i \int_t^s Y_u du)]$ from (i) and the calibrated term structures of survival probabilities $q_{i,t}(s)$ to determine the idiosyncratic component of survival probabilities

$$E_t^Q \left[\left(- \int_t^s X_{i,u} du \right) \right]$$

for all firms i , all dates t and all maturities s using the relation

$$q_{i,t}(s) = E_t^Q \left[\left(-a_i \int_t^s Y_u du \right) \right] \cdot E_t^Q \left[\left(- \int_t^s X_{i,u} du \right) \right]$$

- (iii) Use the idiosyncratic components $E_t^Q[(- \int_t^s X_{i,u} du)]$ from (ii) as input to equation (5) and calculate spreads for the 5 CDO tranches for all dates t (this is what is referred to as the “pricing formula” f in section 3.2).

- (iv) Use the MCMC estimation routine to update the parameters and the path of the common factor Y , and repeat steps (i)-(iv) until convergence.

E 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.

E.1 Conditionals of S, Y, J , and Z

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

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

With the discretization in (8) we have that

$$\begin{aligned} p(Y|\Theta, \Sigma_\epsilon, J, Z) &= \left(\prod_{t=1}^T p(Y_t|Y_{t-1}, \Theta, \Sigma_\epsilon, J, Z) \right) p(Y_0) \\ &= p(Y_0) \prod_{t=1}^T \frac{1}{\sigma \sqrt{\Delta_t Y_{t-1}}} \exp \left(-\frac{1}{2} \frac{[Y_t - (\kappa_0 \Delta_t + (\kappa_1^P \Delta_t + 1)Y_{t-1} + J_t Z_t)]^2}{\sigma^2 \Delta_t Y_{t-1}} \right) \\ &\propto p(Y_0) \sigma^{-T} Y_x^{-\frac{1}{2}} \exp \left(-\frac{1}{2} \sum_{t=1}^T \frac{[Y_t - (\kappa_0 \Delta_t + (\kappa_1^P \Delta_t + 1)Y_{t-1} + J_t Z_t)]^2}{\sigma^2 \Delta_t Y_{t-1}} \right) \end{aligned} \quad (16)$$

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

The conditional posterior of S is found as

$$\begin{aligned}
p(S|\Theta, \Sigma_\epsilon, Y, J, Z) &= \prod_{t=1}^T |\Sigma_\epsilon|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}[S_t - f(\Theta^Q, Y_t)]' \Sigma_\epsilon^{-1} [S_t - f(\Theta^Q, Y_t)]\right) \\
&= |\Sigma_\epsilon|^{-\frac{T}{2}} \exp\left(-\frac{1}{2} \sum_{t=1}^T \hat{e}_t' \Sigma_\epsilon^{-1} \hat{e}_t\right),
\end{aligned} \tag{17}$$

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

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

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

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

Since Z_t is exponentially distributed we have that

$$p(Z|\Theta, \Sigma_\epsilon, Y, J, S) \propto p(S|\Theta, \Sigma_\epsilon, Y, J, Z)p(Z|\Theta, \Sigma_\epsilon, Y, J) \tag{18}$$

$$\begin{aligned}
&\propto p(Y|\Theta, \Sigma_\epsilon, J, Z)p(Z|\Theta, \Sigma_\epsilon, J) \\
&\propto p(Y|\Theta, \Sigma_\epsilon, J, Z) \prod_{t=1}^T \frac{1}{\mu^P} \exp\left(-\frac{Z_t}{\mu^P}\right) \\
&\propto p(Y|\Theta, \Sigma_\epsilon, J, Z)(\mu^P)^{-T} \exp\left(-\frac{Z_\bullet}{\mu^P}\right)
\end{aligned} \tag{19}$$

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

$$p(J|\Theta, \Sigma_\epsilon, Y, Z, S) \propto p(S|\Theta, \Sigma_\epsilon, Y, J, Z)p(J|\Theta, \Sigma_\epsilon, Y, Z) \tag{20}$$

$$\begin{aligned}
&\propto p(Y|\Theta, \Sigma_\epsilon, J, Z)p(J|\Theta, \Sigma_\epsilon, Z) \\
&\propto p(Y|\Theta, \Sigma_\epsilon, J, Z)p(J|\Theta) \\
&\propto p(Y|\Theta, \Sigma_\epsilon, J, Z) \prod_{t=1}^T \left((l^P \Delta_t)^{J_t} (1 - l^P \Delta_t)^{1-J_t}\right) \\
&\propto p(Y|\Theta, \Sigma_\epsilon, J, Z)(l^P \Delta_t)^{J_\bullet} (1 - l^P \Delta_t)^{T-J_\bullet}
\end{aligned} \tag{21}$$

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

E.2 Conditional Posteriors

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

(i) The conditional posterior of the error matrix Σ_{ϵ} is given as

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

The last line follows because $-\frac{1}{2}\sum_{t=1}^T \hat{e}'_t \Sigma_{\epsilon}^{-1} \hat{e}_t = -\frac{1}{2}\sum_{t=1}^T \text{tr}(\hat{e}'_t \Sigma_{\epsilon}^{-1} \hat{e}_t) = -\frac{1}{2}\sum_{t=1}^T \text{tr}(\Sigma_{\epsilon}^{-1} \hat{e}_t \hat{e}'_t) = -\frac{1}{2}\text{tr}(\sum_{t=1}^T \Sigma_{\epsilon}^{-1} \hat{e}_t \hat{e}'_t) = -\frac{1}{2}\text{tr}(\Sigma_{\epsilon}^{-1} \sum_{t=1}^T \hat{e}_t \hat{e}'_t)$. 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.

(ii) The conditional posterior of κ_1^P is found as

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

According to equation (16) we have

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

so

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

where

$$a_t = -\Delta_t Y_{t-1}$$

$$b_t = \kappa_0 \Delta_t + Y_{t-1} + J_t Z_t - Y_t.$$

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

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

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

(iii) For the jump size parameter μ^P the conditional posterior is found as

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

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

- (iv) The same calculations as for the jump-size parameter μ^P yields the conditional posterior of the jump-time parameter l^P as

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

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

- (v) 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

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

Flat priors on both parameters are assumed.

- (vi) 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 Y does not depend on these three parameters. Letting θ_i represent any of the three parameters, the conditional distribution is found as

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

Flat priors on all three parameters are assumed.

- (vii) The latent jump indicators J_t 's are sampled individually from Bernoulli distributions. To see this, note that equation (21) implies that

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

In the actual implementation we use

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

since this is numerically more robust.

- (viii) For the latent jump sizes Z_t we have according to equation (19) that

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

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

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

where $Z_t \geq 0$. Therefore, Z_t is drawn from a $N((Y_t - \kappa_0 \Delta_t) - ((\kappa_1^P + \mu^P \sigma^2) \Delta_t + 1)Y_{t-1}, \sigma^2 \Delta_t Y_{t-1})$ distribution and the draw is rejected if $Z_t < 0$. In practice the number of rejections are small.¹⁹

¹⁹If the draws were frequently rejected the method in Gelfand, Smith, and Lee (1992) could be used.

- (ix) The latent Y_t s are sampled individually by Metropolis-Hastings and for $t = 1, \dots, T-1$ the conditional posterior is

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

For Y_T the conditional posterior is

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

while for Y_0 it is

$$\begin{aligned}
p(Y_0|\Theta, \Sigma_\epsilon, Y_{\setminus Y_0}, J, Z, S) &\propto p(Y_0|\Theta, \Sigma_\epsilon, Y_1, J, Z) \\
&\propto p(Y_1|\Theta, \Sigma_\epsilon, Y_0, J, Z)p(Y_0).
\end{aligned}$$

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

The Fourier inversion in equation (7) is calculated by using Fast Fourier Transform and the number of points used in FFT is 2^{18} . We use Simpson's rule in the Fast Fourier Transform routine as suggested by Carr and Madan (1999), and our results suggest a significant improvement in overall accuracy. 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 60 points. Also, the integration in (6) is done using Gauss-Legendre integration and the number of integration points is 60.

References

- Amato, J. D. and J. Gyntelberg (2005). CDS Index Tranches and the Pricing of Credit Risk Correlations. *BIS Quarterly Review, March*, 73–87.
- Andersen, L., J. Sidenius, and S. Basu (2003). All Your Hedges in one Basket. *Risk*, 67–72.
- Besag, J. (1974). Spatial Interaction and the Statistical Analysis of Lattice Systems. *Journal of the Royal Statistical Association Series B* 36, 192–236.
- BIS (2007). Triennial Central Bank Survey. *Bank for International Settlements*.
- Campbell, J. Y. and G. B. Taksler (2003). Equity volatility and corporate bond yields. *Journal of Finance* 58, 2321–2349.
- Carr, P. and D. B. Madan (1999). Option Valuation Using the Fast Fourier Transform. *Journal of Computational Finance* (2), 61–73.
- 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.
- Collin-Dufresne, P. (2009). A short introduction to correlation markets. *Journal of Financial Econometrics* 7, 12–29.
- Duffee, G. (1999). Estimating the Price of Default Risk. *Review of Financial Studies* 12, 197–226.
- Duffie, D. (1999). Credit Swap Valuation. *Financial Analysts Journal* 55(1), 73–87.
- Duffie, D. (2008). Innovations in credit risk transfer: implications for financial stability. *BIS Working Paper, No. 255*.

- 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*.
- Eckner, A. (2009). Computational Techniques for Basic Affine Models of Portfolio Credit Risk. *Journal of Computational Finance* 13, 63–97.
- Errais, E., K. Giesecke, and L. R. Goldberg (2007). Pricing Credit from the Top Down with Affine Point Processes. *Working Paper*.
- Feldhütter, P. (2006). Can Affine Models Match the Moments in Bond Yields? *Working Paper, Copenhagen Business School*.
- Feldhütter, P. and D. Lando (2008). Decomposing Swap Spreads. *Journal of Financial Economics* 88, 375–405.
- 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.
- Giesecke, K. and L. R. Goldberg (2006). A Top Down Approach to Multi-Name Credit. *Working Paper, Stanford University*.
- Hammersley, J. and P. Clifford (1970). Markov Fields on Finite Graphs and Lattices. *Unpublished Manuscript*.
- Johannes, M. and N. Polson (2006). MCMC Methods for Financial Econometrics. *Handbook of Financial Econometrics* (Chapter).
- 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, 2213–2253.
- Longstaff, F. and A. Rajan (2008). An Empirical Analysis of the Pricing of Collateralized Debt Obligations. *Journal of Finance* 63(2), 529–563.
- Marsaglia, G. and A. Zaman (1991). A New Class of Random Number Generators. *Annals of Applied Probability* 3, 462–480.
- 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.
- Mortensen, A. (2006). Semi-Analytical Valuation of Basket Credit Derivatives in Intensity-Based Models. *Journal of Derivatives*, 8–26.
- Nelson, C. R. and A. F. Siegel (1987). Parsimonious Modeling of Yield Curves. *Journal of Business* 60, 473–489.
- O’Kane, D. (2008). *Modelling Single-Name and Multi-Name Credit Derivatives*. Wiley.
- 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.
- Zhang, B. Y., H. Zhou, and H. Zhu (2009). Explaining Credit Default Swap spreads with the equity volatility and jump risks of individual firms. *Review of Financial Studies* 22, 5099–5131.

Figures

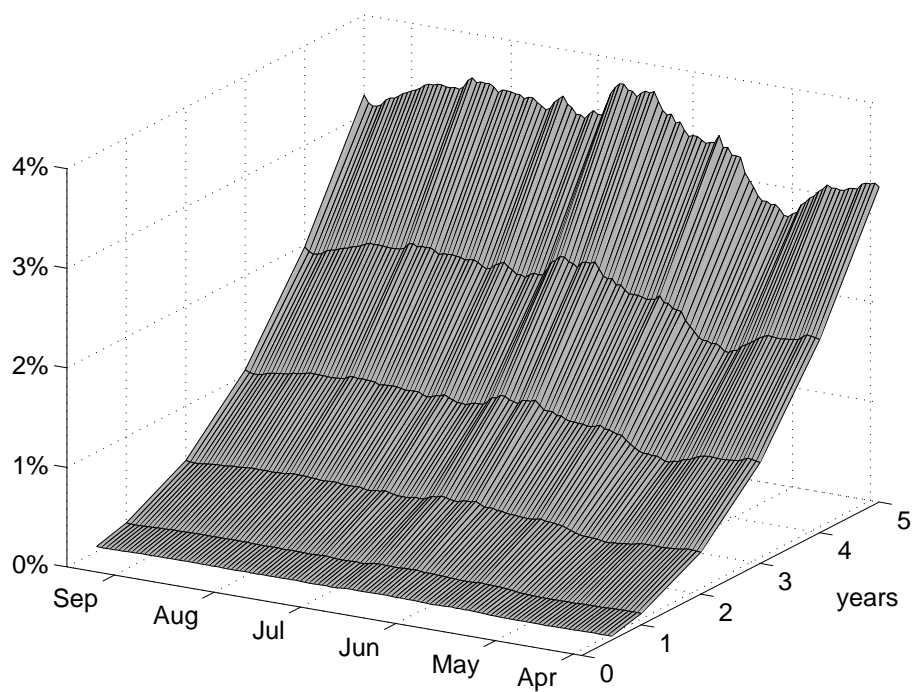


Figure 1: *Default probabilities.* The figure shows the average calibrated term structure of risk-neutral default probabilities for 0 to 5 years over the period March 30, 2006 to September 20, 2006, averaging across all 125 constituents of the CDX NA IG 6 index. Default probabilities are calibrated on a firm-by-firm basis following the procedure outlined in appendix B.




Figure 2: *Common factor sensitivities.* The figure shows the distribution of the estimated common factor sensitivities a_i for the 125 constituents of the CDX NA IG 6 index. The sensitivities are estimated following the procedure outlined in appendix C.

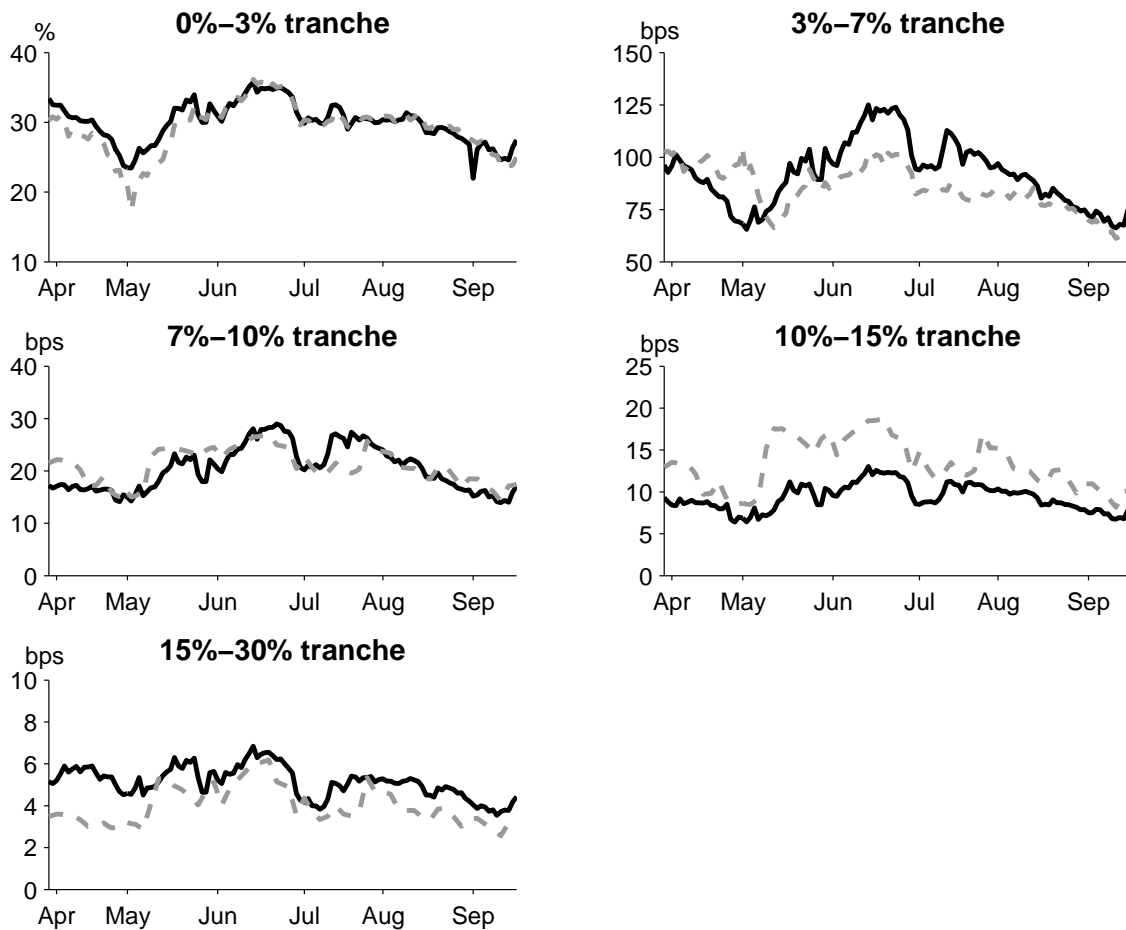


Figure 3: *CDO tranche spreads.* The graphs show the observed (solid black) and model-implied (dashed gray) CDO tranche spreads for the five CDX NA IG 6 tranches: 0% – 3%, 3% – 7%, 7% – 10%, 10% – 15%, and 15% – 30% over the period March 30, 2006 to September 20, 2006. The model-implied spreads are based on the parameter estimates reported in Table 2.

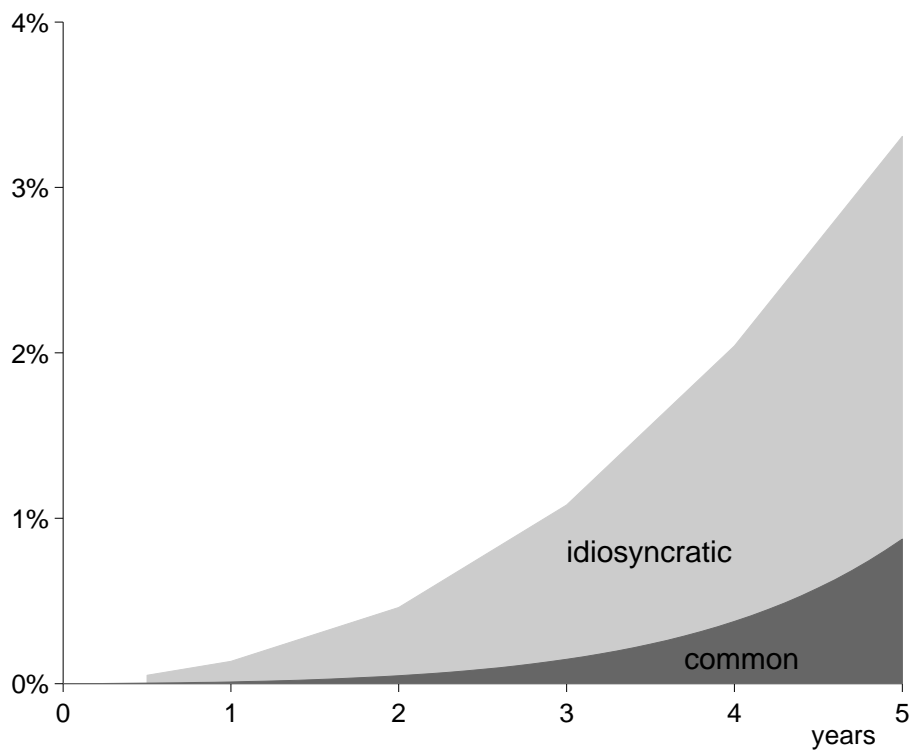


Figure 4: *Average default probabilities.* The figure shows the average term structure of risk-neutral default probabilities, averaging across all 125 constituents of the CDX NA IG 6 index and across all trading days in the period March 30, 2006 to September 20, 2006. The default probabilities are decomposed into their common (dark gray) and idiosyncratic (light gray) parts. The total default probabilities (dark and light gray) are calibrated from CDS spreads (see appendix B), and the common part is calculated using the parameter estimates reported in Table 2.

Tables

Panel A: CDS spreads for CDX NA IG 6 constituents						
Maturity	0.5 yr (in bps)	1 yr (in bps)	2 yrs (in bps)	3 yrs (in bps)	4 yrs (in bps)	5 yrs (in bps)
Mean	6.78	8.75	14.58	21.52	30.52	39.14
Std.	5.77	6.62	11.25	16.74	23.44	29.82
Median	4.86	6.56	10.72	16.07	22.51	29.03
Min.	0.41	1.73	2.65	2.94	3.99	5.45
Max.	56.46	59.82	103.73	140.48	181.60	222.19
Observations	15,000	15,000	15,000	15,000	15,000	15,000

Panel B: CDO tranche spreads for CDX NA IG 6 tranches					
Tranche	0% – 3% (in %)	3% – 7% (in bps)	7% – 10% (in bps)	10% – 15% (in bps)	15% – 30% (in bps)
Mean	29.95	91.83	20.43	9.33	5.13
Std.	2.92	15.37	4.27	1.58	0.74
Median	30.29	92.48	20.31	9.06	5.17
Min.	21.97	65.52	13.96	6.40	3.54
Max.	35.75	125.02	28.97	13.02	6.84
Observations	120	120	120	120	120

Table 1: *Summary statistics.* Panel A reports summary statistics for CDS spreads of the 125 constituents of the CDX NA IG 6 index over the period March 30, 2006 to September 20, 2006. Panel B reports summary statistics for the five CDO tranches: 0% – 3%, 3% – 7%, 7% – 10%, 10% – 15%, and 15% – 30% of the CDX NA IG 6 index over the same period.

$\kappa_0 (\times 10^5)$	κ_1	$\sigma (\times 10^2)$
2.32 (2.15, 2.58)	0.94 (0.90, 0.99)	1.66 (1.48, 1.81)
$l (\times 10^3)$	$\mu (\times 10^2)$	
3.74 (2.54, 4.59)	1.59 (1.11, 2.12)	
κ_1^P	$l^P (\times 10^2)$	$\mu^P (\times 10^{10})$
-3.45 (-15.09, 5.08)	2.54 ($3.40 \cdot 10^{-13}$, $2.18 \cdot 10^4$)	8.34 (8.19 , $1.57 \cdot 10^8$)
$\sqrt{\Sigma_{11}}$	$\sqrt{\Sigma_{22}}$	$\sqrt{\Sigma_{33}}$
0.11 (0.10, 0.33)	0.19 (0.15, 0.43)	0.16 (0.11, 0.51)
$\sqrt{\Sigma_{44}}$	$\sqrt{\Sigma_{55}}$	
0.35 (0.30, 0.64)	0.38 (0.28, 0.67)	

Table 2: *Parameter estimates (common factor).* The table reports point estimates and 95% confidence intervals (in parenthesis) for the parameters of the multi-name default model outlined in section 2.

Tranche	0% – 3% (in %)	3% – 7% (in bps)	7% – 10% (in bps)	10% – 15% (in bps)	15% – 30% (in bps)
Mean	-0.7888	-6.93	0.55	3.96	-1.12
RMSE	1.7907	14.31	2.94	4.42	1.31

Table 3: *CDO pricing errors.* The table reports mean and standard deviation of the daily pricing errors for each of the five CDO tranches: 0%–3%, 3%–7%, 7%–10%, 10%–15%, and 15%–30% of the CDX NA IG 6 index over the period March 30, 2006 to September 20, 2006. The pricing errors are calculated as model-implied minus observed tranche spreads, and the model spreads are based on the parameter point estimates in Table 2.

Maturity	0.5 yr	1 yr	2 yrs	3 yrs	4 yrs	5 yrs
Total DP	0,051%	0,134%	0,460%	1,080%	2,042%	3,309%
Common part of total DP	0,003%	0,010%	0,048%	0,147%	0,376%	0,874%
Common part in % of total DP	5,88%	7,64%	10,40%	13,59%	18,42%	26,41%

Table 4: *Average default probabilities.* The table reports average risk-neutral default probabilities (DP), averaging across all 125 constituents of the CDX NA IG 6 index and across all trading days in the period March 30, 2006 to September, 2006. “Total DP” reports the total default probability and corresponds to the light gray area in Figure 4, and “Common part of total DP” similarly expresses the common factor part of the total default probability corresponding to the dark gray area in Figure 4.

$\kappa_0 (\times 10^6)$	κ_1	σ
9.08	0.80	0.14
0.41	0.87	0.16
(36.61)	(0.24)	(0.08)
$l (\times 10^3)$	μ	
4.48	8.93	
2.68	0.30	
(6.05)	(59.12)	
κ_1^P	$l^P (\times 10^{-2})$	$\mu^P (\times 10^9)$
0.14	1.31	1.66
-0.60	1.30	1.66
(5.79)	(0.34)	(0.07)
$\sqrt{\Sigma_{11}} (\times 10^4)$	$\sqrt{\Sigma_{22}} (\times 10^4)$	$\sqrt{\Sigma_{33}} (\times 10^4)$
1.10	1.19	1.37
1.01	1.04	1.18
(0.26)	(0.42)	(0.52)
$\sqrt{\Sigma_{44}} (\times 10^4)$	$\sqrt{\Sigma_{55}} (\times 10^4)$	
1.40	1.66	
1.23	1.09	
(0.69)	(2.55)	

Table 5: *Parameter estimates (idiosyncratic factors).* The table reports mean, median and standard deviation (in parenthesis) of the 125 parameter point estimates resulting from the idiosyncratic factor estimations in the multi-name default model outlined in section 2.