Correlated Default Modeling with a Forest of Binomial Trees

Santhosh Bandreddi
Merrill Lynch
New York, NY 10080
santhosh_bandreddi@ml.com

Sanjiv Das
Santa Clara University
Santa Clara, CA 95053
srdas@scu.edu

Rong Fan
Gifford Fong Associates
Lafayette, CA 94549
rfan@gfong.com

October 29, 2007

Abstract

This paper exploits the endogenous default function framework of Das and Sundaram (2007) to develop an approach for modeling correlated default on binomial trees usually used for pricing equity options. We show how joint default contracts may be valued on these trees. The model accommodates different correlation assumptions and practical implementation considerations. Credit portfolio characteristics are examined within the model and found to be consistent with stylized empirics. Risk premia for default are computable and shown to be relatively higher for poor quality firms. Equity volatility is shown to impact correlated credit loss distributions substantially. Two kinds of default dependence are explored, one coming from default intensity correlations, and the other from further conditional dependence in defaults after accounting for intensity correlations (residual copula correlation). Both are found to impact credit loss distributions, though the absence of either makes these distributions less sensitive to correlation assumptions; on balance intensity correlations are more critical.

1 Introduction

The modeling of correlated default has become an increasingly important area of research in the recent decade. The dollar volume of assets residing in credit portfolio products has grown tremendously and is of the order of $1 trillion.[1] For 2004, the Bank for International Settlements (BIS) estimated synthetic CDO volumes of $673 billion, and cash CDO issuance of $165 billion. Modeling default correlations is crucial when pricing n-th to default contracts, which pay off when a specified number of firms in a basket default within a pre-specified maturity. In this paper we offer a parsimonious approach to the problem, involving a computational extension of the widely used methodology of equity option binomial trees.

Semi-analytic approaches for the valuation of portfolio credit products have been in use for some time. Andersen, Sidenius and Basu (2003) recently developed a recursion technique for computing loss distributions in the presence of correlated default in a factor model framework. Contemporaneously, detailed work in the same vein by Gregory and Laurent (2003), and Gregory and Laurent (2004) developed and extended the idea using probability generating functions. These models, beginning with the Gaussian copula of Li (2000), and extended to other copulas have become the approach of choice for practitioners (see Burtschell, Gregory and Laurent (2005) for a comparison of copulas in these models). Mortensen (2005) extends the semi-analytic valuation approach to the realm of intensity-based models. A comparison of copulas in an intensity framework is undertaken in Das and Geng (2004), and Luciano (2005) provides a copula comparison under the risk-neutral measure. Longstaff and Rajan (2006) develop a completely new approach to modeling credit portfolio losses using a direct approach that does not require modeling individual correlations. This top-down approach (see also Giesecke and Goldberg (2005); Schönbucher (2005); Sidenius, Piterbarg and Andersen (2004); and Egami and Esteghamat (2006)) is a new innovation that offers many advantages for modeling credit portfolio products.

The basis of existing approaches for the valuation of credit portfolio products (baskets of various types) lies in the manner in which the models account for credit correlations. In the copula approach, this has become proxied by asset correlations, though these are only valid where the Merton (1974) model is applied in a single period static framework (Mortensen (2005)). In intensity-based models, credit correlations are implied from spread correlations when the models are calibrated to default swap markets. The framework developed in Das and Sundaram (2000); Carayannopoulos and Kalimipalli (2003) and Das and Sundaram (2007) suggests a third approach, where equity return correlations may be used to drive intensity correlations. This is the approach explored in this paper. The framework directly uses observable equity prices in an intensity-based model to produce correlated default in an arbitrage-free setting. Further, correlated default may be either imposed in a Cox process framework, where the assumption of doubly stochastic processes is maintained, i.e. where defaults, conditioned on the path of intensities, are independent, or where further conditional dependence is imposed. In the latter case, freedom to choose a copula for imposing conditional dependence is allowed for. Hence, we are able to model the extraction of default intensity functions, and then provide the correlations

[1] ISDA reported in June 2004, that all credit product notionals were in the order of $5.4 trillion. RiskNews reports the notional amount of credit derivatives at $12.43 trillion, by mid 2005. Estimates of volumes vary considerably by source, but the growth rate in this area is clearly high.

necessary to value basket products.

Our approach comprises three steps. One, we develop for all issuers equity binomial trees with default risk. Two, we calibrate these trees to the default swap market for each issuer in the credit portfolio. Three, using exogenously supplied equity return correlations, our collection of trees delivers default correlations. Our correlated default "forest" supports the modeling of several products. The first two steps above are based on ideas in prior related work. The final step, i.e. using a forest of calibrated trees to examine default risk distributions is a new approach that is computationally attractive.

The model enables us to explore many questions of interest, for example: (i) How are equity correlations related to correlations of hazard rates of default? (ii) How does equity correlation (and the consequent hazard rate correlation) impact the pricing of correlated default contracts? (iii) How do the shapes of default swap spread curves for different issuers impact correlated default? (iv) Which is more important, correlation in default intensities or conditional correlations of default?

Das and Sundaram (2000) developed a model of "endogenous" default risk within the Heath, Jarrow and Morton (1990) model. Carayannopoulos and Kalimipalli (2003) and Das and Sundaram (2007) (DS) developed an extended framework for correlated default under the risk-neutral measure using a model in which default risk is modeled jointly with equity and interest rate risk. In these papers, credit risk is specified as a function of the other values on the pricing trees used. It is in this connotation that default is specified as "endogenous" since the hazard function is specified completely as a function of the other information available in the model. In contrast, the models of Davis and Lischka (1999) and Schönbucher (2002) specify hazard rates as a separate exogenous process; such models have more parameters, are harder to callbrate, but may be more flexible. In contrast, the class of endogenous models used here is extremely parsimonious. Since the approach is based on a natural modification of the Cox, Ross and Rubinstein (1979) (CRR) model, it is easy to implement. Recently, Linetsky (2004) has developed a continuous time version of this approach. The discrete time approach in this paper further allows for the pricing of contracts with American style features.

So far the theoretical connection between CDS curves and correlated default is not well researched. For example, we do not as yet know whether credit portfolio risk (which may itself be defined in many ways) increases or decreases if the shape of the CDS curve is humped versus flat. It is known that the slope of the CDS curve depends on the credit quality of the issuer in question (see Helwege and Turner (1999)). High quality issuers tend to have upward sloping CDS curves, medium quality issuers are likely to have humped curves, and the poorest quality issuers have downward sloping credit curves. The reasons for such shapes are based on simple conditioning arguments pertaining to the hazard rates of default (see Duffie and Lando (2001) for example). Whether these theoretical features of the CDS curve's slope affect correlated default is not known, and here, with a simple model, we are able to explore this question directly.

The framework of this paper relies on a three-parameter endogenous hazard function calibrated on a defaultable binomial tree to the default swap market. This function models default using these parameters, given the exogenous stock price and time. Equity levels proxy for default and since they change over time, the default function is dynamic. Parsing the correlation in equity across issuers through this default function results in dynamic hazard rates with correlation.

The primary benefit to this modeling approach is that the three extracted parameters are directly related to the *level*, *slope* and *curvature* of the CDS spread curve. Hence, intensity correlations in the model are also directly related to the shape of CDS curves. The model may be used to determine the sensitivity of default correlations to each of these three facets of the CDS market.

There is an additional benefit to this framework. Since the model is based on equity binomial trees, it provides a link between the credit markets and the equity option markets. There is some discussion in the literature that credit risk may account for some of the option skew (see Hull, Nelken and White (2003), Linetsky (2004), Carr and Wu (2005)), and thus implicitly, co-movements in the option skew must also be related to correlated default risk amongst issuers. This is an open area of research, both theoretical and empirical.

We undertake many calibration and numerical experiments in the paper to examine the features of the model and its relationship to observed phenomena in the credit markets.

1. We demonstrate how the model may be calibrated to the default swap market and the equity options market. We find that the fit of the model to market CDS curves is extremely good for both high and low credit quality firms.

2. The model may be used to generate *forward* default intensities, and we find that correlation across issuer intensities declines as the forward horizon increases. Because forward correlations have a mathematical tendency to decline, *conditional* correlated default (hazard rates) in portfolios is of greater concern in the short-run than in the long-run.

3. The framework is flexible enough to allow portfolio loss generation with the assumption of conditional independence (i.e. a doubly stochastic model) or the case where further conditional correlation is allowed. In the latter case, any copula may be chosen for conditional correlation (we employ the Gaussian copula). We illustrate loss distributions when the doubly stochastic assumption holds and when it is violated.

4. Portfolio losses are found to be less sensitive to conditional (on intensities) correlation than to default intensity correlation. We find that credit portfolios demonstrate much greater sensitivity to correlation assumptions only when both, intensities and conditional default are allowed to be correlated. Das, Duffie, Kapadia and Saita (2007) find that conditional correlation exists in the data; here we quantify the impact of this on credit loss distributions.

5. In the model, increases in equity volatility translate naturally into higher variation in credit portfolio losses. Empirically, it is known that economic epochs in which equity correlations were high were also periods of substantially higher credit correlations and credit portfolio losses (see Das, Freed, Geng and Kapadia (2006)), and the model reproduces this behavior.

6. We calibrated high and low quality firms to the model, and then generated default probabilities for different horizons. Firms with low levels of default risk have *convex* (in maturity) cumulative default probability functions, implying that they revert to being of slightly lower quality, just as evidenced with rating transitions. Conversely, low quality firms have *concave* cumulative probability functions, as they revert to better quality, conditional on survival.

7. We also found that the model endogenously generates lower risk premia for good quality issuers than for bad quality ones; this is of course consistent with investors being risk-averse to default occurrence.

The rest of the paper proceeds as follows. Section 2 presents the basic features of the approach. Section 3 provides examples of use of the model and calibration. Section 4 conducts various basic analyses of default correlation and portfolio losses. In Section 5, the link between equity correlations and default correlations is explored, and the sensitivity of credit portfolios to various aspects of the model is analyzed. Section 6 suggests various extensions that make the model palatable to modeling variations undertaken in practice. Section 7 concludes.

2 Model

The model implementation consists of the following components:

1. An equity binomial tree with embedded default risk for each issuer in a credit portfolio.
2. Calibrating the tree to the default swap market.
3. Using the default *forest* to compute correlated default risk based on equity correlations.

2.1 Building a defaultable binomial tree

We briefly summarize the approach of Das and Sundaram (2007) here.[2] We recognize that equity is a security which receives a zero recovery rate in the event of default.[3] We build a binomial tree with constant interval h (expressed in units of years). The maturity of the tree is T, hence, the number of periods on the tree is $n = T/h$. In a binomial framework with default, the stock price at time t is denoted S_t, and the next period's stock price S_{t+h} may take one of three possible values:

$$S_{t+h} = \begin{cases} u\, S_t & \text{(an up move)} \quad \text{w/prob} \quad q(1-\lambda_t) \\ d\, S_t & \text{(a down move)} \quad \text{w/prob} \quad (1-q)(1-\lambda_t) \\ 0 & \text{(default)} \quad \text{w/prob} \quad \lambda_t \end{cases} \tag{1}$$

As usual, u and d are up and down shift parameters. In the CRR model we have $u = \exp(\sigma h)$, $d = 1/u$, where σ is the stock volatility. The risk-neutral evolution of the stock price on the tree has a probability measure which is a function of q (the parameter that modulates the likelihood of the stock moving up or down) and λ_t, which is the one-period default probability (and also affects the probability of the stock remaining solvent or defaulting).

[2]This is a simplified version of their model.

[3]This is not necessarily the case if there are deviations from the absolute priority rule. We may easily accommodate a modification of the model for violations of APR, where we stipulate a recovery rate for equity as well, in which case it recovers more then zero. Technically, this does not impact the model at all, barring an implicit change in the risk-neutral probability measure governing the evolution of the stock price in the model.

The probability of default over period $(t, t+h)$ is specified as

$$\lambda_t = 1 - \exp(-\xi_t \; h) \tag{2}$$

where ξ_t is the constant default intensity over the interval $(t, t+h)$. We specify the intensity (which we denote also as the default function) with the following three parameter equation:

$$\xi_t = \exp(\alpha + \gamma t) S_t^{-\beta} \tag{3}$$

See Carayannopoulos and Kalimipalli (2003) and Linetsky (2004) for similar (though simpler) choices of default function. The three parameters $\{\alpha, \beta, \gamma\}$ modulate the intensity at all times t. The parameter α may be interpreted as the *level* modulator; β impacts the *curvature*, and γ is a *slope* parameter. This function determines the shape of many credit term structures, all of which are related to each other, such as that of cumulative default probability, survival probabilities, or credit spreads. This default function is calibrated to credit market data, specifically the default swap spread curve. We assume that recovery rates in the CDS market are constant and are denoted ϕ. Under the usual martingale conditions, we require that

$$E(S_{t+h}) = e^{rh} S_t, \tag{4}$$

where r is the constant risk free rate of interest. Expanding the left hand side of this equation using equation (1) and re-arranging, we can show that

$$q = \frac{e^{(r+\xi)h} - d}{u - d}. \tag{5}$$

Therefore, once we are given the parameters of the default function, i.e. $\{\alpha, \beta, \gamma\}$, initial stock price S_0 and volatility σ, the interval h, and maturity T, we can construct the entire tree with the probability measure from the equations above. Equity volatility σ is exogenously available from many sources. S_0 is observable, and h, T are standard choices for equity binomial trees. Thus, implementation of the model only requires fitting the three parameters of the default function.

2.2 Calibrating the parameters

A suitable instrument to proxy for credit risk is a credit default swap (CDS). A CDS contract is one where the buyer pays a fixed stream of payments at regular intervals to the seller who insures the buyer against losses from default of a pre-specified reference instrument, usually a bond. Assume a recovery rate of ϕ. No payments are made by the seller until default, when the loss on default $(1-\phi)$ per dollar is paid by the seller of the CDS to the buyer. CDS are quoted in terms of a spread which is paid periodically by the buyer of the CDS to the seller. We denote this spread as c per annum. Hence, the spread payment per period is $c \times h$. The plot of CDS spreads against time for CDSs of varying maturity is the term structure of credit spreads.

Given at least three credit default swaps spreads of varied maturity, we can determine the values exactly of the parameter set $\{\alpha, \beta, \gamma\}$. If we have more CDS spreads, then a least squares fit may be used. At a practical level, using the known level of the stock volatility allows using just three default swap spreads to extract the parameters of the default function. However, we may also choose to calibrate the model to CDS spreads and options together, thus calibrating the entire parameter set $\{\alpha, \beta, \gamma, \sigma\}$ simultaneously.

Fair pricing of a CDS is characterized by an equality of expected present values of payments by buyer and seller under the risk-neutral probability measure (see Duffie and Garleanu (2001)). The CDS spread c may be determined on the tree by equating the expected present value of payments made by the buyer with the expected present value of the possible loss on default to be paid by the seller.

We assume that a spread premium payment is made at the end of each period only if default has not occurred (an alternative convention is to assume spread payments are made at the start of the period). By putting a spread payment of $c \times h$ at each node of the tree (except the root node and defaulted nodes), we may compute the expected present value of all these payments via backward recursion on the tree. We denote this amount as P_C. Note that since the payments are the same at each node (unless h is not constant on the tree), we may also write this amount as chP_1, where P_1 is the present value of $1 received at all non-default nodes on the tree. (This implies that we assume the notional value of the CDS is $1).

The expected present value of losses is determined by putting the loss amount of $(1-\phi)$ on each defaulted node, and zero elsewhere, and then discounting all cashflows to the root node. We denote this amount as P_L. Setting P_L to be equal to chP_1 which is the condition for a fairly priced CDS, we may compute the quoted spread c. This is equal to

$$c = \frac{P_L}{P_1} \times \frac{100}{h}, \tag{6}$$

where multiplication by 100 expresses the default swap spread in basis points.

2.3 Correlated default from multiple trees

We assume that there are N issuers in a credit portfolio. For each issuer i, we calibrate a defaultable binomial tree. Thus we will have an intensity function for each issuer, i.e.

$$\xi_{it} = \exp(\alpha_i + \gamma_i t) S_{it}^{-\beta_i}, \quad \forall i. \tag{7}$$

From this equation, it is obvious that a simulation of a path of stock prices S_{it} for issuer i implies a simulation of the path of default intensities ξ_{it} for the issuer. Hence, if we simulate a joint process for all stock prices (with the appropriate correlation), then it also implies a simulation of the joint stochastic process for $\xi_{it}, \forall i$.

We begin our simulation positioned at the root node of each issuer. At these nodes we know the values of $S_{it}, \forall i$. Hence, we also know the initial values of all the intensities, ξ_{it}. Our simulation proceeds in two steps:

1. We first check each firm for default. We compute λ_{it} from ξ_{it}. Recall that $\lambda_{it} = 1 - \exp(-\xi_{it} h)$. Then we draw a set of independent uniform random numbers u_i. Firm i defaults if $u_i \leq \lambda_{it}$. Note here that by assuming the draws are independent, we are invoking the assumption of doubly stochastic processes. There is evidence that this assumption is

not fully supported (see Das, Duffie, Kapadia and Saita (2007)). Hence we may draw random variables u_i with correlation from *any* joint distribution using a copula or other appropriate technique; thus it becomes easy to impose conditional dependence as required.

2. In the second step, if the firm has not defaulted, we then wish to move along each tree in our binomial forest, where the stock price moves up or down.

   Suppose the correlation (not covariance) matrix of stock returns, conditional on not defaulting, is denoted Σ. This is identical to the covariance matrix for random variables of mean zero and variance one. We sample with correlation Σ an N-vector of correlated standard normal random variables $x = \{x_1, x_2, \ldots, x_N\}$. We use this vector of random variables to move along each issuer's tree. Given the probability q_i for each issuer as in equation (5), we move up the tree if

$$N(x_i) \leq q_i, \quad N(x_i) = \int_{-\infty}^{x_i} \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{1}{2}z^2\right\} \, dz \tag{8}$$

   If $N(x_i) > q_i$, then the stock moves down along the tree for issuer i. These moves are made for all issuers.

We keep track of where we are in each issuer tree. We repeat these two steps period by period, until we reach time T. It is simple to track where and when an issuer defaults, and we may also trigger payments in the event of a default, as well as present value such payments to value a derivatives contract based on default. Thus, by repeated sampling in steps 1 and 2 above we traverse an entire forest of defaultable binomial trees in a dynamic manner over time, and can simulate default losses in any credit portfolio.[4]

3 Fitting the Model

In this section, we illustrate the application of the model using many different analyses that may be of interest to a modeler or trader in basket default products.

3.1 Sample spread curve

We analyze the various shapes of spread curves that may be generated using the three default function parameters: $\{\alpha, \beta, \gamma\}$. We choose the following baseline values for the three parameters:

$$\alpha = -0.5, \quad \beta = 1, \quad \gamma = 0.1$$

If we choose a stock price of 50, then the default function $\xi = \exp(\alpha + \gamma t)S^{-\beta}$ may be plotted for a range of maturities up to 10 years. This plot is presented in Figure 1. The generated intensity varies from over 0.01 at short maturities to over 0.03 at longer maturities. Since γ is positive, the intensity term structure is upward sloping.

[4]See Duffie and Singleton (1999) for a more detailed exposition on simulating correlated defaults.

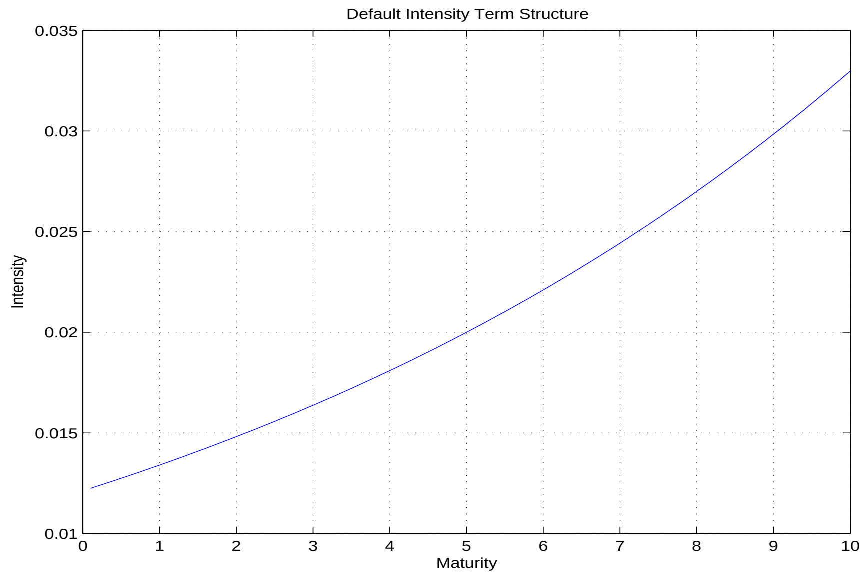

Figure 1: Intensity term structure with base parameters. $\alpha = -0.5$, $\beta = 1$, $\gamma = 0.1$, $S = 50$.

3.2 Calibration

We present calibration examples to demonstrate how the model may be fitted to the data on default swap spreads. We extracted data from the `CreditGrades` web site. Stock volatility σ is given on the site. We also extracted default swap spreads for various maturities (1,2,3,4,5 years) and used these to fit implied values for the default function parameters $\{\alpha, \beta, \gamma\}$. Fitting was undertaken by a sum of least squares minimization of differences between the actual and model CDS spreads. Table 1 shows the results of this exercise. Panel A of this table contains the input values, and Panel B shows the fitted parameters of the model.

In the table, the first three firms have upward sloping spread curves, and the latter three firms have downward sloping curves. Examining the calibrated parameters, we notice that the slope is determined by the variable γ which is the coefficient on time in the default function. For the first three firms, $\gamma > 0$, and for the last three, $\gamma < 0$. The parameter β should be positive for all firms as it signifies that default risk is inversely proportional to the level of the stock price. However, for one firm, GM, $\beta < 0$; this is widely attributed to the credit and equity markets taking divergent views. The equity price is holding up because of the speculation of a buyout but the credit market is pricing it as a junk bond. However, the fit of the model to the spread curve is extremely good. It is likely that since this coefficient is small, and since GM is deep in junk status, the sensitivity of default risk to the stock price is now minimal.

The first three firms in the table have lower levels of default risk than the latter three. Their spread curves slope upwards because on average, with the passage of time, conditional credit quality is skewed to decline. In poorer quality firms, such as the last three in the table, the

Table 1: Calibrated default functions for various firms. The data for GM, SUNW, BARC and PRF is extracted on June 6, 2005; for AMZN and NWAC on June 13, 2005. CreditGrades PD is the probability of default (PD) under the real world probability measure. The model PD is under the risk-neutral measure. Hence, it is higher than that under the physical measure.

Panel A: Input Data								
Ticker	Stk(S)	σ	ϕ	CDS Spreads(bps)				
				1yr	2yr	3yr	4yr	5yr
BARC (UK)	5.2	33.40%	0.5	9.21	12.5	15.7	18.4	20.6
SUNW	3.95	61.20%	0.5	5.07	48.6	109	161	200
AMZN	34.95	59.60%	0.5	0.07	5.0	23.5	50.7	79.1
GM	30.93	34.50%	0.5	3723	1954	1364	1069	892
PRF (Italy)	0.11	86.20%	0.5	16383	8546	5775	4430	3605
NWAC	6.33	65.50%	0.5	5603	2988	2115	1677	1414

Panel B: Calibrated Parameters					
	Parameters			PD (5 Yrs)	
Ticker	α	β	γ	Model	CreditGrades
BARC (UK)	-3.2	1.89	0.15	0.21	0.10
SUNW	-3.93	1.16	0.65	0.52	0.19
AMZN	-2.45	1.44	0.73	0.32	0.08
GM	-0.8	-0.37	-3.17	0.83	0.46
PRF	0.87	0.53	-2.37	0.99	0.77
NWAC	3.9	1.49	-1.93	0.91	0.59

opposite usually occurs, i.e. spread curves are downward sloping. We plot the empirical and fitted CDS curves for the six firms presented in Table 1. These are presented in Figures 2 and 3. In Figure 2 the three firms are chosen from Table 1 with upward sloping spread curves, and in Figure 3 the three firms have downward sloping spread curves. The plots show that the empirical and fitted spread curves are extremely close to each other, signifying a good fit to the default function for each firm.

4 Analysis and Applications

In this section we examine the impact of varying input parameters on the loss distribution in a credit portfolio of 10 identical bonds with the same base value of parameters as we had in the previous section. These analyses look at how various input parameters impact the loss distribution of a credit portfolio.

4.1 Simulation of correlated losses

We undertake numerical experiments to examine what effect matters the most in determining the shape of the loss distribution in a credit portfolio. In order to do this, we assume that we have 10 identical firms with the same base level parameters we used before, i.e.

$$S = 50,\ \sigma = 0.3,\ \alpha = -0.5,\ \beta = 1,\ \gamma = 0.1,\ r_f = 0.03.$$

We also assume that the average correlation between the firms' equity returns is given by parameter ρ. This is consistent with the usual assumption in practice of "flat" correlations, where models are calibrated to a single average correlation parameter. Based on these values, we simulate losses for the credit portfolio using the procedure outlined in Section 2.3. We ran 10,000 simulations of paths of stock prices, generating prices each month for 5 years. These returns are generated with the correct amount of correlation ρ. We thus have

$$S_{it}, \quad i = 1 \ldots 10, \quad t = 1 \ldots 60.$$

Correspondingly, these stock prices are then converted into an equivalent number of intensities, ξ_{it}. The total intensity for each firm for the period of 5 years is

$$\theta_i = \sum_{t=1}^{60} \frac{\xi_{it}}{12}$$

where we divide by 12 because the intensities are in annual terms. The survival probability for each firm over 5 years is then

$$s_i = \exp(-\theta_i). \quad \forall i.$$

We determine defaults by generating independent uniform random variables, u_i, and default occurs if $u_i > s_i$. We repeat this process 10,000 times, each time recording the number of defaults from the 10 firms. The baseline level of correlation is taken to be $\rho = 0.4$. On completion of the simulation, we record the moments of the number of firms that default.

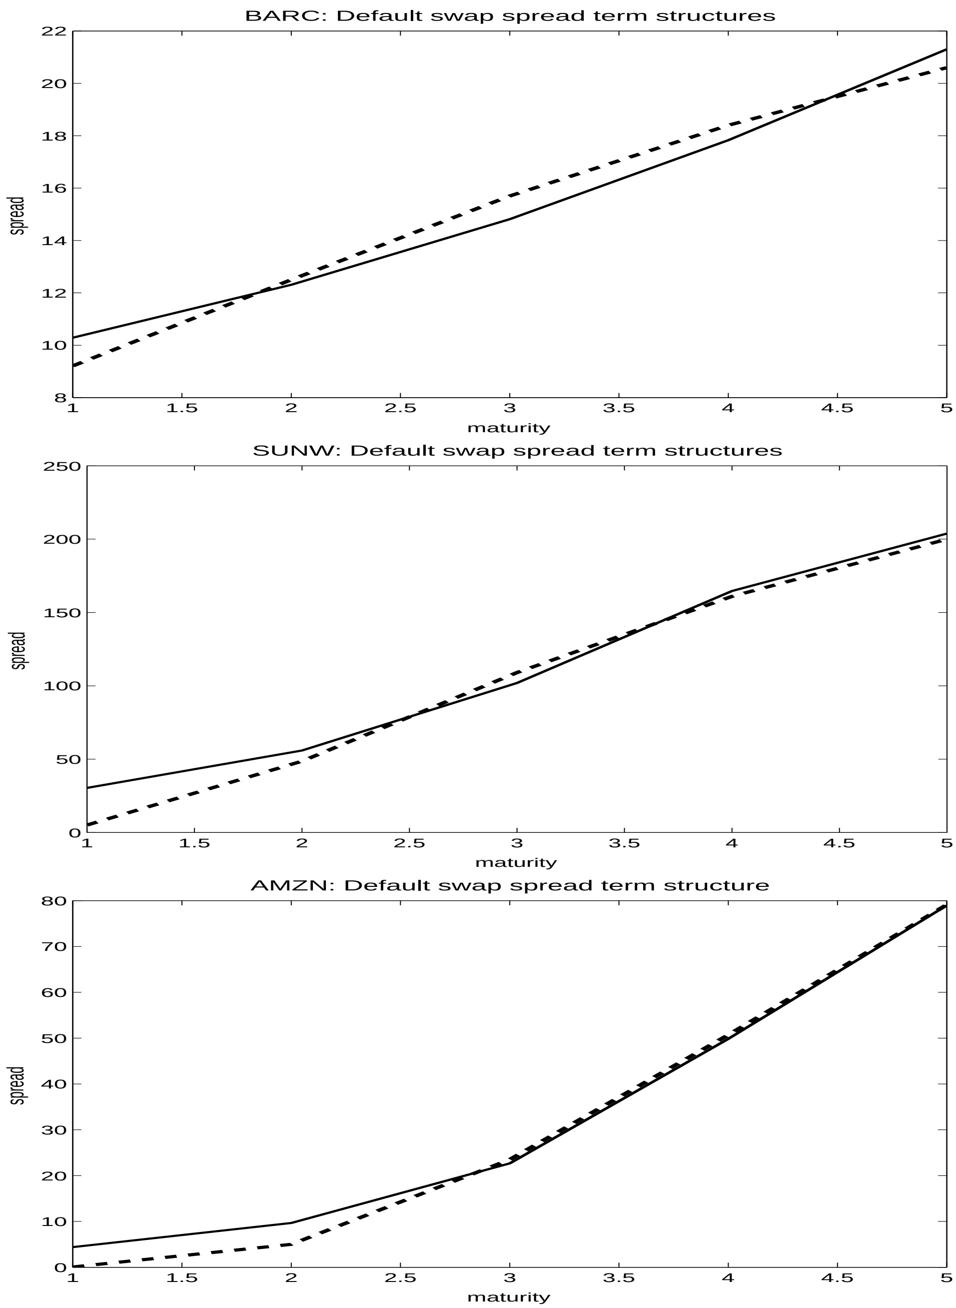

Figure 2: We provide three plots of the empirical CDS spreads against the spreads (in bps) generated from the calibrated model. The three firms chosen here display upward sloping spread curves. The three firms here are (from top to bottom): BARC, SUNW, AMZN. The bold line represents the empirical values, and the dotted line the fitted spread curves.

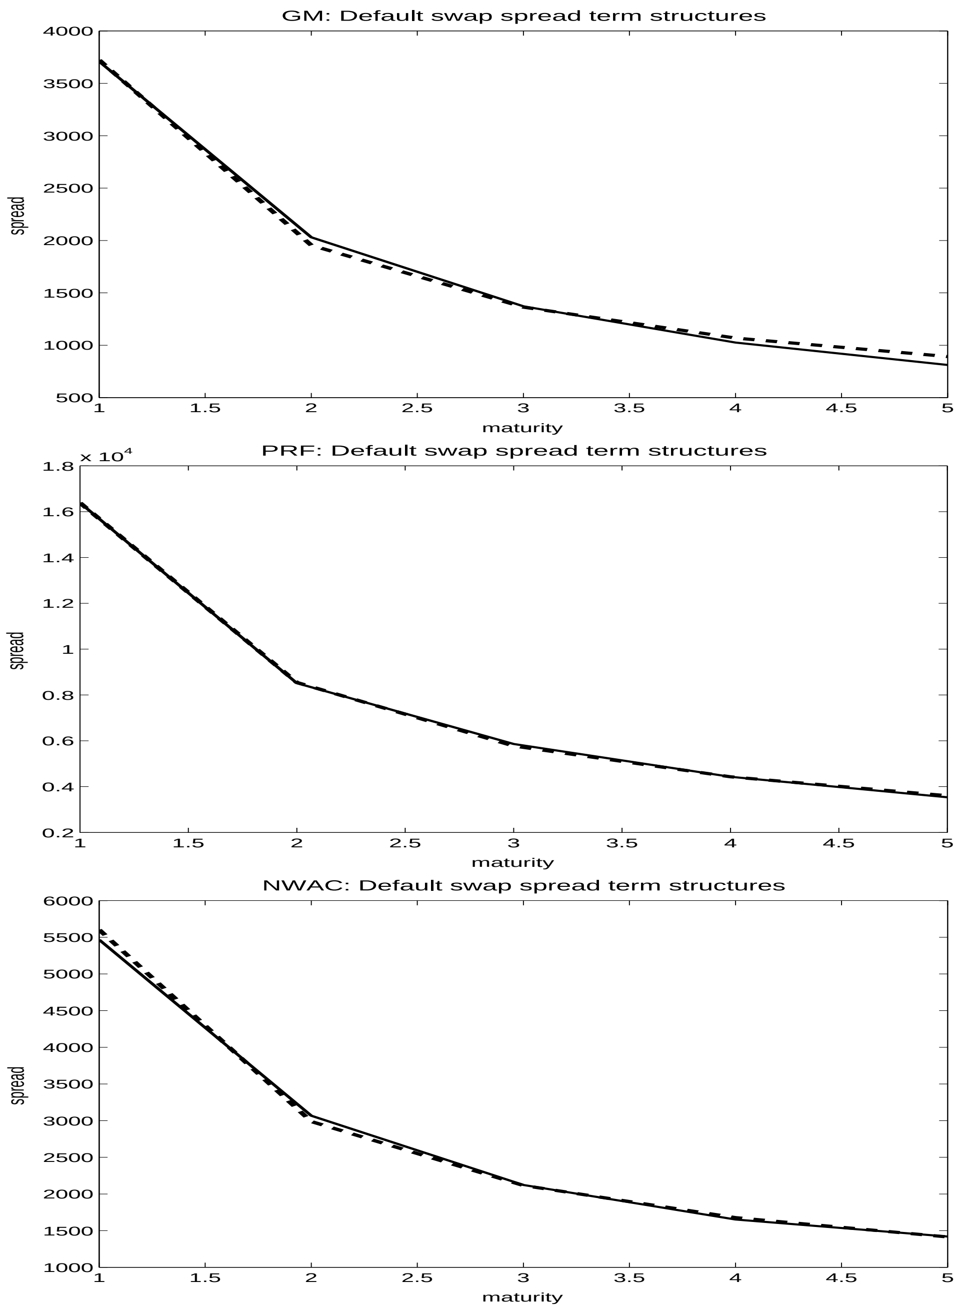

Figure 3: We provide three plots of the empirical CDS spreads against the spreads (in bps) generated from the calibrated model. The three firms chosen here display downward sloping spread curves. The three firms here are (from top to bottom): GM, PRF, NWAC. The bold line represents the empirical values, and the dotted line the fitted spread curves.

Table 2: Moments of the number of defaults when equity correlation is varied. The number of firms used in the credit basket is $n = 10$. The base level of parameters used in the simulation are: $S = 50$, $\sigma = 0.3$, $\alpha = -0.5$, $\beta = 1$, $\gamma = 0.1$, $r_f = 0.03$, $\rho = 0.4$. The simulated loss distributions are computed for a horizon of 5 years, and the simulation step is monthly. The defaults are drawn under the assumption of independent defaults after conditioning on default probabilities, i.e. the doubly stochastic assumption.

Correlation (ρ)	Mean	Variance	Skewness	Kurtosis
0.00	0.68	0.64	1.11	4.12
0.25	0.69	0.64	1.08	4.00
0.50	0.69	0.65	1.08	3.99
0.75	0.69	0.65	1.09	4.04
1.00	0.69	0.65	1.10	5.05

As a first check (results not reported), we simulated the mean number of defaults for a single firm in two ways, to ensure that the discretization using the binomial tree does not result in a loss of accuracy. First, we used the base level parameters of the model as given above to build the binomial tree to a maturity of five years using monthly steps. We then compute the mean probability of default on the tree over five years. Then, we run a simulation using the equity process as the driver, and compute the underlying intensity at each stage. We sum up the intensity and compute the survival probability over the horizon of 5 years and use this to determine the mean probability of default. We then compare the mean default from this method to one based on Monte Carlo over the full possible state space of stock prices. The comparison confirms that the mean level of default is the same across both methods. Hence, simulating paths along the trees (a discrete set of prices) is not materially different from simulating stock price paths from geometric Brownian motions with jumps to default.

4.2 Varying equity correlation

Table 2 shows how the moments of the loss distribution vary when the correlation of equity returns is varied from zero to unity. The mean number of firms (out of 10 firms) defaulting in five years is 0.69. Mean default across the portfolio should be invariant to correlation across the portfolio, and we see that this is the case. We undertake a quick cross check of the simulated default rate to see if it lies within the ballpark of the theoretical expectation. Note that if the stock price were constant at \$50 for all five years, then the intensity for each month n would be $\xi_n = \exp(-0.5 + 0.1t)/50$, where $n = 1 \ldots 60$, and $t = (n-1)/12$. The total intensity for five years would be $I = \sum_{n=1}^{60}(\xi_n/12)$. Finally the probability of default over five years would be $\lambda = 1 - \exp(-I)$. Undertaking these calculations gives a mean value for the five-year default probability of a single firm of $\lambda = 0.0769$, and hence for ten identical firms, the mean number of defaults is 0.769. This is slightly higher than the range observed in Table 1, though in the same approximate vicinity; this is because in the simulation the stock price is positively skewed $(q > \frac{1}{2})$, spending more time higher than 50, resulting in a lower average default probability.

Table 3: Moments of the number of defaults when equity volatility is varied. The number of firms used in the credit basket is $n = 10$. The base level of parameters used in the simulation are: $S = 50$, $\sigma = 0.3$, $\alpha = -0.5$, $\beta = 1$, $\gamma = 0.1$, $r_f = 0.03$, $\rho = 0.4$. The simulated loss distributions are computed for a horizon of 5 years, and the simulation step is monthly.

Volatility (σ)	Mean	Variance	Skewness	Kurtosis
0.2	0.68	0.64	1.12	4.15
0.3	0.69	0.64	1.09	4.02
0.4	0.70	0.66	1.10	4.07
0.5	0.71	0.67	1.08	4.05
0.6	0.72	0.67	1.05	3.91
0.7	0.74	0.70	1.05	3.95
0.8	0.77	0.72	1.03	3.89

It is interesting to note that the moments of the loss distribution are not very sensitive to changes in the correlation assumption, though kurtosis increases with intensity correlation. Of course, the mean loss will not vary as it is an expectation. The other moments do not vary much either. This verifies the known feature of intensity based models where the loss distribution is less sensitive to changes in intensity correlations when defaults are independent once we condition on intensities. We will show later that injecting some conditional dependence makes loss distributions much more sensitive to correlation assumptions. We will also see that conditional (on intensities) dependence in defaults has little impact if intensities are uncorrelated. Overall, correlation in intensities impacts loss distributions more, though it requires some conditional dependence in defaults.

4.3 Varying equity volatility

Next, we fix the level of equity correlation to be $\rho = 0.4$, and examine how the loss distribution is impacted when the other parameters are varied. We begin by varying equity volatility, and the results are shown in Table 3.

Mean credit losses in the portfolio increase when equity volatility increases. There are two effects that drive this result. First, in the binomial tree, holding the risk free rate constant, if volatility is increased, the probability of an up move in the stock price (q) declines as σ increases. With increasing volatility, equity prices skew below the starting value, and result in above average default probability. This results in increasing mean credit losses. Second, the default function is convex in stock prices ($\xi \sim \frac{1}{S}$). This implies that as stock prices fall, default probabilities rise rapidly. On the other hand when stock prices rise, default probabilities drop, but not as fast. The effects of this convexity are exacerbated with increasing equity volatility. Hence, for both these reasons, mean losses increase with equity volatility. This result also has a connection with structural models of default, where increases in firm volatility (and consequently equity volatility) result in increases in default risk, and thus increase mean credit losses.

Table 4: Moments of the number of defaults when the default function parameters are varied. The number of firms used in the credit basket is $n = 10$. The base level of parameters used in the simulation are: $S = 50$, $\sigma = 0.3$, $\alpha = -0.5$, $\beta = 1$, $\gamma = 0.1$, $r_f = 0.03$, $\rho = 0.4$. The simulated loss distributions are computed for a horizon of 5 years, and the simulation step is monthly.

Panel A				
Level (α)	Mean	Variance	Skewness	Kurtosis
-0.5	0.68	0.64	1.09	4.01
0	1.08	0.97	0.80	3.45
0.5	1.67	1.40	0.57	3.12

Panel B				
Curvature (β)	Mean	Variance	Skewness	Kurtosis
0.5	3.78	2.35	0.17	2.85
1	0.69	0.65	1.10	4.07
2	0.01	0.01	8.62	78.04

Panel C				
Slope (γ)	Mean	Variance	Skewness	Kurtosis
-0.1	0.44	0.42	1.42	4.89
0	0.54	0.51	1.24	4.35
0.1	0.69	0.64	1.10	4.07
1	6.32	2.34	-0.16	2.83

We see that the variance of credit losses increases when equity volatility rises. This follows simply from the fact that when stock prices become more variable, so do hazard rates, as well as credit losses. The skewness and kurtosis of credit losses decline when equity volatility increases because at higher volatilities, outlier loss observations are no longer as extreme.

4.4 Varying default function parameters

The default function (see equation 7) contains three parameters, which we had previously identified with its shape as follows: the level parameter (α), the curvature parameter (β), and the slope parameter (γ). In Table 4 we present the moments of the loss distribution as these parameters are varied.

In Panel A of the table, we vary α. As α increases, the hazard rate increases, and the mean loss rate also increases as expected.

The change in the level of losses when the curvature parameter β increases is as expected – mean losses decline. The variance of loss also declines, and this results in increasing skewness and kurtosis of loss. The same effect is noticed with a decrease in the slope parameter γ.

5 Further Analysis

5.1 Intensity correlation and equity correlation

A common assumption for driving a system of correlated default in practice appears to be to use firms' asset correlations or equity correlations as proxies. Given that we have calibrated a firm's default function, we may use equity correlations to extract implications about hazard rate correlations.

Recall that for issuer i, the default function is given by $\xi_i(t) = \frac{\exp(\alpha_i+\gamma_i t)}{S_i^{\beta_i}}$. Hence, $\xi_i(t)$ or default intensity represents the instantaneous rate of default at time t, i.e. the forward rate for default. Using this functional form it is easy to write down the relationship between the covariance of intensities of any two firms and the covariance of their stock prices.

$$\begin{aligned} Cov(\xi_i, \xi_j) &= \frac{\partial \xi_i}{\partial S_i} \frac{\partial \xi_j}{\partial S_j} Cov(S_i, S_j) \\ &= \left[\frac{\beta_i \beta_j}{S_i S_j}\right] \xi_i \xi_j \; Cov(S_i, S_j) \end{aligned}$$

From this equation, the quantities of largest absolute magnitude tend to be the stock prices which reside in the denominator. (It may be that both stock prices are very small, in which case, intensity correlations will be high as expected). Usually though, the covariance of intensities will usually be scaled in magnitude below the covariance of stocks; this parallels the stylized fact that default correlations are lower than asset correlations (Lucas (1995)). Note too that the sign of the correlation of default intensities depends on the signs of β_i, β_j. These are almost always positive so that default correlations will be positive unless the two stocks are negatively correlated. The equation above shows that as intensity increases, correlation also increases, corresponding to the theoretical results of Zhou (2001). Therefore, our endogenous default functions has properties that conform to known stylized facts.

By varying t, we may also examine the time-dependence of conditional default correlation. To do so, we conduct the following simulation experiment. Using the parameters estimated in Table 1, we simulated stock prices S_{it} for each issuer for a given maturity t assuming that stock prices follow the usual geometric Brownian with which the model in this paper is consistent. We assumed that all six firms in Table 1 have pairwise equity return correlations of 0.40. Given the forward time t and S_{it}, we determine the corresponding intensity $\xi_{it}, \forall i$ using the default function. Our simulation comprised 10,000 joint draws of default intensity, from which we computed correlations. We repeated the exercise for forward maturities $t = \{1, 2, 3, 4, 5\}$ years. The resulting correlations matrices are portrayed in Table 5.

An examination of the correlations of intensities in Table 5 confirms the fact that the values are lower than the equity correlation of 40%. Further, the correlations are declining with horizon. Hence, as time progresses, forward values of intensity correlation decline, implying that in a credit portfolio, it is the early years in which the impact of correlated default is more severe. The implication is that credit diversification increases with horizon.

Note here that we are dealing explicitly with *forward* default intensities. As time proceeds, the forward intensities become less correlated. Forward intensities are different from the cumulated

Table 5: Correlation of forward default intensities at maturities t ranging from 1 to 5 years based on the calibrated parameters for the six firms in Table 1. The correlations of intensities are based on an underlying equity return correlation of 0.40 pairwise for all stocks. The patterns in the table remain the same even when this pairwise correlation is different.

	BARC	SUNW	AMZN	GM	PRF	NWAC
			$t = 1$ year			
BARC	1.000	0.350	0.333	-0.371	0.379	0.311
SUNW	0.350	1.000	0.322	-0.355	0.360	0.304
AMZN	0.333	0.322	1.000	-0.327	0.340	0.288
GM	-0.371	-0.355	-0.327	1.000	-0.381	-0.309
PRF	0.379	0.360	0.340	-0.381	1.000	0.342
NWAC	0.311	0.304	0.288	-0.309	0.342	1.000
			$t = 2$ years			
BARC	1.000	0.297	0.263	-0.304	0.309	0.245
SUNW	0.297	1.000	0.281	-0.281	0.311	0.270
AMZN	0.263	0.281	1.000	-0.244	0.270	0.230
GM	-0.304	-0.281	-0.244	1.000	-0.347	-0.225
PRF	0.309	0.311	0.270	-0.347	1.000	0.254
NWAC	0.245	0.270	0.230	-0.225	0.254	1.000
			$t = 3$ years			
BARC	1.000	0.250	0.220	-0.268	0.278	0.207
SUNW	0.250	1.000	0.210	-0.250	0.285	0.199
AMZN	0.220	0.210	1.000	-0.205	0.218	0.208
GM	-0.268	-0.250	-0.205	1.000	-0.331	-0.208
PRF	0.278	0.285	0.218	-0.331	1.000	0.252
NWAC	0.207	0.199	0.208	-0.208	0.252	1.000
			$t = 4$ years			
BARC	1.000	0.216	0.150	-0.250	0.291	0.115
SUNW	0.216	1.000	0.101	-0.169	0.207	0.124
AMZN	0.150	0.101	1.000	-0.164	0.214	0.083
GM	-0.250	-0.169	-0.164	1.000	-0.306	-0.117
PRF	0.291	0.207	0.214	-0.306	1.000	0.137
NWAC	0.115	0.436	0.083	-0.117	0.137	1.000
			$t = 5$ years			
BARC	1.000	0.100	0.228	-0.223	0.219	0.100
SUNW	0.100	1.000	0.098	-0.123	0.122	0.070
AMZN	0.228	0.098	1.000	-0.161	0.188	0.082
GM	-0.223	-0.123	-0.161	1.000	-0.281	-0.091
PRF	0.219	0.122	0.188	-0.281	1.000	0.118
NWAC	0.100	0.070	0.082	-0.091	0.118	1.000

total intensity. The correlation of cumulated total intensity will be somewhat higher than that of forward intensities because it is a time average.

5.2 Risk-neutral and statistical intensities

Using the calibrated defaultable binomial tree, we can compute the cumulative probability of default for any horizon. This is undertaken by summing up the probabilities of all paths that lead to a default up to the horizon in question. These are under the risk-neutral measure.

We plotted these term structures of cumulative default probabilities against those from `CreditGrades`, which are under the physical probability measure. Figures 4 and 5 show that the risk-neutral probabilities of default are higher than those under the physical measure, as required for risk averse investors. For firms with low default probability (our first 3 firms) the cumulative probability curve tends to be convex. This means that with time the probability of default is expected to increase, as conditionally, good firms revert to being of lesser quality, a stylized fact that is well recognized in the manner in which credit ratings are known to revert to the mean. Consistent with this mean-reversion logic, for high default risk firms (the latter three), the cumulative probability curve is concave.

Another important feature of the comparison between risk-neutral and statistical intensities is that the difference between the term-structures of intensity increases as default risk increases. This is consistent with increasing risk premia for default risk, which accounts for the difference between the curves.

5.3 Pricing nth to default contracts

In this section, we examine the valuation of n-th to default contracts. We take a simple case where there are 2 firms with identical parameters. For each firm we simulate a path of stock prices each month for one year. The stock price along the path is used to determine the default intensity for each month. By integrating the default intensity across all months, we obtain the total intensity for one year; this may be then used to simulate default. We generate 100,000 sample paths. For each path we simulate default in two ways:

1. *Conditional Independence*: Assuming that correlated default comes solely from correlations amongst default probabilities, and no further dependence in default after this conditioning. This is consistent with a Cox process or doubly stochastic process. We do this by drawing uniform random numbers without correlation at each point in the simulation path and compare these to the default probabilities to determine if default occurs or not.

2. *Conditional Dependence*: Drawing random normal numbers with correlation based on that of default probabilities, and then comparing the CDF of these numbers with the default probabilities to determine default occurrence. This approach is similar to a copula approach where default correlations are driven by correlated draws of random numbers for comparison with default probabilities. Effectively, in this approach, additional residual copula correlation is injected into default times, over and above that emanating from default

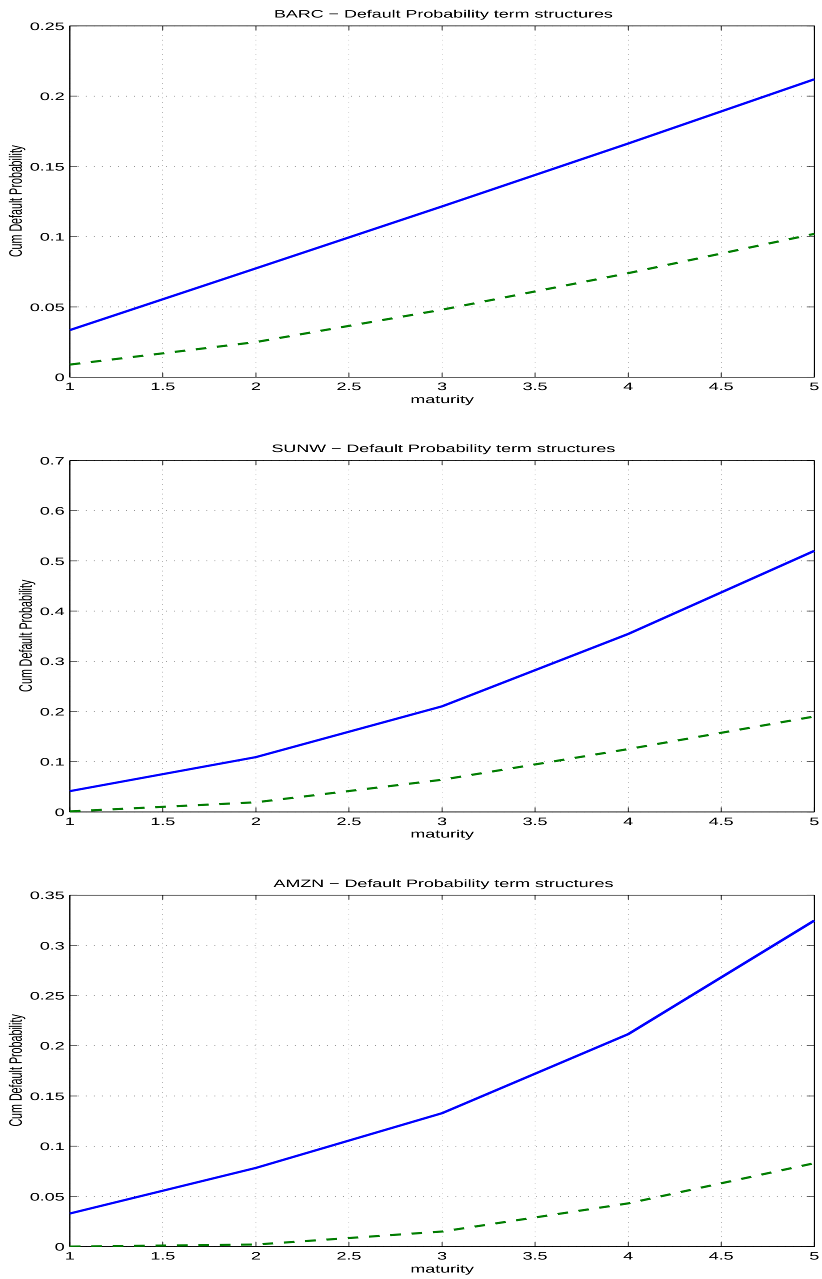

Figure 4: Cumulative default probabilities. The bold line presents the risk-neutral values and the dashed line those under the physical measure. The three firms here are (from top to bottom): BARC, SUNW, AMZN.

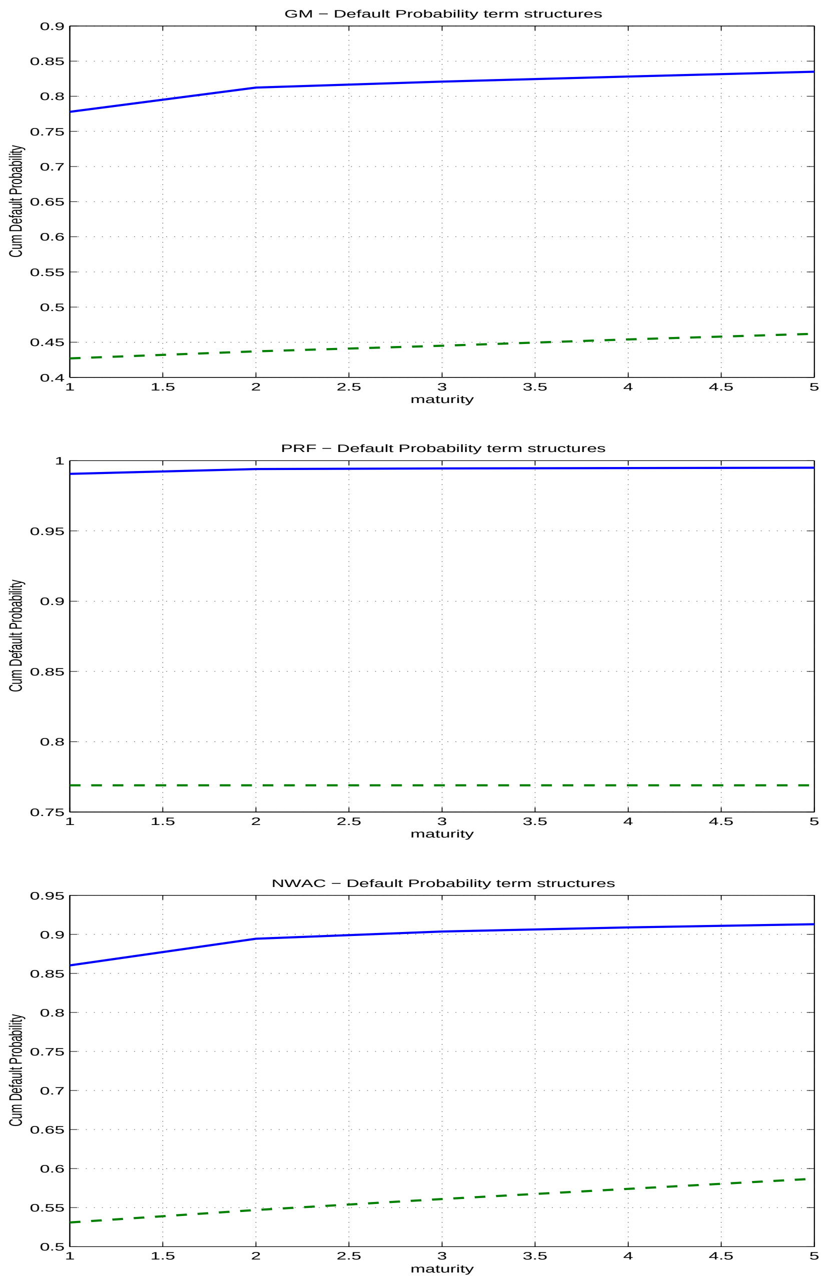

Figure 5: Cumulative default probabilities. The bold line presents the risk-neutral values and the dashed line those under the physical measure. The three firms here are (from top to bottom): GM, PRF, NWAC.

Table 6: Probability of triggering n-th to default contracts. The table presents the probability of a pay off on a first-to-default contract and a second-to-default contract. We assume there are 2 firms each with the following parameters: $S_0 = \$50$, $\sigma = 0.3$, $T = 1$, $\alpha = -0.5$, $\beta = 1$, $\gamma = 0.1$, $r_f = 0.03$. We simulated paths of stock prices and intensities for 1 year, with a period interval of 1 month ($h = 1/12$). The sum of intensities for the year is the sum of 12 monthly intensities along the simulation path divided by 12. We then converted this annual intensity (I) into a probability of default: $p = 1 - \exp(-I)$. We used two approaches to simulate default after determining the values of p for each firm: (a) assuming independent defaults after conditioning on p (i.e. respecting the doubly stochastic assumption), and (b) assuming defaults drawn from a Gaussian copula with correlation given by the average correlation of default probabilities. We report below the correlation of stock returns (Corr S), correlations of default probabilities (Corr PD), and the probabilities of first and second to default outcomes.

Corr S	Corr PD	1st def	2nd def
Panel A: Intensity Corr $= \rho$, Conditional Corr $= \rho$			
0.00	-0.0006	0.025660	0.000158
0.40	0.3952	0.024665	0.001228
0.80	0.7965	0.021124	0.005046
Panel B: Intensity Corr $= \rho$, Conditional Corr $= 0$			
0.00	-0.0015	0.025605	0.000139
0.40	0.3953	0.025753	0.000167
0.80	0.7971	0.025726	0.000162
Panel C: Intensity Corr $= 0$, Conditional Corr $= \rho$			
0.00	0.0005	0.025611	0.000177
0.40	-0.0004	0.026013	0.000178
0.80	0.0004	0.025556	0.000169

intensities. Das, Duffie, Kapadia and Saita (2007) find that residual copula correlation is of the magnitude of upto 5% in U.S. data.

We demonstrate the effect of these two kinds of dependence using a simulation under different settings for both the approaches. The results are presented in Table 6. We first allow for both types of correlation (results shown in Panel A). In Panel B, we suppress conditional dependence, only allowing for correlation in intensities. In Panel C of the table, we suppress intensity correlations, allowing only for conditional dependence with a copula.

From Panel A, we see that the probability of triggering a first to default contract declines with an increase in correlation. This is because at a correlation of zero, there is a greater chance for at least one default given that the firm's defaults are unrelated. If the correlation is high, then either both are likely to default or both are unlikely to default, making the probability of any one defaulting less likely. In contrast, under the second to default contract, the probability of a payoff naturally increases with an increase in correlation. This is an interesting aspect of these contracts - the former is inversely related to correlation whereas the latter is positively related to correlation. Most important, conditional dependence is an important driver of the value of these contracts.

From a comparison of Panels B and C with Panel A, we see that when we switch off either intensity correlation or conditional correlation, the probabilities of triggering either contract do not seem sensitive to correlation assumptions. This shows that while intensity correlations are important to account for, without conditional correlation, the contract prices would not display adequate sensitivity. However, more importantly, in practice, conditional copulas are used, assuming static intensities. We see that unless default intensities are dynamic, these contracts may not be very sensitive to correlation assumptions.

5.4 The shape of the CDS curve and portfolio losses

In this section, we examine how the slope of the spread curve impacts the pricing of n-th to default contracts. We repeat the analysis of the previous subsection with one change. The second firm's CDS curve slope is switched to negative by flipping the sign of the variable γ for the second firm. In all other respects the two firms remain identical. The results are presented in Table 7.

Since one firm now has a downward sloping spread curve, it also implies a downward sloping forward intensity curve. Thus, the probability of triggering both, first and second to default contracts will be reduced. The pattern of correlations and trigger probabilities however remains similar across Tables 6 and 7. Thus the slope of the CDS curve has relatively minor impact on correlated default.

Looking at nth to default contracts is useful as it points out how correlations matter. In the next section we look at a larger credit portfolio to gain a better understanding of which correlation (intensity or copula) may be more important.

5.5 The impact of conditional default distributions

Conditional on the level of default intensities, the events of default may be correlated via a wide selection of joint distributions. The approach in this paper can easily accommodate all sorts of variations. Here, we analyze twelve separate cases for the loss distribution using the six firms we calibrated in Table 1. Our simulation approach will be as follows. (i) Fixing the stock return correlation to be ρ, we generate a path of stock prices and default intensities using monthly (time step $h = 1/12$) steps for all 6 stocks to a horizon of 5 years (i.e. 60 monthly periods). (ii) We sum up the intensities and divide by 12 to get the total intensity for the period for each firm. (iii) We convert this intensity into the total default probability for the 5-year period. (iv) We draw a random number for each firm and compare the CDF (a number between zero and one) of these numbers to the default probability for each firm to determine if default is triggered or not. We maintain a count of how many firms default in each simulation run. Aggregating across all runs, we build up the frequency distribution of the number of losses.

The twelve cases we analyze have to do with step (iv) of the preceding paragraph. We draw

Table 7: Probability of triggering n-th to default contracts. The table presents the probability of a pay off on a first-to-default contract and a second-to-default contract. We assume there are 2 firms each with the following parameters: $S_0 = \$50$, $\sigma = 0.3$, $T = 1$, $\alpha = -0.5$, $\beta = 1$, $r_f = 0.03$. However, the first firm has $\gamma = 0.1$ and the second firm has $\gamma = -0.1$. We simulated paths of stock prices and intensities for 1 year, with a period interval of 1 month ($h = 1/12$). The sum of intensities for the year is the sum of 12 monthly intensities along the simulation path divided by 12. We then converted this annual intensity (I) into a probability of default: $p = 1 - \exp(-I)$. We used two approaches to simulate default after determining the values of p for each firm: (a) assuming independent defaults after conditioning on p (i.e. respecting the doubly stochastic assumption), and (b) assuming defaults drawn from a Gaussian copula with correlation given by the average correlation of default probabilities. We report below the correlation of stock returns (Corr S), correlations of default probabilities (Corr PD), and the probabilities of first and second to default outcomes.

Corr S	Corr PD	1st def	2nd def
Panel A: Intensity Corr = ρ, Conditional Corr = ρ			
0.00	0.0003	0.024390	0.000156
0.40	0.3950	0.023342	0.001048
0.80	0.7973	0.020019	0.004790
Panel B: Intensity Corr = ρ, Conditional Corr = 0			
0.00	0.0009	0.024306	0.000154
0.40	0.3945	0.024374	0.000132
0.80	0.7965	0.024404	0.000153
Panel C: Intensity Corr = 0, Conditional Corr = ρ			
0.00	-0.0000	0.024273	0.000158
0.40	0.0006	0.024263	0.000160
0.80	0.0001	0.024829	0.000152

random numbers in the following different ways using the student's T distribution. We use three correlation levels in the joint PD distribution $\{0, 0.4, 0.8\}$ and four degrees of freedom of the T distribution $\{5, 10, 20, 30\}$. In these resulting twelve cases, for correlation parameter ρ, we are generating stock returns with common correlation ρ, and are drawing conditional defaults with copula correlation also equal to ρ. Table 8 shows the loss distributions under these scenarios. The cases with both intensity correlation and positive conditional correlation are presented in Panel A of Table 8 and those with positive intensity correlation but no conditional correlation are explored in Panel B of the same table. Panel C shows the case with no intensity correlation and positive conditional correlation. The three panels enable us to examine the impact on loss distributions of allowing for different combinations of intensity and conditional correlation.

First, we note from Panel A that as correlation ρ increases, the tails of the loss distribution become thicker - there is a greater frequency of few losses as well as many losses. Hence correlation certainly matters in determining the riskiness of a credit portfolio.

Second, we find that as the degrees of freedom in the T distribution increase, the tails of the loss distribution are fatter. This is an artifact of taking the CDF of a Gaussian random variable under the T distribution, which results in a higher value than the CDF under the normal distribution (because the T distribution has fatter tails), thereby making it less likely to fall below the default probability value and trigger a default. This effect dissipates at the degrees of freedom increase, making the T distribution closer to a normal. Hence, the tails of the loss distribution get fatter when degrees of freedom increase. The mathematical niceties are less relevant than the economic interpretation which suggests that the loss distribution will depend also on the marginal distributions of loss for each individual credit in the portfolio.

Third, while we have demonstrated that credit portfolio risk is dependent on the level of correlation and the shape of the marginal loss distribution, a critical question to ask is: which correlation matters more, (a) the correlation of default intensities, or (b) the correlation of defaults conditional on intensities? To examine this question, we look at Panels B and C in Table 8 and compare them to Panel A. In Panel B, we only allow default intensities to have common correlation ρ, and set conditional correlation to be zero in all cases. Comparing Panels A and B, we see that the loss distributions have thinner tails in Panel B (as would be expected), but the thinning of the tails is not substantial. Hence, switching off conditional correlation does not seem to change the loss distributions substantially. In Panel C, the intensity correlations are set to zero in all cases, and only conditional correlation is set to ρ. Now, in comparison to Panels A and C, we see a dramatic thinning in the tails of the loss distribution. Hence, conditional correlation does not seem to impact the loss distribution as much as we might expect.

Overall, we may conclude the following from our analysis in this subsection. (i) The loss distribution is sensitive to default correlations, (ii) to marginal distributions in a copula, and (iii) is more sensitive to intensity correlations than conditional correlations. There is a substantial debate about whether doubly stochastic models of default (also known as Cox process models) are empirically justified (see the tests in Das, Duffie, Kapadia and Saita (2007)). Our analysis here suggests that residual conditional correlation matters in the presence of intensity correlation.

Table 8: Loss distributions under various correlation assumptions. We ran 10,000 paths of stock prices for all the six firms we calibrated in Table 1. These prices are generated from paths of stock returns drawn with correlation ρ. Given a generated path of intensities we compute the five year probability of default and use it to determine which firms defaulted on that path by drawing random numbers from a multivariate normal distribution with correlation ρ. We compared the CDF of these numbers to the five-year default probabilities. (This proxies the use of an induced multivariate T distribution through a Gaussian copula). The table shows the frequency distribution of losses from $n = 0$ defaults to $n = 6$ (all firms default). Thus, in Panel A, the intensity paths have correlation ρ *and* the conditional defaults are drawn with correlation ρ as well. In Panels B and C one of the correlations is set to zero.

Panel A: Non-zero Intensity correlation, non-zero conditional correlation								
ρ	d.o.f	$n=0$	$n=1$	$n=2$	$n=3$	$n=4$	$n=5$	$n=6$
0.0	5	16	435	2487	3352	2665	976	69
0.0	10	11	516	2484	3246	2681	988	74
0.0	20	25	491	2493	3220	2733	960	78
0.0	30	25	491	2514	3286	2638	963	83
0.4	5	63	653	2613	2861	2456	1212	142
0.4	10	91	696	2482	2739	2594	1220	178
0.4	20	109	712	2495	2714	2527	1228	215
0.4	30	119	733	2464	2679	2522	1272	211
0.8	5	119	784	2678	2347	2387	1504	181
0.8	10	195	777	2648	2358	2317	1432	273
0.8	20	225	827	2524	2270	2335	1482	337
0.8	30	247	780	2591	2232	2285	1507	358
Panel B: Non-zero Intensity correlation, zero conditional correlation								
ρ	d.o.f	$n=0$	$n=1$	$n=2$	$n=3$	$n=4$	$n=5$	$n=6$
0.0	5	17	443	2498	3308	2671	998	65
0.0	10	18	474	2433	3303	2725	989	58
0.0	20	29	444	2447	3271	2715	1011	83
0.0	30	34	513	2390	3312	2688	988	75
0.4	5	75	651	2550	2969	2546	1086	123
0.4	10	88	691	2532	2819	2551	1136	183
0.4	20	100	693	2490	2761	2541	1240	175
0.4	30	126	705	2499	2711	2546	1222	191
0.8	5	112	769	2688	2350	2446	1430	205
0.8	10	155	796	2694	2397	2245	1420	293
0.8	20	218	729	2674	2297	2324	1401	357
0.8	30	216	815	2659	2273	2289	1391	357
Panel C: Zero Intensity correlation, non-zero conditional correlation								
ρ	d.o.f	$n=0$	$n=1$	$n=2$	$n=3$	$n=4$	$n=5$	$n=6$
0.0	5	15	426	2567	3279	2698	960	55
0.0	10	16	437	2592	3278	2674	935	68
0.0	20	23	465	2483	3285	2714	953	77
0.0	30	23	500	2479	3255	2683	987	73
0.4	5	15	451	2598	3223	2710	949	54
0.4	10	17	438	2581	3250	2671	965	78
0.4	20	25	494	2493	3292	2684	918	94
0.4	30	31	463	2521	3292	2664	938	91
0.8	5	13	411	2528	3255	2821	928	44
0.8	10	25	440	2624	3242	2650	941	78
0.8	20	22	496	2510	3237	2697	962	76
0.8	30	18	465	2547	3262	2759	884	65

6 Extensions

In this section we briefly discuss some issues that arise in practical settings; we discuss how our framework is easily extendable to more complicated approaches.

6.1 Using factor models for default simulation

We note that in practice it is often impractical to model the correlation amongst stock returns especially when the credit portfolio comprises thousands of issuers. For example, in the case of CDO collateral comprising over a thousand issuers, we would end up dealing with a correlation matrix of enormous size. Drawing samples with correlation from a system this size can become numerically very taxing. Hence, practitioners resort to reducing the dimensionality of the system by using a parsimonious factor model comprising a few factors only.

By modeling stock returns using a factor model, we then only simulate the factors. The factor values generate stock returns which in turn translate into default intensities. A very simple factor model is the CAPM, which models every stock's return as being a function of the market index. The simulation would be conducted as follows. First, determine the relationship between the stock and the market index using a regression model. This also gives us the idiosyncratic risk of the stock. Second, simulate the index value. Third, embed this in the regression and simulate the idiosyncratic portion of the return to determine the total return for the stock. Fourth, use the simulated stock return (and resultant stock price) to generate the default intensity using the default function. By repeating this period after period, we obtain a time series of default intensities, and resultant default probabilities, which may be used to simulate default.

What is feasible in a one-factor model is just as feasible in a model with many factors. For example, we may wish t o use the Fama-French three-factor model instead. The important point of course, is that the dimension of the simulation will come just from a 3×3 correlation matrix, which is highly economical in simulation run time.

6.2 Multistage Default Simulation

There are two ways in which we may run the default simulations using the model in this paper: single-stage and multi-stage. In the previous examples we undertook default simulation in a single-stage approach by generating the sample path of default intensities for a year, integrating these to obtain the total intensity for the year and then using this to determine whether the issuer defaulted during the year or not.

Sometimes we may wish to proceed period by period (as in the case of cashflow CDOs), determining each month whether default has occurred or not. While simulating default times is also possible from the total intensity for the year, it ignores path-dependence in the underlying factors that drive intensity. Hence, in the presence of path-dependence, the correct approach to simulating default is to proceed period by period.

For example, if each period is a month ($h = 1/12$), then the multistage simulation would be as follows: (i) generate the stock price, (ii) compute the intensity (ξ) from the default function based on the stock price, (iii) convert the intensity into a default probability, i.e. $p = 1 - \exp(-\xi h)$, (iv) draw a uniform random variable u, and (v) assign default if $u \leq p$. Continue this way each month until maturity.

It is important to note that the moments of the default occurrence for a single issuer remain the same irrespective of whether the single-stage or multistage approach is used (this may be easily shown by running a simulation both ways). Path-dependence matters only for correlated default, in the setting where the doubly stochastic assumption is not imposed.

7 Conclusions

We implement a simple version of the model of Das and Sundaram (2007) to drive a system of default intensities from underlying stock return processes, similar to that used for equity options, extended to be calibrated to default swap spreads. We use this "forest of binomial trees" framework to explore features of correlated default in credit portfolios.

Various calibration and numerical analyses revealed some useful results. One, we found that correlations of forward default intensity tend to be lower than that of stock return correlations. Two, we explored two kinds of dependence: that emanating from intensity correlations, and further dependence coming from conditional correlation of defaults. Our results show that both kinds of dependence are important in determining credit loss distributions, and no one singly provides much impact. On balance, intensity correlations may be more important. Three, increases in equity volatility significantly drive the variation is credit portfolio losses. Four, the model endogenously generates higher credit risk premia for poor quality firms than for high quality ones. The model may be used in a factor framework and in either a single-stage or a multi-stage simulation setting.

References

Andersen, Lief., J. Sidenius, and S. Basu (2003). "All Your Hedges in One Basket," *Risk*, November, 67-72.

Burtschell, X., J. Gregory and J-P. Laurent (2005). "A Comparative Analysis of CDO Pricing Models," working paper, BNP-Paribas.

Cox, John., Stephen Ross, and Mark Rubinstein (1979). "Option Pricing: A Simplified Approach," *Journal of Financial Economics* 7, 229-263.

Carayannopoulos, Peter., and Madhu Kalimipalli (2003). "Convertible Bonds and Pricing Biases," *Journal of Fixed Income*, December.

Carr, Peter., and Liuren Wu (2005). "Stock Options and Credit Default Swaps: A Joint Framework for Valuation and Estimation," working paper, Baruch College, NY.

Das, Sanjiv., and Rangarajan Sundaram (2001). "A Discrete-Time Approach to Arbitrage-Free Pricing of Credit Derivatives," *Management Science* 46(1), 46-62.

Das, Sanjiv., Laurence Freed, Gary Geng, and Nikunj Kapadia (2006). "Correlated Default Risk," *Journal of Fixed Income*, Fall, 7-32.

Das, Sanjiv., Darrell Duffie, Nikunj Kapadia and Leandro Saita (2007). "Common Failings: How Corporate Defaults are Correlated," *Journal of Finance* 62, 93-117.

Das, Sanjiv., and Rangarajan Sundaram (2002). "A Simple Model for Pricing Derivative Securities with Equity, Interest-rate and Default Risk," Working paper, Santa Clara University.

Das, Sanjiv., and Rangarajan Sundaram (2007). "An Integrated Model for Pricing Hybrid Securities," *Management Science* 53, 1439-1451.

Das, Sanjiv., and Gary Geng (2004). "Correlated Default Processes: A Criterion-Based Copula Approach," *Journal of Investment Management* 2(2), 44-70.

Davis, Mark., and Fabian R. Lischka (1999). "Convertible Bonds with Market Risk and Credit Risk," working paper, Tokyo-Mitsubishi Intl, PLC.

Duffie, Darrel, J., and Nicole Garleanu (2001). "Risk and Valuation of Collateralized Debt Obligations," *Financial Analysts Journal* 57(1), (January-February), 41-59.

Duffie, Darrell., and David Lando (2001). "Term Structure of Credit Spreads with Incomplete Accounting Information" *Econometrica* 69, 633-664.

Duffie, Darrel, J., and Kenneth J. Singleton (1999). " Simulating Correlated Defaults," working paper, Stanford University, Graduate School of Business.

Egami, Masahiko., and Kian Esteghamat (2006). "An Approximation Method for Analysis and Valuation of Credit Correlation Derivatives," *Journal of Banking and Finance* 30, 341-364.

Giesecke, Kay, and Lisa Goldberg (2005). "A Top-Down Approach to Multi-Name Credit," Working paper, Stanford University.

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

Gregory, J., and J-P. Laurent (2004). "In the Core of Correlation," *Risk*, October, 87-91.

Heath, David., Robert A. Jarrow, and Andrew Morton (1990) "Bond Pricing and the Term Structure of Interest Rates: A Discrete Time Approximation," *Journal of Financial and Quantitative Analysis* 25(4), 419-440.

Helwege, Jean, and Christopher M. Turner, "The Slope of the Credit Yield Curve for Speculative-Grade Issuers," *Journal of Finance* 54(5), 1869-1884.

Hull, John., Izzy Nelken, and Alan White (2003). "Merton's Model, Credit Risk, and Volatility Skews," working paper, University of Toronto.

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

Linetsky, Vadim (2004). "Pricing Equity Derivatives Subject to Bankruptcy," forthcoming *Mathematical Finance*.

Longstaff, Francis, and Arvind Rajan (2006). "An Empirical Analysis of the Pricing of Collateralized Debt Obligations," forthcoming *Journal of Finance*.

Lucas, Douglas (1995). "Default Correlation and Credit Analysis," *Journal of Fixed Income*, March, 76-87.

Luciano, Elisa (2005). "Calibrating Risk-Neutral Default Correlation," working paper no 12/2005, International Center for Economic Research, University of Turin.

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

Mortensen, Allan (2005). "Semi-Analytical Valuation of Basket Credit Derivatives in Intensity-Based Models," working paper, Copenhagen Business School.

Schönbucher, Philipp. (2002). "A Tree Implementation of a Credit Spread Model for Credit Derivatives," *Journal of Computational Finance* 6(2), 1-38.

Schönbucher, Philipp. (2005). "Portfolio Losses and the Term Structure of Loss Transition Rates: A New Methodology for the Pricing of Portfolio Credit Derivatives," Working paper, ETH, Zurich.

Sidenius, Jakob., Vladimir Piterbarg, and Leif Andersen (2004). "A New Framework for Dynamic Credit Portfolio Loss Modeling," Working paper, Royal Bank of Scotland.

Zhou, Chunsheng (2001) "An Analysis of Default Correlation and Multiple Defaults," *Review of Financial Studies* 14, pp. 555-576.