

Common Failings: How Corporate Defaults Are Correlated

SANJIV R. DAS, DARRELL DUFFIE, NIKUNJ KAPADIA, and LEANDRO SAITA*

ABSTRACT

We test the doubly stochastic assumption under which firms' default times are correlated only as implied by the correlation of factors determining their default intensities. Using data on U.S. corporations from 1979 to 2004, this assumption is violated in the presence of contagion or "frailty" (unobservable explanatory variables that are correlated across firms). Our tests do not depend on the time-series properties of default intensities. The data do not support the joint hypothesis of well-specified default intensities and the doubly stochastic assumption. We find some evidence of default clustering exceeding that implied by the doubly stochastic model with the given intensities.

WHY DO CORPORATE DEFAULTS CLUSTER IN TIME? Several explanations have been explored. First, firms may be exposed to common or correlated risk factors whose co-movements cause correlated changes in conditional default probabilities. Second, the event of default by one firm may be "contagious," in that one such event may directly induce other corporate failures, as with the collapse of Penn Central Railway in 1970. Third, learning from default may generate default correlation. For example, to the extent that the defaults of Enron and World-Com revealed accounting irregularities that could be present in other firms, they may have had a direct impact on the conditional default probabilities of other firms.

Our primary objective is to examine whether cross-firm default correlation that is associated with observable factors determining conditional default probabilities (the first channel on its own) is sufficient to account for the degree of time clustering in defaults that we find in the data.

*Sanjiv Das is with Santa Clara University; Darrell Duffie is with Stanford University; Nikunj Kapadia is with the University of Massachusetts, Amherst; and Leandro Saita is with Lehman Brothers, New York. This research is supported by a fellowship grant from the Federal Deposit Insurance Corporation (FDIC). We received useful comments from participants at the FDIC Center for Financial Research conference, the Quantitative Work Alliance for Applied Finance, Education and Wisdom, San Francisco, Citigroup, the Quant Congress, Derivatives Securities Conference, Moodys-London Business School Credit Risk Conference, Federal Reserve Bank of New York, Bank of International Settlements and Deutsche Bundesbank workshop on Concentration Risk, Wilfrid Laurier University, National Bureau of Economic Research, the Q-Group, and the Western Finance Association Meeting. We are grateful to the editor and referees, as well as to Mark Flannery, Jean Helwege, Robert Jarrow, Edward Kane, Paul Kupiec, Dan Nuxoll, Neal Pearson, George Pennacchi, Louis Scott, Philip Shively, and Haluk Unal for their suggestions. We are also grateful to Moody's Investors Services, Gifford Fong Associates, and Professor Ed Altman for data and research support for this paper. The first author is grateful for the support of a Breetwor Fellowship.

Specifically, we test whether our data are consistent with the standard doubly stochastic model of default. Under this model, conditional on the paths of risk factors that determine all firms' default intensities, firm defaults are independent Poisson arrivals with these conditionally deterministic intensity paths. While this model is particularly convenient for computational and statistical purposes, its empirical relevance for default correlation has been unresolved in the literature. We develop a new test of the doubly stochastic assumption and apply it to default intensity and default time data for U.S. corporations over the period 1979 to 2004. The data do not support the joint hypothesis of well-specified default intensities and the doubly stochastic assumption. That is, we find evidence of default clustering beyond that predicted by the doubly stochastic model and our data.

Understanding how corporate defaults are correlated is particularly important for the risk management of portfolios of corporate debt. For example, to back the performance of their loan portfolios, banks retain capital at levels designed to withstand default clustering at extremely high confidence levels, such as 99.9%. Some banks determine their capital requirements on the basis of models in which default correlation is assumed to be captured by common risk factors determining conditional default probabilities, as in Gordy (2003) and Vasicek (1987). (Note that banks do attempt to capture the effects of contagion that arise from parent subsidiary and other direct contractual links.) If defaults are more heavily clustered in time than envisioned in these default risk models, then significantly greater capital might be required in order to survive default losses, especially at high confidence levels. An understanding of the sources and degree of default clustering is also crucial for the rating and risk analysis of structured credit products, such as collateralized debt obligations (CDOs) and options on portfolios of default swaps, that are exposed to correlated default. This is especially true given the rapid growth in these markets. For example, the Bank of International Settlements reports that synthetic CDO volumes reached \$673 billion in 2004.¹

While there is some empirical evidence regarding the average default correlation (see Akhvein, Kocagil, and Neugebauer (2005), Lucas (1995), and deServigny and Renault (2002)) and correlated changes in corporate default probabilities (Das et al. (2006)), there is relatively little evidence regarding the presence of clustered defaults. In particular, there is no extant work on whether the degree of default clustering in the data can be reasonably captured by doubly stochastic models. Collin-Dufresne, Goldstein, and Helwege (2003) and Zhang (2004) find that default events are associated with significant increases in the credit spreads of other firms, consistent with default clustering in excess of that suggested by the doubly stochastic model, or at least a failure of the doubly stochastic model under risk-neutral probabilities. This suggests that their findings may be due to default-induced increases in the conditional default probabilities of other firms, or to default-induced increases in the default

¹ Data are provided in the BIS Annual Report, 2005, and mention cash CDO volumes of \$163 billion.

risk premia² of other firms, as argued by Kusuoka (1999). That is, both effects could be at play.

Explicitly considering a failure of the doubly stochastic hypothesis, Collin-Dufresne, Goldstein, and Helwege (2003), Giesecke (2004), Jarrow and Yu (2001), and Schönbucher (2003) explore learning-from-default interpretations, based on the statistical modeling of frailty, where default intensities include the expected effect of unobservable covariates. In a frailty setting, the arrival of a default causes (via Bayes's Rule) a jump in the conditional distribution of hidden covariates, and therefore a jump in the conditional default probabilities of any other firms whose default intensities depend on the same unobservable covariates. For example, the collapses of Enron and WorldCom could have caused a sudden reduction in the perceived precision of accounting leverage measures of other firms. Indeed, Yu (2005) finds empirical evidence that, other things equal, a reduction in the measured precision of accounting variables is associated with a widening of credit spreads. Lang and Stulz (1992) explore evidence of default contagion in equity prices.

In theory, banks and other managers of credit portfolios could extend the doubly stochastic model if it were found to be seriously deficient. In practice, however, few methods used to measure loan portfolio credit risk allow for contagion or frailty. For example, when applied in practice, the Merton (1974) model and its variants imply that default correlation is captured by co-movement in the observable default covariates (primarily leverage, normalized for volatility) that determine conditional default probabilities.³ Ratings-based transition models have sometimes been applied to the task of credit portfolio risk management, again based on the doubly stochastic assumption that credit rating transitions intensities are based on commonly observable covariates.

The doubly stochastic property, sometimes referred to as “conditional independence,” also underlies the standard econometric duration models used for event forecasting, including default prediction models, such as those of Couderc and Renault (2004), Shumway (2001), and Duffie, Saita, and Wang (2005). This property implies that the likelihood function that is to be maximized when estimating the coefficients of an intensity model can be expressed as the product of the covariate-conditional likelihood functions of the firms' default-survival events in the data. One of our objectives is to provide a tool with which to verify whether this tractability is achieved at the expense of mis-specification associated with a failure of the doubly stochastic property.

Before describing our data, methods, and results in detail, we offer a brief synopsis. Our default intensity estimates are from Duffie et al. (2005) and

² Collin-Dufresne, Goldstein, and Hugonnier (2004) provide a simple method for incorporating the pricing impact of failure, under risk-neutral probabilities, of the doubly stochastic hypothesis. Other theoretical work on the impact of contagion on default pricing includes that of Cathcart and Jahel (2002), Davis and Lo (2001), Giesecke (2004), Jarrow, Lando, and Yu (2005), Kusuoka (1999), Schönbucher and Schubert (2001), Terentyev (2004), Yu (2003), and Zhou (2001).

³ Das et al. (2006) report that leverage and volatility are the two largest factors empirically explaining covariation in conditional default probabilities.

are based on two firm-specific covariates (distance to default and the trailing 1-year stock return), and two macro-covariates (the current 3-month Treasury rate and the trailing 1-year Standard and Poors 500 return). The data cover the period January 1979 to October 2004. Default times are correlated in this model both through correlated changes across firm-level covariates as well as through common dependence of default intensities on the two macro-covariates. The default-time data come from Moody's (and are slightly augmented as needed with information from Compustat and Bloomberg). The firm-specific covariates are based on data from Compustat and CRSP. We describe the data further in Section II. After excluding financial firms and dropping firms for which we have missing data matched across the data sources, our results cover 2,770 firms, 495 defaults, and 392,404 firm-months. The out-of-sample default intensities provide default prediction accuracy ratios averaging 88% during 1993 to 2004, exceeding those of any other available model. Broadly speaking, based on these default intensity data, we reject the joint hypothesis of correctly measured default intensities and the doubly stochastic property.

We exploit the following new result, developed in Section I. Consider a change of time scale under which the passage of one unit of "new time" coincides with a period of calendar time over which the cumulative total of all firms' default intensities increases by one unit. This is the calendar time period that, at current intensities, would include one default in expectation. Under the doubly stochastic assumption and under this new time scale, the cumulative number of defaults to date defines a standard (constant mean arrival rate) Poisson process. For example, with successive time periods each lasting for some fixed amount c of new time (corresponding to calendar periods that each include an accumulated total default intensity, across all firms, of c), the number of defaults in successive time intervals (X_1 defaults in the first interval lasting for c units, X_2 defaults in the second interval, and so on) are independent Poisson distributed random variables with mean c . This time-changed Poisson process is the basis of most of our tests, outlined as follows:

1. We apply a Fisher dispersion test for consistency of the empirical distribution of the numbers X_1, \dots, X_k, \dots of defaults in successive time bins of a given accumulated intensity c , with the theoretical Poisson distribution of mean c implied by the doubly stochastic model. The null hypothesis that defaults arrive according to a time-changed Poisson process is rejected at traditional confidence levels for all of the bin sizes that we study (2, 4, 6, 8, and 10).
2. We test whether the mean of the upper quartile of our sample X_1, X_2, \dots, X_K of numbers of defaults in successive time bins of a given size c is significantly larger than the mean of the upper quartile of a sample of like size drawn independently from the Poisson distribution with parameter c . An analogous test is based on the median of the upper quartile. These tests are designed to detect default clustering in excess of that implied by the default intensities and the doubly stochastic assumption. We also extend this test so as to simultaneously treat a number of bin

- sizes. The null is rejected at traditional confidence levels at bin sizes 2, 4, and 10, and is rejected in a joint test covering all bins. That is, at least insofar as this test implies, the data suggest excess clustering of defaults.
3. Taking another perspective, some of our tests are based on the fact that, in the new time scale, the inter-arrival times of default are independent and identically distributed exponential variables with parameter 1. We provide the results of a test due to Prahl (1999) for clustering of default arrival times (in our new time scale) in excess of that associated with a Poisson process. The null is rejected, which again provides evidence of clustering of defaults in excess of that suggested by the assumption that default correlation is captured by co-movement of the default covariates used for intensity estimation.
 4. Fixing the size c of time bins, we test for serial correlation of X_1, X_2, \dots by fitting an autoregressive model. The presence of serial correlation would imply a failure of the independent-increments property of Poisson processes, and, if the serial correlation were positive, could lead to default correlation in excess of that associated with the doubly stochastic assumption. The null is rejected in favor of positive serial correlation for all bin sizes except $c = 2$.

Because these tests do not depend on the joint probability distribution of the firms' default intensity processes, including their correlation structure, they allow for both generality and robustness. We find that the data are broadly consistent with a rejection at standard confidence intervals of the joint hypothesis of correctly specified intensities and the doubly stochastic hypothesis.

Such rejection could be due to mis-specification associated with missing covariates. For example, if the true default intensities depend on macroeconomic variables that are not used to estimate the *measured* intensities, then even after the change of time scale based on the measured intensities, the default times could be correlated. For instance, if the true default intensities decline with increasing gross domestic product (GDP) growth, even after controlling for the other covariates, then periods of low GDP growth would induce more clustering of defaults than would be predicted by the measured default intensities. Indeed, we find mild evidence that U.S. industrial production (IP) growth is a missing covariate. Even reestimating intensities after including this covariate, however, we continue to reject the nulls associated with the above tests (albeit at slightly larger p -values). Nonetheless, it remains possible that missing covariates, rather than a failure of the doubly stochastic property, could be responsible for some of the poor fit of the joint hypothesis that we test.

In order to gauge the degree of default correlation that is not captured by correlations among estimated default intensities, we calibrate a version of the intensity-conditional copula model of Schönbucher and Schubert (2001). The associated intensity-conditional Gaussian copula correlation parameter is a measure of the amount of additional default correlation that must be added, on top of the default correlation already implied by co-movement in default intensities, in order to match the degree of default clustering observed in the data. This

estimated incremental copula correlation ranges from 1% to 4% depending on the length of time window used. To place these estimates in perspective, Akhavan et al. (2005) estimate a Gaussian copula correlation parameter of approximately 19.7% within sectors and 14.4% across sectors, by calibration to empirical default correlations, that is, before “removing” the correlation associated with covariance in default intensities. Although this is a rough comparison, it indicates that default intensity correlation accounts for a large fraction, but not all, of the default correlation.

The rest of the paper comprises the following. In Section I, we derive the property that the total default arrival process is a Poisson process with constant intensity under a new time scale measured in units of the cumulative aggregate default intensity to date. This provides our testable implications. Section II describes our data. Section III presents various tests of the doubly stochastic hypothesis. Section IV estimates the degree of residual default correlation, above that implied by covariation in intensities, in terms of the incremental Gaussian copula correlation. Section V.A addresses the presence of serial independence of increments of the time-changed process governing default arrivals. In Section V.B, we test our default intensity data for missing macroeconomic covariates, and examine whether these may be responsible for the rejection of the doubly stochastic hypothesis. Section VI concludes.

I. Time Rescaling for Poisson Defaults

In this section, we define the doubly stochastic default property that rules out default correlation beyond that implied by correlated default intensities, and we provide testable implications of this property.

We start by fixing a probability space (Ω, \mathcal{F}, P) and an observer’s information filtration $\{\mathcal{F}_t : t \geq 0\}$ satisfying the usual conditions. This and other standard technical definitions that we rely on may be found in Protter (2003). We suppose that, for each firm i , $i \in 1, \dots, n$, default occurs at the first jump time τ_i of a nonexplosive counting process N_i with stochastic intensity process λ_i . (Here, N_i is (\mathcal{F}_t) -adapted and λ_i is (\mathcal{F}_t) -predictable.)

The key question at hand is whether the joint distribution of, and in particular any correlation among, the default times τ_1, \dots, τ_n is determined by the joint distribution of the intensities. Violation of this assumption means, in essence, that even after conditioning on the paths of the default intensities $\lambda_1, \dots, \lambda_n$ of all firms, the default times can be correlated.

A standard version of the assumption that default correlation is captured by co-movement in default intensities is the assumption that the multidimensional counting process $N = (N_1, \dots, N_n)$ is doubly stochastic. That is, conditional on the path $\{\lambda_t = (\lambda_{1t}, \dots, \lambda_{nt}) : t \geq 0\}$ of all intensity processes, as well as the information \mathcal{F}_T available at any given stopping time T , the counting processes $\hat{N}_1, \dots, \hat{N}_n$ defined by $\hat{N}_i(u) = N_i(u + T)$ are independent Poisson processes with respective (conditionally deterministic) intensities $\hat{\lambda}_1, \dots, \hat{\lambda}_n$ defined by $\hat{\lambda}_i(u) = \lambda_i(u + T)$. In this case, we also say that (τ_1, \dots, τ_n) is doubly stochastic with intensity $(\lambda_1, \dots, \lambda_n)$. In particular, the doubly stochastic assumption implies that the default times τ_1, \dots, τ_n are independent given the intensities.

We test the following key implication of the doubly stochastic assumption.

PROPOSITION: Suppose that (τ_1, \dots, τ_n) is doubly stochastic with intensity $(\lambda_1, \dots, \lambda_n)$. Let $K(t) = \#\{i : \tau_i \leq t\}$ be the cumulative number of defaults by t , and let $U(t) = \int_0^t \sum_{i=1}^n \lambda_i(u) \mathbf{1}_{\{\tau_i > u\}} du$ be the cumulative aggregate intensity of surviving firms to time t . Then $J = \{J(s) = K(U^{-1}(s)) : s \geq 0\}$ is a Poisson process with rate parameter 1.

Proof: Let $S_0 = 0$ and $S_j = \inf\{s : J(s) > J(S_{j-1})\}$ be the jump times, in the new time scale, of J . By Billingsley (1986), Theorem 23.1, it suffices to show that the inter-jump times $\{Z_j = S_j - S_{j-1} : j \geq 1\}$ are i.i.d. exponential with parameter 1. Let $T(j) = \inf\{t : K(t) \geq j\}$. By construction,

$$Z_j = \int_{T_{j-1}}^{T_j} \sum_{i=1}^n \lambda_i(u) \mathbf{1}_{\{\tau_i > u\}} du.$$

By the doubly stochastic assumption, given $\{\lambda_t = (\lambda_{1t}, \dots, \lambda_{nt}) : t \geq 0\}$ and \mathcal{F}_{T_j} , we know that $\tilde{N}_{j+1} = \{\tilde{N}(u) = \sum_{i=1}^n N_i(u + T_j) \mathbf{1}_{\{\tau_i > T_j\}} u \geq T_j\}$ is a sum of independent Poisson processes, and therefore is itself a Poisson process with intensity $\tilde{\lambda}_{j+1}(u) = \sum_{i=1}^n \lambda_i(u + T_j) \mathbf{1}_{\{\tau_i > T_j\}}$. Thus, Z_{j+1} is exponential with parameter 1.

In order to check the independence of Z_1, Z_2, \dots , consider any integer $k > 1$ and any bounded Borel functions f_1, \dots, f_k . By the doubly stochastic property and the law of iterated expectations applied recursively,

$$\begin{aligned} & E[f_1(Z_1) f(Z_2) \cdots f_{k-1}(Z_{k-1}) f_k(Z_k)] \\ &= E[f_1(Z_1) f(Z_2) \cdots f_{k-1}(Z_{k-1}) E[f_k(Z_k) | \lambda, \mathcal{F}_{T_{k-1}}]] \\ &= E[f_1(Z_1) f(Z_2) \cdots f_{k-1}(Z_{k-1})] \int_0^\infty f_k(z) e^{-z} dz \\ &\vdots \\ &= \prod_{i=1}^k \int_0^\infty f_i(z) e^{-z} dz. \end{aligned}$$

Thus, Z_1, Z_2, \dots are indeed independent, and J is a Poisson process with parameter 1, completing the proof.

Using this result, some of the properties of the doubly stochastic assumption that we test are based on the following characterization.

COROLLARY: Under the conditions of the proposition, for any $c > 0$, the successive numbers of defaults per bin,

$$J(c), J(2c) - J(c), J(3c) - J(2c), \dots,$$

are i.i.d. Poisson distributed with parameter c .

That is, by dividing our sample period into non-overlapping time “bins” that each contain an equal cumulative aggregate default intensity of c , we can test the doubly stochastic assumption by testing whether the numbers of defaults

in the successive bins are independent Poisson random variables with common parameter c . Other tests based on the implications of the Proposition will also be applied.

II. Data

The default-intensity data used in this paper are from Duffie, Saita, and Wang (2005),⁴ which estimates the default intensity of firm i at time t according to

$$\lambda_i(t) = e^{\beta_0 + \beta_1 X_{i1}(t) + \beta_2 X_{i2}(t) + \gamma_1 Y_1(t) + \gamma_2 Y_2(t)}, \quad (1)$$

where

$X_{i1}(t)$ is the distance to default of firm i , an estimate of the number of standard deviations by which the assets of the firm exceed a measure of liabilities.⁵ $X_{i2}(t)$ is the trailing 1-year stock return of firm i , a covariate shown by Shumway (2001) to provide significant explanatory power.

$Y_1(t)$ is the U.S. 3-month Treasury bill rate.

$Y_2(t)$ is the trailing 1-year return of the Standard and Poors 500 stock index.

We obtain data on corporate defaults and bankruptcies from two sources, namely, Moody's Default Risk Service and CRSP/Compustat. Moody's Default Risk Service provides detailed issue and issuer information on rating, default, and bankruptcy date and type (e.g., distressed exchange, missed interest payment, and so on), tracking 34,984 firms as of 1938. CRSP/Compustat provides reasons for deletion and year and month of deletion (data items AFTNT35, AFTNT34, and AFTNT33, respectively). Firm-specific financial data come from the CRSP/Compustat database. Stock prices are from CRSP's monthly file. We obtain short-term and long-term debt from Compustat's annual (data items DATA5, DATA9, and DATA34) and quarterly files (DATA45, DATA49, and DATA51), respectively. We construct the S&P 500 index trailing 1-year returns

⁴ The initial version of this paper was based instead on intensities derived from a smaller data set of default probabilities ("PDs") that were developed by Moody's Investor Services, as described in Sobehart et al. (2000).

⁵ Distance to default, the sole relevant default covariate in the model proposed by Merton (1974), is the number of standard deviations of annual asset growth by which assets exceed a measure of book liabilities. In order to estimate distance to default, DTD , the initial asset value, A_t , is taken to be the sum of S_t (end-of-month stock price times number of shares outstanding, from the CRSP database) and L_t (the sum of short-term debt and one-half long-term debt, from Compustat). The risk-free interest rate, r_t , is taken to be the 3-month T-bill rate. One solves for the asset value A_t and asset volatility σ_A by iteratively applying the equations:

$$A_t = S_t \Phi(d_1) - L_t e^{-r_t} \Phi(d_2) \\ \sigma_A = \text{sdev}(\ln(A_t) - \ln(A_{t-1})), \quad (6)$$

where Φ is the standard normal cumulative distribution function, and d_1 and d_2 are defined by

$$d_1 = \frac{\ln\left(\frac{A_t}{L_t}\right) + \left(r_t + \frac{1}{2}\sigma_A^2\right)}{\sigma_A},$$

$$d_2 = d_1 - \sigma_A.$$

Bharath and Shumway (2005) show that the estimated default intensity is relatively robust to various alternative approaches to estimating distance to default.

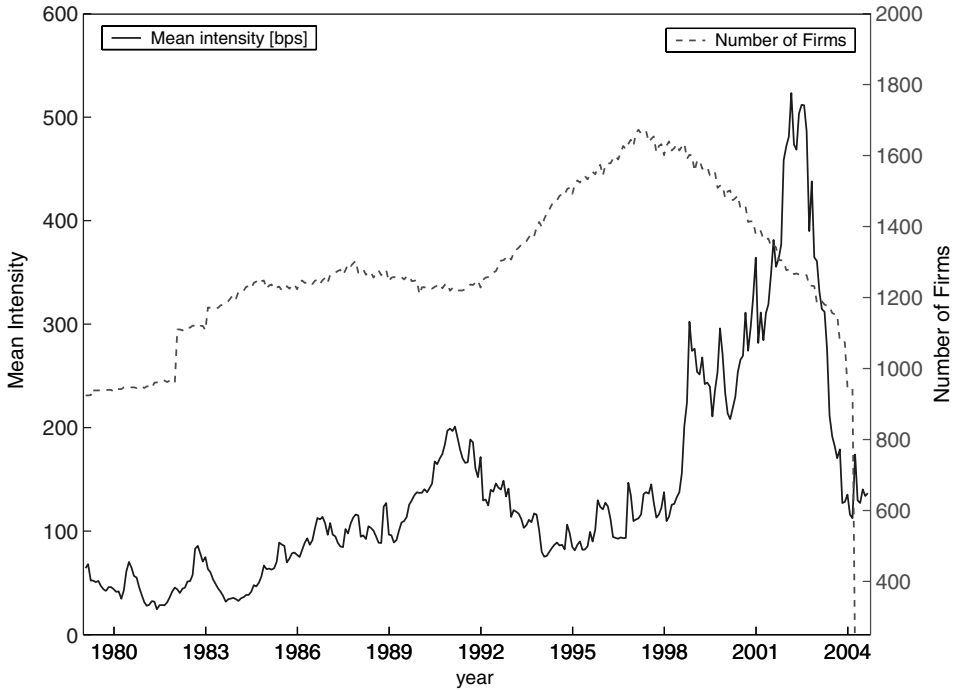


Figure 1. Firms and intensities. Cross-sectional sample mean of annualized default intensities and the number of firms covered, 1979 to 2004.

from monthly CRSP data. Treasury rates come from the website of the U.S. Federal Reserve Board of Governors. Included firms are those in Moody's "Industrial" category sector for which we have a common firm identifier for the Moody's, CRSP, and Compustat databases. This includes essentially all matchable U.S.-listed nonfinancial nonutility firms. We restrict attention to firms for which we have at least 6 months of monthly data both in CRSP and Compustat. Since Compustat provides only quarterly and yearly data, for each month we take debt to be the value provided for the corresponding quarter.

Using the selection procedure above, our sample consists of 2,770 firms, covering 392,404 firm-months of monthly data for the period January 1979 to October 2004. Our data set includes 495 defaults. The coefficients β_0 , β_1 , β_2 , γ_1 , and γ_2 , are estimated by full maximum likelihood, as detailed in Duffie et al. (2005).

Figure 1 shows that the cross-sectional mean of estimated default intensities increases markedly with the U.S. recession of 2000 to 2001. Figure 2 illustrates the number of defaults over this period on a month-by-month basis, (ranging from 0 to a maximum of 12), as well as a plot of the total across firms of the estimated monthly default intensities. If the default intensities are correctly measured, then the number of defaults in a given month is a random variable whose conditional mean, given the total intensity path, is the average of the total intensity path for the month.

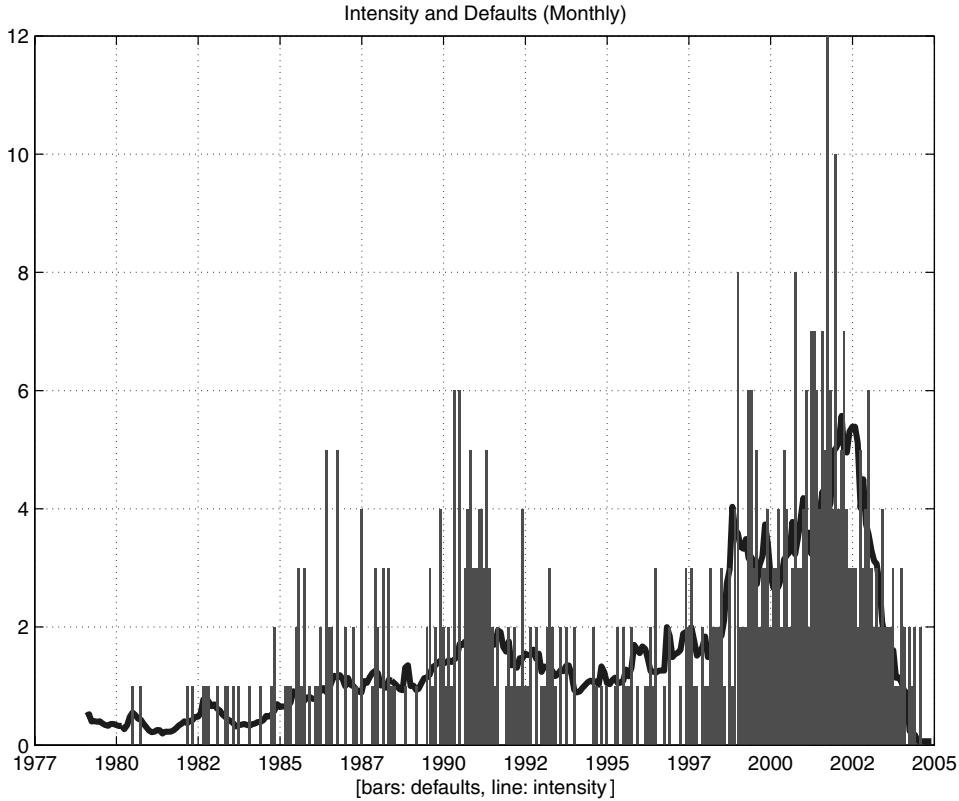


Figure 2. Intensities and defaults. Aggregate (across firms) of monthly default intensities and number of defaults by month, from 1979 to 2004. The vertical bars represent the number of defaults, and the line depicts the intensities.

III. Goodness-of-Fit Tests

Having estimated the annualized default intensity λ_{it} for each firm i and each date t (with λ_{it} taken to be constant within months), and letting $\tau(i)$ denote the default time of firm i , we let $U(t) = \int_0^t \sum_{i=1}^n \lambda_{is} \mathbf{1}_{\{\tau(i) > s\}} ds$ be the total accumulative default intensity of all surviving firms. In order to obtain time bins that each contain c units of accumulative default intensity, we construct calendar times t_0, t_1, t_2, \dots such that $t_0 = 0$ and $U(t_i) - U(t_{i-1}) = c$. We then let $X_k = \sum_{i=1}^n \mathbf{1}_{\{t_k \leq \tau(i) < t_{k+1}\}}$ be the number of defaults in the k^{th} time bin. Figure 3 illustrates the time bins of size $c = 8$ over the years 1995 to 2001.

Table I presents a comparison of the empirical and theoretical moments of the distribution of defaults per bin, for each of several bin sizes.⁶ The actual bin sizes vary slightly from the integer bin sizes shown because of granularity in

⁶ Under the Poisson distribution, $P(X_i = k) = \frac{e^{-c} c^k}{k!}$. The associated moments of X_k are a mean and variance of c , a skewness of $c^{-0.5}$, and a kurtosis of $3 + c^{-1}$.

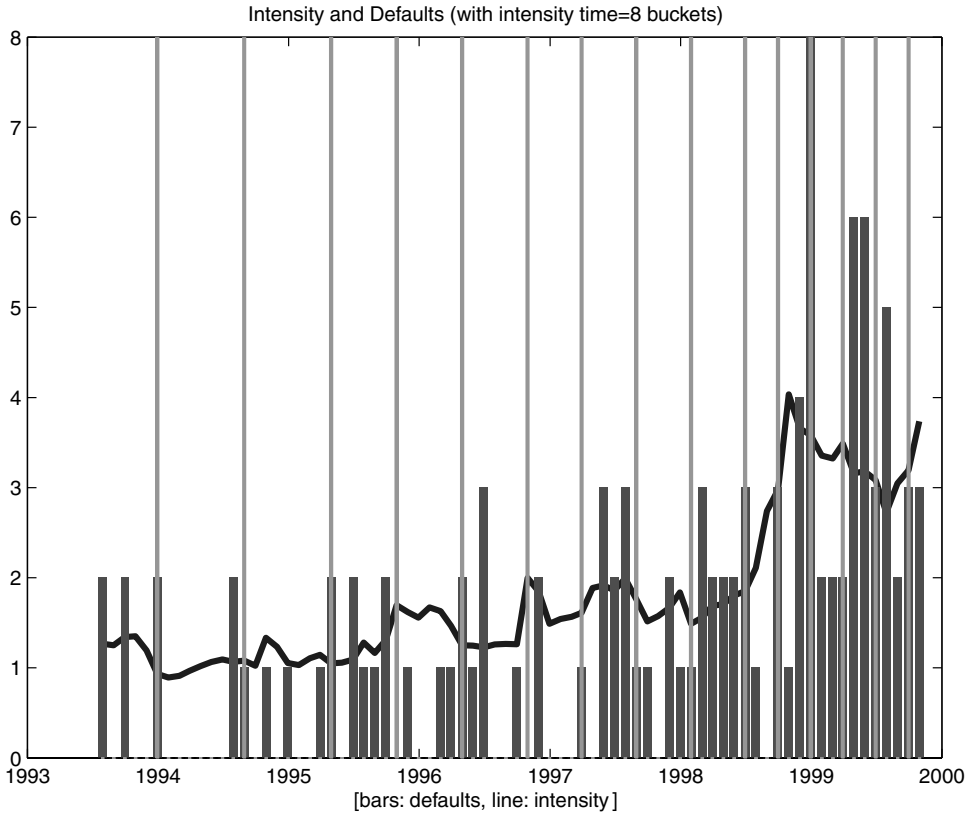


Figure 3. Time rescaled intensity bins. Aggregate intensities and defaults by month, 1994–2000, with time bin delimiters marked for intervals that include a total accumulated default intensity of $c = 8$ per bin. The vertical bars represent the number of defaults, and the line depicts the intensities.

the construction of the binning times t_1, t_2, \dots . The approximate match between a bin size and the associated sample mean $(X_1 + \dots + X_K)/K$ of the number of defaults per bin offers some confirmation that the underlying default intensity data are reasonably well estimated. However, this is somewhat expected given the within-sample nature of the estimates. To place this issue in context, the total number of defaults that is expected conditional on the paths of all default intensities is 470.6, whereas the actual number of defaults is 495. For larger bin sizes, Table I shows that the empirical variances are larger than their theoretical counterparts under the null of a correct doubly stochastic default intensity model.

Figures 4 and 5 present the observed default frequency distribution and the associated theoretical Poisson distribution for bin sizes 2 and 8, respectively. For bins of sizes larger than 4, there is a tendency for multimodality (multiple peaks), as opposed to the unimodal theoretical Poisson distribution associated with the hypothesis of doubly stochastic defaults. To the extent that defaults

Table I
Empirical and Theoretical Moments

This table presents a comparison of empirical and theoretical moments for the distribution of defaults per bin. The number K of bin observations is shown in parentheses under the bin size. The upper-row moments are those of the theoretical Poisson distribution under the doubly stochastic hypothesis; the lower-row moments are the empirical counterparts.

Bin Size	Mean	Variance	Skewness	Kurtosis
2	2.04	2.04	0.70	3.49
(230)	2.12	2.92	1.30	6.20
4	4.04	4.04	0.50	3.25
(116)	4.20	5.83	0.44	2.79
6	6.04	6.04	0.41	3.17
(77)	6.25	10.37	0.62	3.16
8	8.04	8.04	0.35	3.12
(58)	8.33	14.93	0.41	2.59
10	10.03	10.03	0.32	3.10
(46)	10.39	20.07	0.02	2.24

depend on unobservable covariates, or at least on relevant covariates that are not included whether observable or not, violations of the Poisson distribution would tend to be larger for larger bin sizes because of the time necessary to build up a significant incremental impact of the missing covariates on the probability distribution of the number of defaults per bin.

A. Fisher's Dispersion Test

Our first goodness-of-fit test of the hypothesis of correctly measured default intensities and the doubly stochastic property is Fisher's dispersion test of the agreement of the empirical distribution of defaults per bin, for a given bin size c , to the theoretical Poisson distribution with parameter c .

Fixing the bin size c , a simple test of the null hypothesis that X_1, \dots, X_K are independent Poisson distributed variables with mean parameter c is Fisher's dispersion test (Cochran (1954)). Under this null,

$$W = \sum_{i=1}^K \frac{(X_i - c)^2}{c} \quad (2)$$

is distributed as a χ^2 random variable with $K - 1$ degrees of freedom. An outcome for W that is large relative to a χ^2 random variable of the associated number of degrees of freedom would generate a small p -value, meaning a surprisingly large amount of clustering if the null hypothesis of doubly stochastic default (and correctly specified conditional default probabilities) applies. The p -values shown in Table II indicate that at standard confidence levels such as 95%, for all bin sizes we reject this null hypothesis.

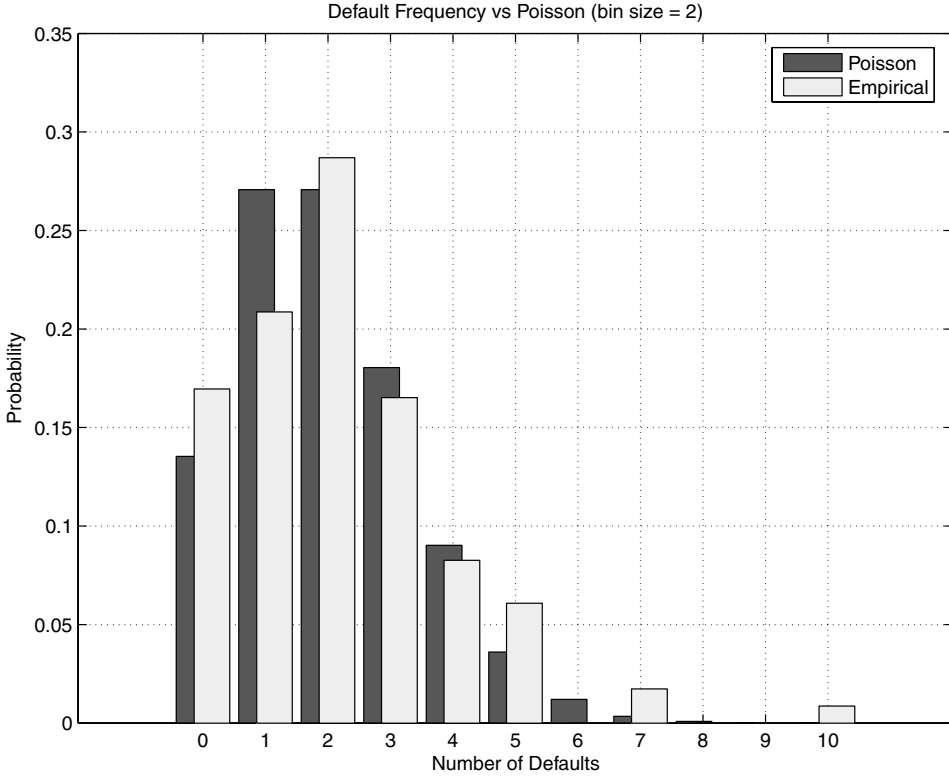


Figure 4. Default distributions. The empirical and theoretical distributions of defaults for bin size 2. The theoretical distribution is Poisson.

B. Upper Tail Tests

If defaults are more positively correlated than would be suggested by the co-movement of intensities, then the upper tail of the empirical distribution of defaults per bin could be fatter than that of the associated Poisson distribution. We use a Monte Carlo bootstrap test of the “size” (mean or median) of the upper quartile of the empirical distribution against the theoretical size of the upper quartile of the Poisson distribution as follows.

For a given bin size c , suppose there are K bins. We let M denote the sample mean of the upper quartile of the empirical distribution of X_1, \dots, X_K . By Monte Carlo simulation, we generate 10,000 data sets, each consisting of K i.i.d. Poisson random variables with parameter c . The p -value is estimated as the fraction of the simulated data sets whose sample upper-quartile size (mean or median) is above the actual sample mean M . For four of the five bin sizes, the sample p -values presented in Table III suggest fatter upper-quartile tails than those of the theoretical Poisson distribution. That is, for these bins, our one-sided tests imply rejection of the null at typical confidence levels. The joint test across all bin sizes also implies a rejection of the null at standard confidence levels.

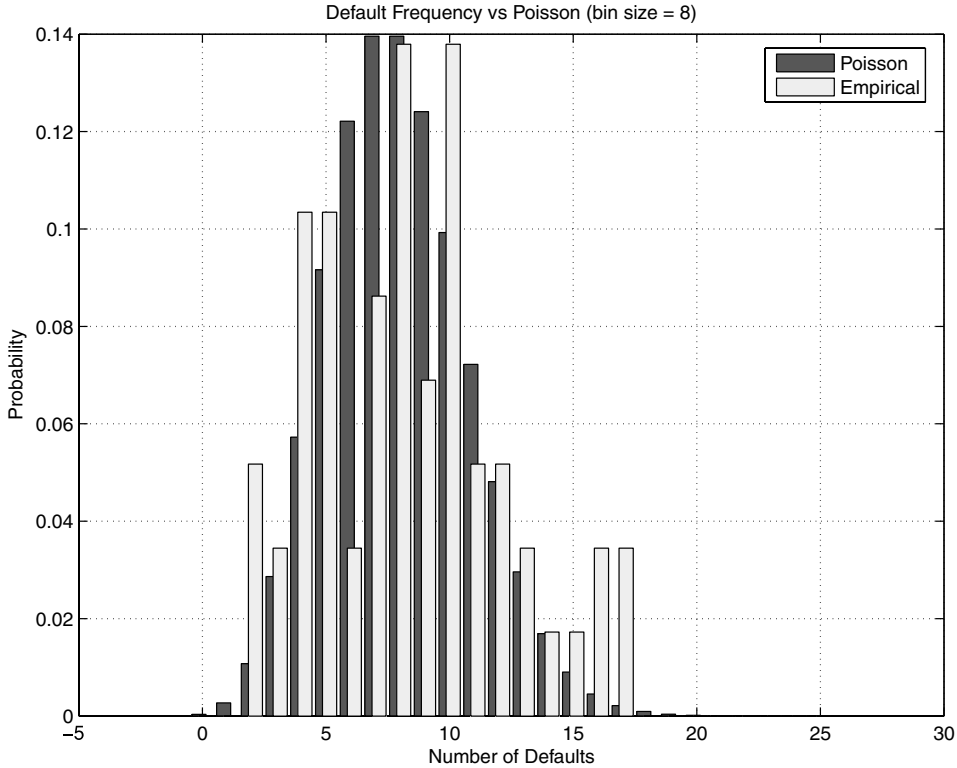


Figure 5. Default distributions. The empirical and theoretical distributions of defaults for bin size 8. The theoretical distribution is Poisson.

C. Prah1's Test of Clustered Defaults

Fisher's dispersion and our tailored upper-tail test undertaken for each bin size do not exploit the information available across all bin sizes well. In this section, we apply a test for "bursty" default arrivals due to Prah1 (1999). Prah1's test is sensitive to clustering of arrivals in excess of those of a theoretical Poisson process. This test is particularly suited for detecting clustering of defaults that may arise from more default correlation than would be suggested by co-movement of default intensities alone. Prah1's test statistic is based on the fact that the inter-arrival times of a standard Poisson process are i.i.d. standard exponential. Under the null, Prah1's test is therefore applied to determine whether, after the time change associated with aggregate default intensity accumulation, the inter-default times Z_1, Z_2, \dots are i.i.d. exponential with parameter 1. (Because of data granularity, our mean is slightly smaller than 1.)

Table IV provides the sample moments of inter-default times in the intensity-based time scale. This table also presents the corresponding sample moments of the unscaled (actual calendar) inter-default times, after a linear scaling of time that matches the mean of the inter-default time distribution to that of

Table II
Fisher's Dispersion Test

The table presents Fisher's dispersion test for goodness of fit of the Poisson distribution with mean equal to bin size. Under the joint hypothesis that default intensities are correctly measured and the doubly stochastic property, the statistic W is χ^2 -distributed with $K - 1$ degrees of freedom, and is provided in equation (2).

Bin Size	K	W	p -Value
2	230	336.00	0.0000
4	116	168.75	0.0008
6	77	132.17	0.0001
8	58	107.12	0.0001
10	46	91.00	0.0001

Table III
Mean and Median of Default Upper Quartiles

This table presents tests of the mean and median of the upper quartile of defaults per bin against the associated theoretical Poisson distribution. The last row of the table, "All," indicates the estimated probability, under the hypothesis that time-changed default arrivals are Poisson with parameter 1, that there exists at least one bin size for which the mean (or median) of number of defaults per bin exceeds the corresponding empirical mean (or median).

Bin Size	Mean of Tails			Median of Tails		
	Data	Simulation	p -Value	Data	Simulation	p -Value
2	4.00	3.69	0.00	4.00	3.18	0.00
4	7.39	6.29	0.00	7.00	6.01	0.00
6	9.96	8.95	0.02	9.00	8.58	0.06
8	12.27	11.33	0.08	11.50	10.91	0.19
10	16.08	13.71	0.00	16.00	13.25	0.00
All			0.0018			0.0003

the intensity-based time scale. A comparison of the moments indicates that conditioning on intensities removes a large amount of default correlation, in the sense that the moments of the inter-arrival times in the default-intensity time scale are much closer to the corresponding exponential moments than are those of the actual (calendar) inter-default times.

Letting C^* denote the sample mean of Z_1, \dots, Z_n , Prahla shows that

$$M = \frac{1}{n} \sum_{\{k : Z_k < C^*\}} \left(1 - \frac{Z_k}{C^*}\right) \quad (3)$$

is asymptotically (in n) normally distributed with mean $\mu_n = e^{-1} - \alpha/n$ and variance $\sigma_n^2 = \beta^2/n$, where

$$\alpha \simeq 0.189$$

$$\beta \simeq 0.2427.$$

Table IV
Moments of the Distribution of Inter-default Times

This table presents selected moments of the distribution of inter-default times. Under the joint hypothesis of doubly stochastic defaults and correctly measured default intensities, the inter-default times in intensity-based time units are exponentially distributed. The inter-arrival time empirical distribution is also shown in calendar time, after a linear scaling of time that matches the first moment, mean inter-arrival time.

Moment	Intensity Time	Calendar Time	Exponential
Mean	0.95	0.95	0.95
Variance	1.17	4.15	0.89
Skewness	2.25	8.59	2.00
Kurtosis	10.06	101.90	6.00

Using our data, for $n = 495$ default times,

$$M = 0.4055$$

$$\mu_n = \frac{1}{e} - \frac{\alpha}{n} = 0.3675$$

$$\sigma_n = \frac{\beta}{\sqrt{n}} = 0.0109.$$

The test statistic M measured from our data is 3.48 standard deviations from the asymptotic mean associated with the null hypothesis of i.i.d. exponential inter-default times (in the new time scale), indicating some evidence of default clustering in excess of that associated with the default intensities under the doubly stochastic model. (In the calendar time scale, the same test statistic M is 11.53 standard deviations from the mean μ_n under the null of exponential inter-default times.)

We also report a Kolmogorov–Smirnov (KS) test of goodness of fit of the exponential distribution of inter-default times in the new time scale. The associated KS statistic is 3.14 (which is \sqrt{n} times the usual D statistic, where n is the number of default arrivals), for a p -value of 0.000, leading to a rejection of the joint hypothesis of correctly specified conditional default probabilities and the doubly stochastic nature of correlated default. (In calendar time, the corresponding KS statistic is 4.03.) Figure 6 shows the empirical distribution of inter-default times before and after rescaling time in units of cumulative total default intensity, compared to the exponential density.

IV. Calibrating the Residual Copula Correlation

In order to gauge the degree to which default correlation is not captured by the doubly stochastic property with our data, we calibrate the intensity-conditional copula model of Schönbucher and Schubert (2001). We estimate the amount of copula correlation that must be added, after conditioning on the intensities, to match the upper-quartile moments of the empirical distribution of defaults per time bin. This measure of residual default correlation depends on the specific copula model. Here, we employ the industry standard “flat Gaussian copula,”

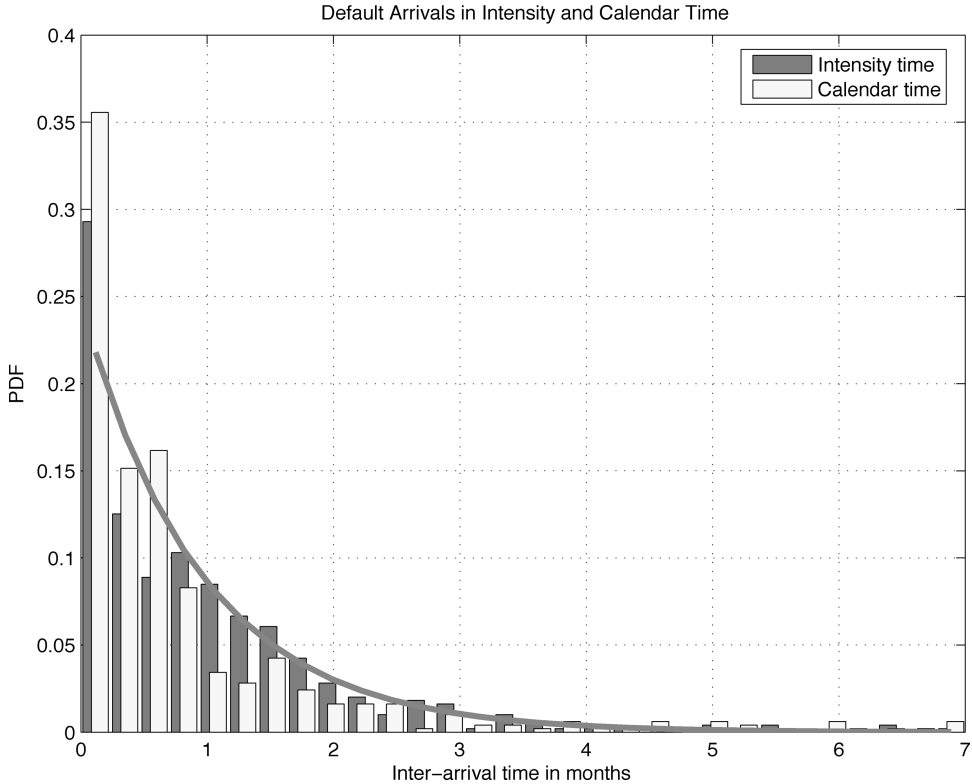


Figure 6. Inter-default times. The empirical distribution of inter-default times after scaling time change by total intensity of defaults, compared to the theoretical exponential density implied by the doubly stochastic model. The distribution of default inter-arrival times is provided both in calendar time and in intensity time. The line depicts the theoretical probability density function for the inter-arrival times of default under the null of an exponential distribution.

which is used for example to price structured credit products such as collateralized debt obligations. In the intensity time scale, the calibrated Gaussian copula correlation is a measure of the degree of correlation in default times that is not captured by co-movement in default intensities. The calibrating algorithm is provided in Appendix A. The results are reported in Table V.

As anticipated by our prior results, the calibrated residual Gaussian copula correlation r is nonnegative for all time bins, and ranges from 0.01 to 0.04. The largest estimate is for bin size 10; the smallest is for bin size 2. We can place these “residual” copula correlation estimates in perspective by referring to Akhavin et al. (2005), who estimate a Gaussian copula correlation parameter of approximately 19.7% within sectors and 14.4% across sectors by calibrating with empirical default correlations (i.e., before “removing,” as we do, the correlation associated with covariance in default intensities.)⁷ Although only a rough

⁷ Their estimate is based on a method suggested by deServigny and Renault (2002). Akhavin, Kocagil, and Neugebauer (2005) provide related estimates.

Table V
Residual Gaussian Copula Correlation

Using a Gaussian copula for intensity-conditional default times and equal pairwise correlation r for the underlying normal variables, we estimate by Monte Carlo the mean of the upper quartile of the empirical distribution of the number of defaults per bin, according to an algorithm described in the Appendix. We set in bold the correlation parameter r at which the Monte Carlo-estimated mean best approximates the empirical counterpart. (Under the null hypothesis of correctly measured intensity and the doubly stochastic assumption, the theoretical residual Gaussian copulation r is zero.)

Bin Size	Mean of Upper Quartile (data)	Mean of Simulated Upper Quartile Copula Correlation				
		$r = 0.00$	$r = 0.01$	$r = 0.02$	$r = 0.03$	$r = 0.04$
2	4.00	3.87	4.01	4.18	4.28	4.48
4	7.39	6.42	6.82	7.15	7.35	7.61
6	9.96	8.84	9.30	9.74	10.13	10.55
8	12.27	11.05	11.73	12.29	12.85	13.37
10	16.08	13.14	14.01	14.79	15.38	16.05

comparison, this indicates that correlation of default intensities accounts for a large fraction, but not all of the default correlation.

V. Tests for Missing Default Covariates

Thus far, we document violations of the joint hypothesis of correctly specified default probabilities and the doubly stochastic property. We now investigate a potential cause of these violations. In particular, the underlying default prediction model may be missing covariates that would, if present, introduce more correlation across firms in measured intensities. In general, adding more intensity covariates (that are not spurious) increases the amount of default correlation that a doubly stochastic model can capture.

A. Testing for Independent Increments

Although all of the above tests depend to some extent on the independent increments property of Poisson processes, we test specifically for serial correlation of the numbers of defaults in successive bins. That is, under the null hypothesis of doubly stochastic defaults, fixing an accumulative total default intensity of c per time bin, the numbers of defaults X_1, X_2, \dots, X_K in successive bins are independent and identically distributed. We test for independence by estimating an autoregressive model for X_1, X_2, \dots , where X_k evolves according to

$$X_k = A + BX_{k-1} + \epsilon_k \quad (4)$$

for coefficients A and B and for i.i.d. innovations $\epsilon_1, \epsilon_2, \dots$. Under the joint hypothesis of correctly specified default intensities and the doubly stochastic

Table VI
Excess Default Autocorrelation

Estimates of the autoregressive model in equation (4) of excess defaults in successive bins, for a range of bin sizes (t -statistics are shown below the parameter estimates). We test specifically for serial correlation of the numbers of defaults in successive bins. That is, under the null hypothesis of doubly stochastic defaults, fixing an accumulative total default intensity of c per time bin, the numbers of defaults X_1, X_2, \dots, X_K in successive bins are independent and identically distributed. The parameter A is the intercept in the AR1 model and B is the autoregression coefficient. T -statistics for A are presented for the test $A = c$, not $A = 0$.

Bin Size	No. of Bins	A (t_A)	B (t_B)	R^2
2	230	2.091	0.019	0.0004
		0.506	0.286	
4	116	2.961	0.304	0.0947
		-2.430	3.438	
6	77	4.705	0.260	0.0713
		-1.689	2.384	
8	58	5.634	0.338	0.1195
		-2.090	2.733	
10	46	7.183	0.329	0.1161
		-1.810	2.376	

property, $A = c$, $B = 0$, and $\epsilon_1, \epsilon_2 \dots$ are i.i.d. demeaned Poisson random variables. A significantly positive estimate for the autoregressive coefficient B would be evidence against the null hypothesis. Possibly, this could reflect missing covariates, whether they are unobservable (frailty) or are observable but missing from the estimated intensity model. For example, if a business cycle covariate should be included but is not, and if this missing covariate is persistent across time, then defaults per bin would be fatter-tailed than the Poisson distribution, and there would be serial correlation in defaults per bin.

Table VI presents the results of this autocorrelation analysis. The estimated autoregressive coefficient B is mildly significant for bin sizes of 4 and larger (with t -statistics ranging from 2.37 to 3.43). Next, we investigate whether this serial correlation can be “cured” by extending the list of covariates used to estimate the intensities.

B. Macroeconomic Covariates

A measured violation of the doubly stochastic assumption that is due to frailty (unobservable covariates that are correlated across firms) could be caused by the existence of default covariates that are in fact observable, but are not used to estimate intensities. In other words, missing covariates play the same role as do unobservable covariates.

Prior work by Lo (1986), Lennox (1999), McDonald and Van de Gucht (1999), Duffie et al. (2005), and Couderc and Renault (2004) suggests that macroeconomic performance is an important explanatory variable in default prediction. In this section, we explore the potential role of missing macroeconomic default

covariates. In particular, we examine (1) whether the inclusion of U.S. gross domestic product (GDP) or industrial production (IP) growth rates helps explain default arrivals after controlling for the default covariates that are already used to estimate our default intensities, and if so, (2) whether these variables can potentially explain the estimated violations of the doubly stochastic assumption. We find that industrial production offers some explanatory power, but GDP growth rates do not.

Under the null hypothesis of no mis-specification, fixing a bin size of c , the number of defaults in a bin in excess of the mean, $Y_k = X_k - c$, is the increment of a martingale and therefore should be uncorrelated with any variable in the information set available prior to the formation of the k^{th} bin. Consider the regression

$$Y_k = \alpha + \beta_1 GDP_k + \beta_2 IP_k + \epsilon_k, \quad (5)$$

where GDP_k and IP_k are the growth rates of U.S. gross domestic product and industrial production observed in the quarter and month, respectively, that ends immediately prior to the beginning of the k^{th} bin. In theory, under the null hypothesis of correct specification of the default intensities, the coefficients α , β_1 , and β_2 are equal to zero. Table VII reports estimated regression results for a range of bin sizes.

We report the results for the multiple regression as well as for GDP and IP separately. For all bin sizes, GDP growth is not statistically significant, and is unlikely to be a candidate for explaining the residual correlation of defaults. Industrial production enters the regression with sufficient significance to warrant its consideration as an additional explanatory variable in the default intensity model. For each of the bins, the sign of the estimated IP coefficient is negative. That is, significantly more than the number of defaults predicted by the intensity model occur when industrial production growth rates are lower than normal.

It is also useful to examine the role of missing macroeconomic factors when defaults are much higher than expected. Table VIII provides the results of a test of whether the excess upper-quartile number of defaults (the mean of the upper quartile less the mean of the upper quartile for the Poisson distribution of parameter c , as examined previously in Table III) are correlated with GDP and IP growth rates. We report two sets of regressions; the first set is based on the prior period's macroeconomic variables, and the second set is based on the growth rates observed within the bin period.⁸

We report results for those bin sizes, 2 and 4, for which we have a reasonable number of observations. Once again, we find some evidence that industrial production growth rates help explain default rates, even after controlling for estimated intensities.

⁸ The within-period growth rates are computed by compounding over the daily growth rates that are consistent with the reported quarterly growth rates.

Table VII
Macroeconomic Variables and Default Intensities

For each bin size c , ordinary least squares coefficients are reported for the regression of the number of defaults in excess of the mean, $Y_k = X_k - c$, on the previous quarter's GDP growth rate (annualized), and the previous month's growth in (seasonally adjusted) industrial production (*IP*). The number of observations is the number of bins of size c . Standard errors are corrected for heteroskedasticity; t -statistics are reported in parentheses.

Bin Size	No. Bins	Intercept	GDP	IP	R^2 (%)
2	230	0.28 (1.59)	-7.19 (-1.43)		1.06
		0.15 (1.21)		-41.96 (-2.21)	1.93
		0.27 (0.17)	-4.57 (-0.83)	-35.70 (-1.68)	2.31
4	116	0.46 (1.11)	-10.61 (-0.91)		1.14
		0.40 (1.60)		-109.28 (-2.88)	5.49
		0.53 (1.41)	-5.08 (-0.50)	-103.27 (-2.51)	5.73
6	77	1.12 (1.84)	-30.72 (-2.12)		4.99
		0.41 (-1.00)		-155.09 (-1.89)	7.55
		0.91 (1.58)	-18.09 (-1.18)	-124.09 (-1.42)	8.98
8	58	0.80 (0.85)	-19.64 (-0.74)		1.81
		1.35 (2.40)		-357.23 (-3.65)	18.63
		1.35 (1.77)	-0.08 (-0.00)	-357.20 (-3.47)	18.63
10	46	1.81 (1.57)	-49.00 (-1.62)		5.89
		0.45 (0.59)		-231.26 (-2.07)	7.66
		1.96 (1.80)	-41.45 (-1.38)	-205.15 (-2.08)	11.78

C. Augmenting the Covariates

In light of the possibility that a missing covariate, U.S. industrial production growth (*IP*), is responsible for rejections in our tests of the doubly stochastic property, we reestimate default intensities after extending Duffie et al.'s (2005) specification (1) to include *IP*. Indeed, *IP* shows up as a significant covariate, with a coefficient that is approximately 2.2 times its standard error. (The original four covariates in (1) have greater significance.) Using the estimated default intensities associated with this extended specification, we repeat all of the tests reported earlier.

Table VIII
Upper-tail Regressions

For each bin size c , ordinary least squares coefficients are shown for the regression of the number of defaults observed in the upper quartile less the mean of the upper quartile of the theoretical distribution (with Poisson parameter equal to the bin size) on the previous and current GDP and industrial production (IP) growth rates. The number of observations is the number K of bins. Standard errors are corrected for heteroskedasticity; t -statistics are reported in parentheses.

Bin Size	K	Intercept	Previous Qtr GDP	Previous Month IP	R^2 (%)
2	77	0.28	1.40		0.00
		(1.55)	(0.22)		
		0.36		-57.75	4.92
		(2.08)		(-2.46)	
4	48	0.16	8.99	-76.80	6.94
		(1.04)	(1.04)	(-2.11)	
		0.41	-6.19		0.97
		(1.24)	(-0.71)		
4	48	0.29		-65.83	3.88
		(-1.26)		(-1.64)	
		0.29	-22.15	-65.26	3.88
		(0.79)	(-0.02)	(-1.14)	

Bin Size	K	Intercept	Current Bin GDP	Current Bin IP	R^2 (%)
2	77	0.45	-5.98		1.03
		(1.67)	(-0.82)		
		0.38		-47.20	2.82
		(2.04)		(-2.07)	
4	48	0.36	0.98	-50.28	2.84
		(1.23)	(0.10)	(-1.56)	
		0.83	-23.29		12.67
		(1.60)	(-2.44)		
4	48	0.48		-77.93	17.88
		(1.90)		(-3.07)	
		0.63	-7.85	-62.55	18.63
		(1.78)	(-0.74)	(-2.30)	

Our primary conclusion remains unchanged. Albeit with slightly higher p -values, the results of all tests are consistent with those reported for the original intensity specification (1), and lead to a rejection of the estimated doubly stochastic model. For example, the goodness-of-fit test rejects the Poisson assumption for every bin size; the upper-tail tests analogous to those of Table III result in a rejection of the null at the 5% level for three of the five bins, and at the 10% level for the other two. The Prahl test statistic using the extended specification is 3.25 standard deviations from its null mean (as compared with 3.48 for the original model). The calibrated residual Gaussian copula correlation parameter r is the same for each bin size as that reported in Table V. Overall, even with the augmented intensity specification, the tests suggest more clustering than implied by correlated changes in the modeled intensities.

VI. Conclusions and Discussion

Defaults cluster in time because firms' default intensity processes are correlated and because, even after conditioning on these intensities, there could be contagion or frailty (unobserved covariates that are correlated across firms). The latter channels are not admitted in a doubly stochastic setting with intensities that are based on all available information. While the doubly stochastic assumption forms the current basis of risk management practice, no test of its validity has been undertaken. This paper makes the following contributions.

1. We introduce a time change technique that, under the doubly stochastic hypothesis, reduces the process of cumulative defaults to a standard (unit intensity) Poisson process. Based on this technique, we provide newly developed tests of the joint hypothesis that default intensities are correctly measured and the doubly stochastic property holds.⁹ We are particularly interested in whether defaults are indeed independent after conditioning on intensities.
2. The null of correctly measured intensities and the doubly stochastic property is rejected in various tests of the hypothesis that the numbers of defaults that occur in successive time periods, all containing the same cumulative total of default intensity, are i.i.d. Poisson.
3. The null is also rejected in a test for exponentially distributed inter-default times, where time is measured in units of cumulative total default intensity.
4. Introducing a measure of residual Gaussian copula correlation, we find that the excess default clustering in our data above and beyond that implied by the factors that cause correlated default intensities can be matched by injecting moderately small amounts of "extra copula correlation."
5. We consider whether the excess degree of default correlation can be explained by missing macroeconomic covariates. While we find some evidence that growth rates of U.S. industrial production (IP) do provide some incremental explanatory role, even after controlling for IP the resulting doubly stochastic model of default correlation is rejected by the data.

These results address the ability of commonly applied credit risk models to capture the tails of the probability distribution of portfolio default losses, and may therefore be of particular interest to bank risk managers and regulators. For example, the level of economic capital necessary to support levered portfolios of corporate debt at high confidence levels is heavily dependent on the degree to which the doubly stochastic property that we test actually applies in practice. This may be of special interest with the advent of more quantitative portfolio credit risk analysis in bank capital regulations, arising under the proposed Basel II (BIS) accord on regulatory capital (see Gordy (2003), Allen and

⁹ Giesecke and Goldberg (2005) provide some new and related results based on Meyer (1971).

Saunders (2003), and Kashyap and Stein (2004)). The results also present a challenge to develop more realistic models of default correlation.

Appendix A: Residual Gaussian Copula Correlation

We estimate the residual Gaussian copula correlation by the following algorithm.

1. We fix a particular correlation parameter r and cumulative-intensity bin size c .
2. For each name i and each bin number k , we calculate the increase in cumulative intensity $C_i^{c,k}$ for name i that occurs in this bin. (The intensity for this name stays at zero until name i appears, and the cumulative intensity stops growing after name i disappears, whether by default or otherwise.)
3. For each scenario j of 5,000 independent scenarios, we draw one of the bins, say k , at random (equally likely), and draw joint standard normal X_1, \dots, X_n with $\text{corr}(X_i, X_m) = r$ whenever i and m differ.
4. For each i , we let $U_i = F(X_i)$ be the standard normal cumulative distribution function $F(\cdot)$ evaluated at X_i , and draw “default” for name i in bin k if $U_i > \exp(-C_i^{c,k})$.
5. We report in Table III the mean of the upper quartile of the simulated distribution (across scenarios j) of the number of defaults per bin.
6. A correlation parameter r is that “calibrated” to the data for bin size c , to the nearest 0.01, if the associated upper-quartile mean best approximates the upper-quartile mean of the actual data, also reported in Table III.

REFERENCES

- Akhavain, Jalal D., Ahmet E. Kocagil, and Matthias Neugebauer, 2005, A comparative empirical study of asset correlations, Working paper, Fitch Ratings, New York.
- Allen, Linda, and Anthony Saunders, 2003, A survey of cyclical effects in credit risk measurement models, BIS Working paper 126, Basel Switzerland.
- Bharath, Sreedhar, and Tyler Shumway, 2005, Forecasting default with the KMV-Merton model, Working paper, University of Michigan.
- Billingsley, Patrick, 1986, *Probability and Measure* (Wiley, New York, II).
- Cathcart, Lara, and Lina El Jahel, 2002, Defaultable bonds and default correlation, Working paper, Imperial College.
- Cochran, W.G., 1954, Some methods of strengthening χ^2 tests, *Biometrics* 10, 417–451.
- Collin-Dufresne, Pierre, Robert Goldstein, and Jean Helwege, 2003, Is credit event risk priced? Modeling contagion via the updating of beliefs, Working paper, Haas School, University of California, Berkeley.
- Collin-Dufresne, Pierre, Robert Goldstein, and Julien Huggonier, 2004, A general formula for valuing defaultable securities, *Econometrica* 72, 1377–1407.
- Couderc, Fabien, and Olivier Renault, 2004, Times-to-default: Life cycle, global and industry cycle impacts, Working paper, University of Geneva.
- Das, Sanjiv, Laurence Freed, Gary Geng, and Nikunj Kapadia, 2006, Correlated default risk, *Journal of Fixed Income* 16, 7–32.
- Davis, Mark, and Violet Lo, 2001, Infectious default, *Quantitative Finance* 1, 382–387.

- deServigny, Arnaud, and Olivier Renault, 2002, Default correlation: Empirical evidence, Working paper, Standard and Poors.
- Duffie, Darrell, Leandro Saita, and Ke Wang, 2005, Multi-period corporate default prediction with stochastic covariates, Working paper, Stanford University.
- Giesecke, Kay, 2004, Correlated default with incomplete information, *Journal of Banking and Finance* 28, 1521–1545.
- Giesecke, Kay, and Lisa Goldberg, 2005, A top-down approach to multiname credit, Working paper, Cornell University.
- Gordy, Michael, 2003, A risk-factor model foundation for ratings-based capital rules, *Journal of Financial Intermediation* 12, 199–232.
- Jarrow, Robert, and Fan Yu, 2001, Counterparty risk and the pricing of defaultable securities, *Journal of Finance* 56, 1765–1800.
- Jarrow, Robert, David Lando, and Fan Yu, 2005, Default risk and diversification: Theory and applications, *Mathematical Finance* 15, 1–26.
- Kashyap, Anil, and Jeremy Stein, 2004, Cyclical implications of the Basel-II capital standards, Federal Reserve Bank of Chicago, *Economic Perspectives* 28, 18–31.
- Kusuoka, Shigeo, 1999, A remark on default risk models, *Advances in Mathematical Economics* 1, 69–82.
- Lang, Larry, and Rene Stulz, 1992, Contagion and competitive intra-industry effects of bankruptcy announcements, *Journal of Financial Economics* 32, 45–60.
- Lennox, Clive, 1999, Identifying failing companies: A reevaluation of the logit, probit, and DA approaches, *Journal of Economics and Business* 51, 347–364.
- Lo, Andrew, 1986, Logit versus discriminant analysis: Specification test and application to corporate bankruptcies, *Journal of Econometrics* 31, 151–178.
- Lucas, Douglas J., 1995, Default correlation and credit analysis, *The Journal of Fixed Income* 5, 76–87.
- McDonald, Cynthia, and Linda M. Van de Gucht, 1999, High yield bond default and call risks, *Review of Economics and Statistics* 81, 409–419.
- Merton, Robert C., 1974, On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance* 29, 449–470.
- Meyer, Paul-André, 1971, Demonstration simplifiée d'un theorem Knight, in Séminaire de Probabilités V, Lecture Note in Mathematics 191, Springer-Verlag Berlin, 191–195.
- Prahl, Juergen, 1999, A fast unbinned test on event clustering in Poisson processes, Working paper, University of Hamburg.
- Protter, Philip, 2003, *Stochastic Integration and Differential Equations* (Springer, New York, II).
- Schönbucher, Philipp, 2003, Information driven default contagion, Working paper, Eidgenössische Technische Hochschule, Zurich.
- Schönbucher, Philipp, and Dirk Schubert, 2001, Copula dependent default risk in intensity models, Working paper, Bonn University.
- Shumway, Tyler, 2001, Forecasting bankruptcy more accurately: A simple hazard model, *Journal of Business* 74, 101–124.
- Sobehart, Jorge, Roger Stein, V. Mikityanskaya, and L. Li, 2000, Moody's public firm risk model: A hybrid approach to modeling short term default risk, Moody's Investors Service, Global Credit Research, Rating Methodology, March.
- Terentyev, Sergey, 2004, Asymmetric counterparty relations in default modeling, Working paper, Stanford University.
- Vasicek, Oldrich, 1987, Probability of loss on loan portfolio, Working paper, KMV Corporation.
- Yu, Fan, 2003, Default correlation in reduced form models, *Journal of Investment Management* 3, 33–42.
- Yu, Fan, 2005, Accounting transparency and the term structure of credit spreads, *Journal of Financial Economics* 75, 53–84.
- Zhang, Gaiyan, 2004, Intra-industry credit contagion: Evidence from the credit default swap market and the stock market, Working paper, University of California, Irvine.
- Zhou, Chunsheng, 2001, An analysis of default correlation and multiple defaults, *Review of Financial Studies* 14, 555–576.

