Common Failings: How Corporate Defaults are Correlated[1]

Sanjiv R. Das
Santa Clara University
Santa Clara, CA 95053

Darrell Duffie
Stanford University
Stanford, CA 94305

Nikunj Kapadia
University of Massachusetts
Amherst, MA 01003.

February, 2005

Abstract

Common Failings: How
Corporate Defaults are Correlated

We develop, and apply to data on U.S. corporations from 1987-2000, tests of the standard doubly-stochastic assumption under which firms' default times are correlated only as implied by the correlation of factors determining their default intensities. This assumption is violated in the presence of contagion or "frailty" (unobservable covariates for default 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, although we provide evidence that this may be due to mis-specification of the default intensities, which do not include macroeconomic default-prediction covariates. Despite this rejection, there is no evidence of significant default clustering in excess of that implied by the doubly-stochastic model and correlation of observable firm-specific default covariates.

1   Introduction

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 changes over time in conditional default probabilities that are correlated across firms. Second, the event of default by one firm may be "contagious," in that this event itself can push other firms toward default. For example, there could be a "domino" or cascade effect, under which corporate failures directly induce other corporate failures, as with the collapse of Penn Central Railway in 1970. A third channel for default correlation is learning from defaults. For example, the defaults of Enron and WorldCom may have revealed accounting irregularities that could be present in other firms, and thus 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 via observable factors determining conditional default probabilities, that is, the first channel on its own, is sufficient to account for the degree of time-clustering of defaults that we find in the data.

Specifically, we test whether our data are consistent with the standard doubly-stochastic model of default, under which, conditional on the path of risk factors determining all firms' default intensities, defaults are independent Poisson arrivals with these (conditionally deterministic) intensity paths. This model is particularly convenient for computational and statistical purposes, although its empirical relevance for default correlation has been unresolved. We develop, and apply to default data for U.S. corporations during the period 1987-2000, a new test of the doubly-stochastic assumption. The data do not support the joint hypothesis of well specified default intensities and the doubly-stochastic assumption, although we provide evidence that this rejection may be due to mis-specification of the default intensities, which do not include macroeconomic default-prediction covariates. These missing macroeconomic covariates may be responsible for some clustering of defaults. Despite the rejection based on goodness-of-fit tests, we do not find substantial 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, as backing for 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 do so on the basis of models in which default correlation is captured by common risk factors determining conditional default probabilities, as in Gordy [2003] and Vasicek [1987]. (Banks do, however, 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 currently envisioned in these default-risk models, however, then significantly greater capital might be required in order to survive default losses 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 that are exposed to correlated default, such as collateralized debt obligations (CDOs) and options on portfolios of default swaps. The Bank of America has reported that synthetic CDO volumes reached over $500 billion in 2003, an annual growth rate of over 130%.

While there is some empirical evidence regarding average default correlation (Lucas [1995] and deServigny and Renault [2002]) and correlated changes in corporate default probabilities (Das, Freed, Geng, and Kapadia [2001]), there is relatively little evidence regarding the presence of clustered defaults. In particular, there has been no prior 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 a clustering effect in excess of that suggested by the doubly-stochastic model, or at least a failure of the doubly-stochastic model under risk-neutral probabilities. That is, their findings may be due to default-induced increases in the conditional default probabilities of other firms, or could be due to default-induced increases in default risk premia[1] of other firms, as envisioned by Kusuoka [1999]. 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, under which default intensities include the expected effect of unobservable covariates. In a frailty setting, the arrival of a default causes, via Bayes' 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 [2004] finds 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.

Banks and other managers of credit portfolios could in theory extend the doubly-stochastic model if it were found to be seriously deficient. At this point, there are few if any methods applied in practice to measure loan portfolio credit risk that allow a role 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) that determine

---

[1]Collin-Dufresne, Goldstein, and Huggonier [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 El Jahel [2002], Davis and Lo [2001], Giesecke [2004], Kusuoka [1999], Schönbucher and Schubert [2001], Terentyev [2003], Yu [2003], and Zhou [2001].

conditional default probabilities. 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-ratings transitions intensities are based on commonly observable covariates.

The doubly stochastic property, sometimes called "conditional independence," also underlies the standard econometric duration models used for event forecasting, including default prediction models, such as Shumway [2001] and Duffie and Wang [2003]. The property allows the likelihood function that is to be maximized when estimating the coefficients of an intensity model to be expressed as the product of the likelihood functions of each of the underlying events in the data. One of our objectives is to provide a tool with which to check 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 data on actual default times and on monthly estimates of conditional probabilities of default within one year (PDs) were provided to us by Moodys, and cover the period January, 1987 to October, 2000. These data are described in Section 3, with further details in Appendix A. After dropping firms for which we had missing data, we were left with 241 individual issuer defaults among a total of 1,990 firms over 216,859 firm-months of data.

From the time-series of PD data for each firm, we estimate default intensities for each firm, using a simple time-series model of intensities. For this, we assume that the default intensity process for each firm is a Feller diffusion (also known as a Cox-Ingersoll-Ross process, or a square-root diffusion). The fitting procedure is outlined in Section 3.2. The current intensity level measured from the one-year default probability is relatively robust to mis-specification of the Feller diffusion model, since intensities and one-year conditional default probabilities are relatively close for a wide range of alternative intensity models and reasonable parameters. We rely on this fitting procedure only for the levels of default intensities, and not for the associated implied probability distributions of intensity movements.

We then exploit the following new result, developed in Section 2. 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, roughly speaking, 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. Thus, the doubly-stochastic property implies, once fixing any scalar $c > 0$ and considering successive non-overlapping time intervals each lasting for c units of new time (corresponding to periods that include an accumulated total default intensity, across all firms, of c), that the number of defaults in the 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 $X_1, \ldots, X_k, \ldots$ 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 mildly rejected in some cases, at traditional confidence levels.

2. We test whether the mean of the upper quartile of our sample $X_1, X_2, \ldots, X_K$ 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. For larger bin sizes, the null is rejected.

3. In order to avoid reliance on specific bin sizes, 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. In this case, the null is not rejected.

4. Fixing the size c of time bins, we test for serial correlation of $X_1, X_2, \ldots$ 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 is positive, could lead to default clustering in excess of that associated with the doubly-stochastic assumption. For certain specifications, the null is rejected at traditional confidence levels.

An appealing feature of these tests is that they do not depend on the joint probability distribution of the default intensity processes of the firms, including their correlation structure, allowing both generality and robustness. We find the data broadly consistent with a rejection of the joint hypothesis of correctly specified intensities and the doubly-stochastic hypothesis, at standard confidence levels. In light of this, we also test for the presence of missing covariates in the PD model, which was estimated from only firm-specific covariates such as leverage, asset volatility, and credit rating. We are especially concerned about missing default covariates that might be associated with default clustering, such as business-cycle variables, which were not included in the statistical model used by Moody's to estimate our default probability data. Indeed, we find evidence, in some tests, that certain macroeconomic business-cycle variables should probably have been included as default-prediction covariates. For example, the number of defaults in a given bin, in excess of its conditional mean, is in theory uncorrelated with any variables in the information set of the observer

before the time bin begins. Among other related results, however, we find some evidence of correlation between X_k, the number of defaults in bin k, and macroeconomic variables such as GDP growth and industrial production (IP) that were observed before bin k begins. It is thus indeed plausible that missing covariates, rather than a failure of the doubly-stochastic property, is responsible for the rejection of the joint hypothesis that we test.

In order to gauge the degree of default correlation that is not captured by default intensity correlation, we calibrate a standard version of the Gaussian copula model of default correlation, estimating the amount of additional correlation that must be added (in the sense of the Gaussian copula correlation parameter), on top of the correlation already present in default intensities, in order to match the degree of default clustering observed in the data. Consistent with our other results, this incremental correlation estimate is relatively small, at most 2% depending on the length of time window used.

The rest of the paper comprises the following. In Section 2, we derive the property that the total default arrival process is a Poisson process with constant intensity under a time rescaling based on aggregate default intensity accumulation. This provides our testable implications. Section 3 describes our data, comprising default probabilities and default times over a period of fourteen years. This section also describes the conversion of default probabilities into intensities. Section 4 provides various tests of the doubly-stochastic hypothesis, and Section 5.1 addresses the question of independence of increments of the time-changed process governing default arrival. In Section 5.2, we test our default intensity data for missing macroeconomic covariates. (This test does not depend on the doubly-stochastic property.) Section 6 concludes. The appendices contain further details on the data and estimation procedures.

2 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 some testable implications of this property.

We fix a probability space $(\Omega, \mathcal{F}, P)$ 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 of n firms, default occurs at the first jump time τ_i of a non-explosive 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, in particular any correlation among, the default times $\tau_1, \ldots, \tau_n$ 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 $\lambda_1, \ldots, \lambda_n$ of all firms, the times of default 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 multi-dimensional counting process $N = (N_1, \ldots, N_n)$ is doubly stochastic. That is, conditional on the path $\{\lambda_t = (\lambda_{1t}, \ldots, \lambda_{nt}) : t \geq 0\}$ of all intensity processes, as well as the information $\mathcal{F}_T$ available at any given stopping time T, the counting processes $\hat{N}_1, \ldots, \hat{N}_n$, defined by $\hat{N}_i(u) = N_i(u + T)$, are independent Poisson processes with respective (conditionally deterministic) intensities $\hat{\lambda}_1, \ldots, \hat{\lambda}_n$ defined by $\hat{\lambda}_i(u) = \lambda_i(u+T)$. In this case, we also say that $(\tau_1, \ldots, \tau_n)$ is doubly-stochastic with intensity $(\lambda_1, \ldots, \lambda_n)$. In particular, the doubly-stochastic assumption implies that the default times $\tau_1, \ldots, \tau_n$ are independent given the intensities.

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

Proposition. *Suppose that $(\tau_1, \ldots, \tau_n)$ is doubly stochastic with intensity $(\lambda_1, \ldots, \lambda_n)$. Let $K(t) = \#\{i : \tau_i \leq t\}$ be the cumulative number of defaults by t, and let $U(t) = \int_0^t \sum_{i=1}^n \lambda_i(u) 1_{\{\tau_i > u\}} \, du$ 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 *iid* exponential with parameter 1. Let $T(j) = \inf\{t : K(t) \geq j\}$. By construction,

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

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

In order to check the independence of $Z_1, Z_2, \ldots$, consider any integer $k > 1$ and any bounded Borel functions $f_1, \ldots, f_k$. 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, $Z_1, Z_2 \ldots$ 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 shall test are based on the following characterization.

Corollary (Poisson property): *Under the conditions of the proposition, for any $c > 0$, the successive numbers of defaults per bin,*

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

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.

3 Data

Our empirical tests are based on a dataset of default probabilities and default events, both of which were developed by Moody's Investor Services.

3.1 Description of the Data

The data on default probabilities consists of a monthly time series of estimated conditional one-year default probabilities for public non-financial North American firms over the period January, 1987 to October, 2000. These default probabilities are the output of a logit model estimated from the history of firm-specific financial covariates and default times. A key covariate is the 'distance-to-default' measure suggested by the Merton [1974] model, which is an estimate of the number of standard deviations of annual asset growth by which assets exceed a measure of book liabilities. Other covariates include financial statement information and Moody's rating, when available. Details of the model and its econometric fit and performance are described in Sobehart, Stein, Mikityanskaya and Li [2000] and Sobehart Keenan and Stein [2000]. This database of estimated default probabilities was part of Moody's *RiskCalc* system. (Moody's subsequently distributed a related default probability estimate, the Moody's KMV EDF, also based on distance to default.)

Key advantages of this PD dataset include: *(i)* it is relatively comprehensive, and *(ii)* it is consistent with Moody's database of historical defaults over the sample period. In particular, the database that we use, extracted from Moody's overall database, covers 1,990 firms, and includes almost all firms that have been rated by Moody's over this period.

Using a separate database of defaults also obtained from Moody's, we identify a total of 241 defaults of the rated firms in our database. Much of the matching of
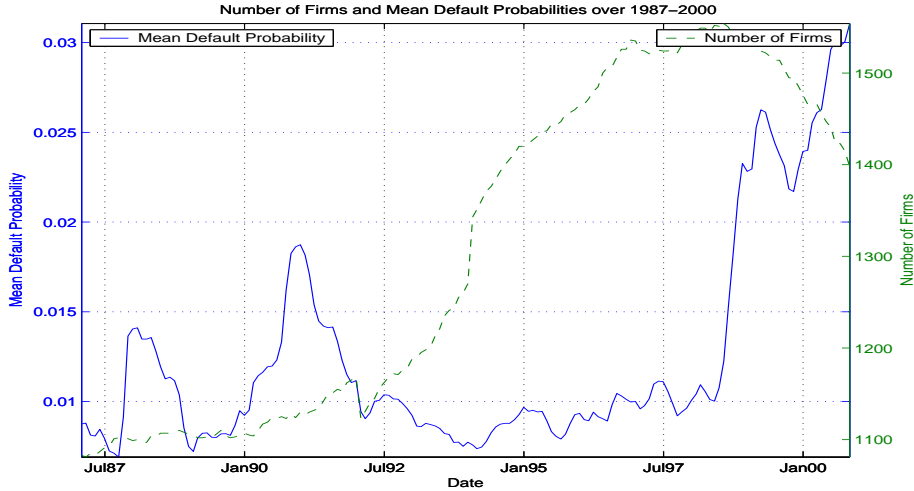
Figure 1: Cross-sectional sample mean of one-year conditional default probabilities, and the number of firms covered, January 1987 to October 2000.

firms across the two databases is done manually. In the end, there is a close match between the mean number of defaults implied by the default probabilities and the actual number of defaults. We discuss this in more detail in our analysis to follow. Appendix A provides further details on the construction of the database.

Figure 1 shows a plot of the monthly cross-sectional sample mean of estimated one-year conditional default probabilities. The plot shows evidence of positive correlation of default intensities, in that the cross-sectional mean of one-year conditional probabilities of default ranges from 0.69% to 3.11%, and increases markedly with the U.S. recession that occurred around 2000-2001. The number of firms in our sample at a given time increases from a low of 1,081 firms at the beginning of the sample period in 1987 to a high of 1,554 firms in the second half of 1998. Figure 2 shows a plot of the number of defaults over this period, month by month, ranging from 0 to a maximum of 8 per month, as well as a plot of the total of the estimated default intensities of all sampled firms. We turn next to the estimation of these intensities from one-year conditional default probabilities.

3.2   From PDs to Default Intensities

In order to test the doubly-stochastic assumption using the new-time-scale Poisson process described in the Proposition of Section 2, we estimate default intensities, firm by firm, from the PD data of one-year conditional default probabilities, as follows.

For a given firm, the default intensity process λ_t is assumed to satisfy a stochastic

differential equation of the form

$$d\lambda_t = k(\theta - \lambda_t)\,dt + \sigma\sqrt{\lambda_t}\,dz_t, \tag{1}$$

where z is a standard Brownian motion, and where k, θ, and σ are positive numbers. The doubly-stochastic assumption implies that the T-maturity survival probability at time t, for a currently surviving firm, is

$$s_t(T) = E_t\left[\exp\left(-\int_t^{t+T}\lambda_u\,du\right)\,\Big|\,\lambda_t\right]. \tag{2}$$

Cox, Ingersoll, and Ross [1985] have provided the well-known solution:

$$s_t(T) = A(T)\exp\left[-\lambda_t B(T)\right], \tag{3}$$

where

$$A(T) = \left(\frac{2\gamma e^{(k+\gamma)T/2}}{(k+\gamma)(e^{\gamma T}-1)+2\gamma}\right)^{\frac{2k\theta}{\sigma^2}} \tag{4}$$

$$B(T) = \frac{2e^{\gamma T}-1}{(k+\gamma)(e^{\gamma T}-1)+2\gamma} \tag{5}$$

$$\gamma = \sqrt{k^2 + 2\sigma^2}. \tag{6}$$

Inverting equation (3), we get, for any time horizon T,

$$\lambda_t = -\frac{1}{B(T)}\,\ln\left[\frac{s_t(T)}{A(T)}\right]. \tag{7}$$

Our PD data are monthly observations of the one-year default probability, $1-s_t(1)$. We estimate the parameters $\{k,\theta,\sigma\}$, and the default intensities of each firm, by a method-of-moments estimator provided in Appendix B. The estimator matches the time-series behavior of λ_t implied by the Feller diffusion, using the relationship between default intensity and PD given by (7). Maximum likelihood estimation has also been used in similar settings, and is efficient in large samples, but is notoriously biased in small samples. Our method-of-moments estimator is robust and computationally efficient, usually able to fit a given firm's default intensity model in a couple of seconds. In any case, the fit is relatively robust to mis-specification of the time-series model and to fitting error, as intensities are relatively close to one-year default probabilities (except for cases of extreme volatility or drift parameters). We use this fitting procedure only for the implied levels of default intensities, and not for the associated probability transition distributions implied by the fitted Feller diffusion parameters. Figure 2 shows the total of the estimated intensities of all firms, as well as the monthly arrivals of defaults.
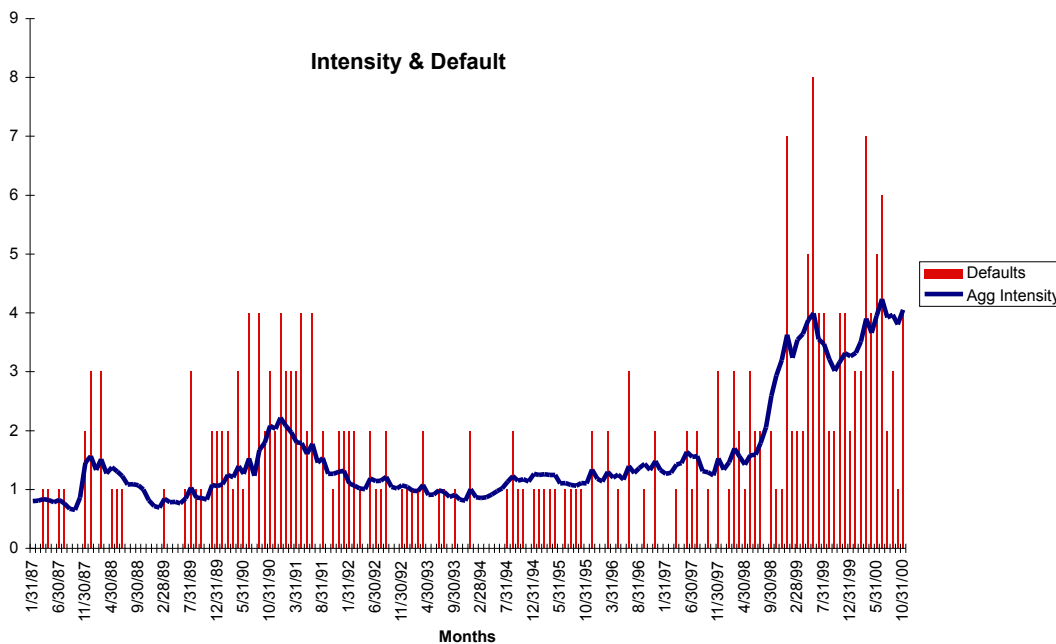
Figure 2: Aggregate (across firms) default intensities and firm defaults from 1987-2000.

4   Goodness-of-Fit Tests

Having estimated default intensities 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 name i, we let $U(t) = \int_0^t \sum_{i=1}^n \lambda_{is} 1_{\{\tau(i) > s\}} \, ds$ 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 $t_0, t_1, t_2, \ldots$ such $t_0 = 0$ and $U(t_i) - U(t_{i-1}) = c$. We then let $X_k = \sum_{i=1}^n 1_{\{t_k \leq \tau(i) < t_{k+1}\}}$ be the number of defaults in the k-th time bin. Figure 3 illustrates the time bins of size $c = 8$ over the last five calendar years of our data set.

Table 1 presents a comparison of the empirical and theoretical moments of the distribution of defaults per bin, for each of several bin sizes.[2] The actual bin sizes vary slightly from the integer bin sizes shown because of granularity in the construction of the binning times $t_1, t_2, \ldots$. The approximate match between a bin size and the associated sample mean $(X_1 + \cdots + X_K)/K$ of the number of defaults per bin offers some confirmation that the underlying PD data are reasonably well estimated, however this is somewhat expected given the within-sample nature of the estimates. For larger bin sizes, Table 1 shows that the empirical variances are bigger than their theoretical counterparts under the null of correctly specified doubly-stochastic intensity

---

[2]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 $3 + c^{-1}$.
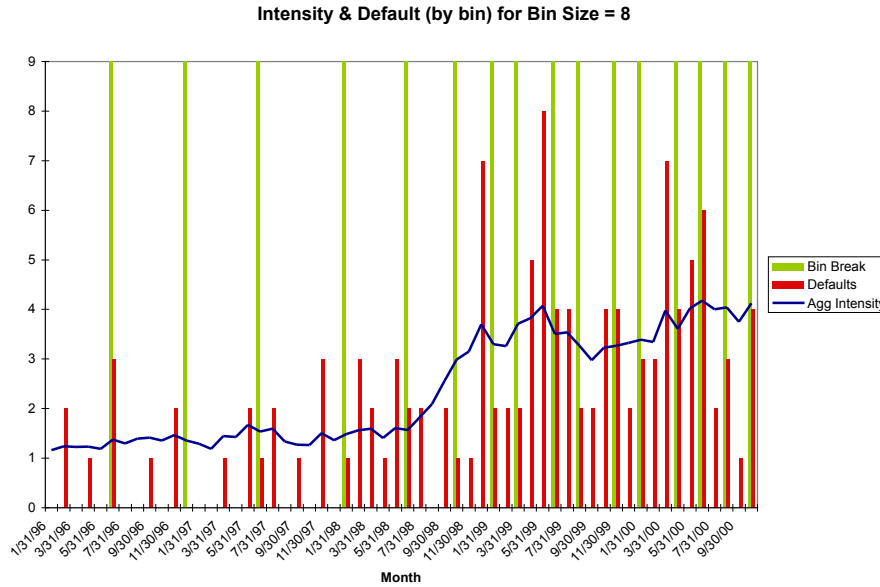
Figure 3: Aggregate intensities and defaults by month, 1996-2000, with time bin delimiters marked for intervals that include a total accumulated default intensity of $c = 8$ per bin.

model of defaults.

Figure 4 presents the observed default frequency distribution, and the associated theoretical Poisson distribution, for bin sizes 2 and 8. For bins of size larger than 4, there is a tendency for bi-modality (two peaks), as opposed to the unimodal theoretical Poisson distribution associated with the hypothesis of doubly-stochastic defaults. To the extent that the measured intensities are based on an incomplete set of covariates, one might suspect that violations of the Poisson distribution are 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.

4.1 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 $X_1, \ldots, X_K$ are independent Poisson distributed variables with mean parameter c is Fisher's dispersion
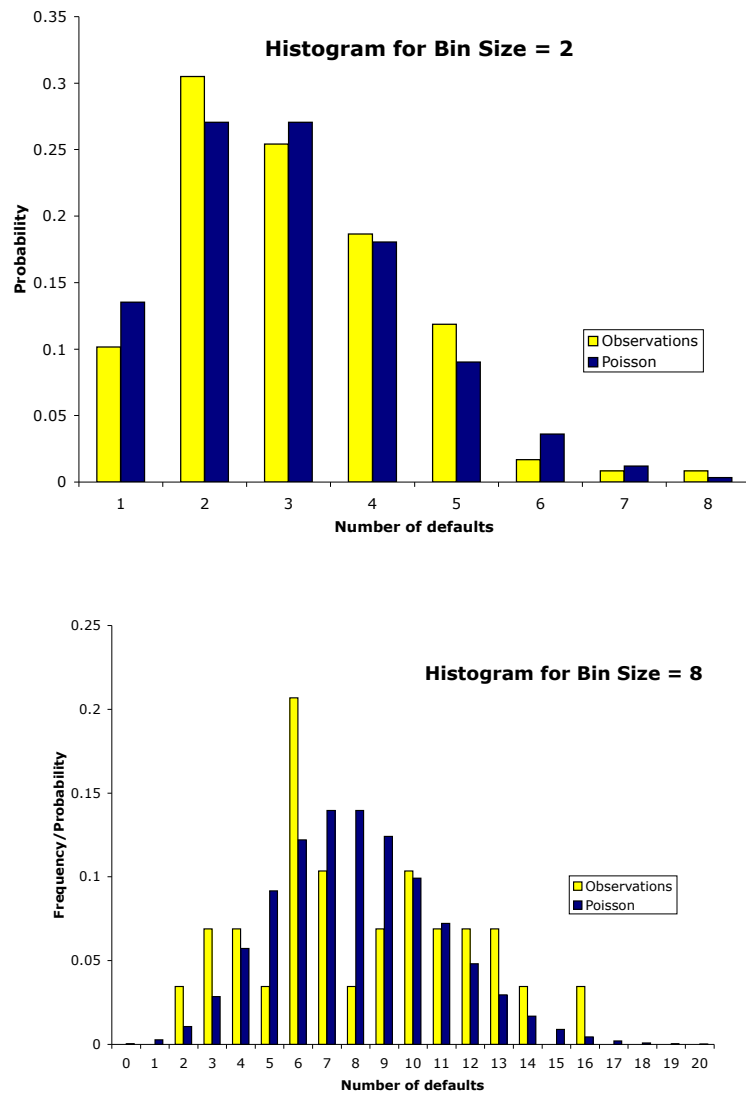
Figure 4: Comparison of the empirical and Poisson distributions of defaults for bin sizes 2 and 8.

Table 1: Comparison of empirical and theoretical moments for the distribution of defaults per bin. The number 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	2.00	2.00	0.71	3.50
(118)	2.04	1.89	0.71	3.52
4	4.00	4.00	0.50	3.25
(59)	4.07	4.00	0.41	2.06
6	6.00	6.00	0.41	3.17
(39)	6.08	8.07	0.41	2.19
8	8.00	8.00	0.35	3.12
(29)	8.14	13.12	0.26	2.07
10	10.00	10.00	0.32	3.10
(24)	10.04	15.43	0.82	2.25

test (Cochran [1954]). Under this null,

$$W = \sum_{i=1}^{K} \frac{(X_i - c)^2}{c},$$ (8)

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 cause 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 2 indicate that, at standard confidence levels such as 95%, there is a borderline rejection of this null hypothesis for bin sizes 8 and 10.

4.2   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 $X_1, \ldots, X_K$. By Monte Carlo simulation, we generated 10,000 data sets, each consisting of K *iid* Poisson random

Table 2: 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, W is χ^2-distributed with $K-1$ degrees of freedom.

Bin Size	K	W	p-value
2	118	110.5	0.65
4	59	58.0	0.47
6	39	51.2	0.07
8	29	46.0	0.02
10	24	35.5	0.05

Table 3: Tests of median and mean of the upper upper quartile of defaults per bin, against the associated theoretical Poisson distribution. The last line in the table, denoted "All" is the 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		p-value	Median of Tails		p-value
	Data	Simulation		Data	Simulation	
2	3.62	3.63	0.58	3.00	3.18	0.25
4	6.71	6.25	0.21	6.00	5.90	0.17
6	10.00	8.81	0.05	9.50	8.42	0.07
8	12.75	11.12	0.03	12.50	10.69	0.03
10	16.00	13.71	0.02	16.50	13.26	0.00
All			0.70			0.44

variables with parameter c. We then compute the fraction p of the simulated data sets whose sample upper-quartile size (mean or median) is above the actual sample mean M. Under the null hypothesis that the distribution of the actual sample is Poisson with parameter c, the expected p-value would be approximately 0.5.

The sample p-values presented in Table 3 suggest, for larger bin sizes, fatter upper-quartile tails than those of the theoretical Poisson distribution. (That is, our one-sided tests imply rejection of the null, for larger bins, at typical confidence levels.)

We corroborated these results with an analysis of the tail distributions using the Pearson χ^2 statistic for the theoretical tail distribution associated with the corresponding theoretical Poisson distribution. The results (not reported) imply a rejection of a Poisson-distributed upper-quartile distribution at standard confidence levels.

Table 4: 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	1.07	1.07	1.07
Variance	1.19	2.19	1.16
Skewness	2.13	3.87	2.00
Kurtosis	7.46	22.01	6.00

4.3   Prahl's Test of Clustered Defaults

Fisher's dispersion and our tailored upper-tail test, undertaken for each bin size, do not exploit well the information available across all bin sizes. In this section, we apply a test for "bursty" default arrivals due to Prahl [1999]. Prahl'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. Prahl's test statistic is based on the fact that the inter-arrival times of a standard Poisson process are *iid* standard exponential. Under the null, Prahl's test is therefore applied to test whether, after the time change associated with aggregate default intensity accumulation, the inter-default times $Z_1, Z_2, \ldots$ are *iid* exponential with parameter 1. (Because of data granularity, our mean is slightly larger than 1.)

The sample moments of inter-default times in the intensity-based time scale are provided in Table 4. This table also presents the corresponding sample moments of the un-scaled (actual calendar) inter-default times, after a linear scaling of time chosen to match the mean of the inter-default time distribution to that of 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 $Z_1, \ldots, Z_n$, Prahl shows that

$$M = \frac{1}{n} \sum_{\{Z_k < C^*\}} \left(1 - \frac{Z_k}{C^*}\right).$$
(9)

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

$$\alpha \quad \simeq \quad 0.189$$

$$\beta \;\simeq\; 0.2427.$$

Using our data, for $n = 240$ default times,

$$
\begin{aligned}
M &= 0.3681 \\
\mu(M) &= \frac{1}{e} - \frac{\alpha}{n} = 0.3671 \\
\sigma(M) &= \frac{\beta}{\sqrt{n}} = 0.0156.
\end{aligned}
$$

Because the test statistic M measured from our data is within one tenth of its standard deviation from the asymptotic mean associated with the null hypothesis of *iid* exponential inter-default times (in the new time scale), this test provides no notable evidence of default clustering in excess of that associated with the default intensities under the doubly stochastic model.[3]

We conducted the same test for inter-arrivals of defaults in calendar time (as opposed to intensity time). The test statistic is $M = 0.4356$, which is 4.38 standard deviations from the mean expected value. Hence, defaults do indeed cluster to a highly statistically significant extent, if one does not first "remove," as we have in our earlier tests, the correlation induced by co-movement of default intensities.

We also report a direct Kolmogorov-Smirnov test of goodness of fit of the exponential distribution of inter-default times in the new time scale. Figure 5 shows the empirical distribution of inter-default times before and after scaling time change by total intensity of defaults, compared to the exponential density implied by the doubly-stochastic model. The associated K-S statistic is 1.8681 (this is $\sqrt{n}$ times the usual D statistic, where n is the number of default arrivals), for a p-value of only 0.002, 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 K-S statistic is 2.0716, with a p-value of 0.0004.)

In summary, despite the rejection of overall fit implied by the K-S and our earlier tests, Prahl's test shows that after conditioning on intensities, defaults are not significantly clustered.

4.4   Calibrating the residual copula correlation

Although Prahl's test indicates no significant degree of clustering after conditioning on default intensities, we have already seen some evidence of a lack of fit of the doubly-stochastic assumption at the measured default intensities.

---

[3]Even if there is no clustering at the level of the economy, it is possible that there may be clustering at the industry level. We implemented Prahls' test, as above, within each sector, classifying firms on the basis of their broad SIC code. We found no evidence of statistically significant clustering, in intensity time, for any sector.
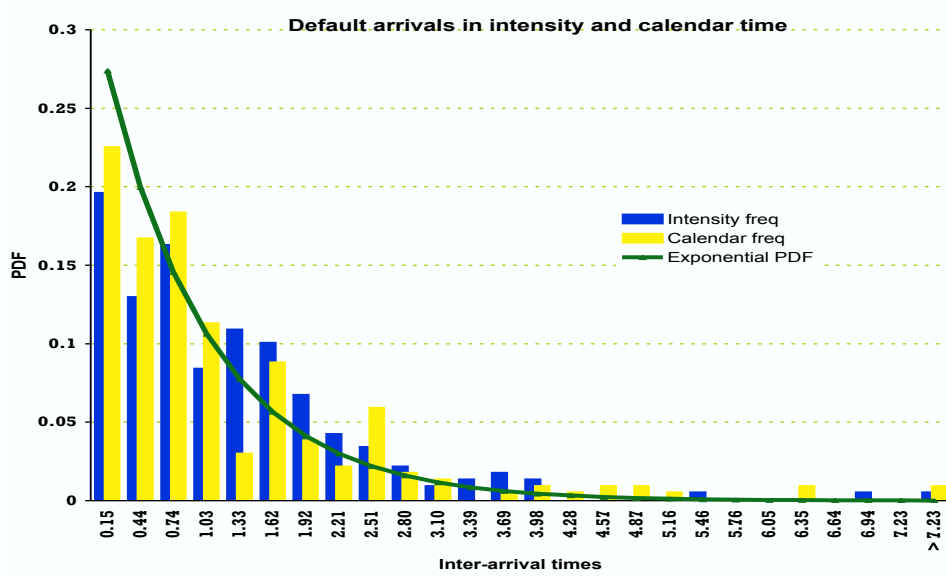
Figure 5: The empirical distribution of inter-default times after scaling time change by total intensity of defaults, compared to the theoretical exponential density implied by the doubly-stochastic model. The distribution of default inter-arrival times is provided both in calendar time and in intensity time.

In this section, we provide a rough estimate of the residual degree of correlation in default arrivals after conditioning on default intensities, based on the intensity-conditional copula approach of Schönbucher and Schubert [2001]. We estimate the amount of copula correlation that must be added, after conditioning on the intensities, to match the excess of 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; we use the industry standard "flat Gaussian copula," used for example to price structured credit products such as collateralized debt obligations.

The magnitude of the calibrated Gaussian copula correlation, in the intensity time scale, is a measure of the degree of correlation in default times that is not captured by co-movement in default intensities. The calibrating algorithm, provided in Appendix C, is applied to cumulative intensity bin sizes $2, 4, 6, 8$, and 10. The results are reported in Table 5.

As anticipated by our prior results, the calibrated residual Gaussian copula correlation r is non-negative in all cases. The largest estimate of r, for bin sizes 8 and 10, is close to 0.02. For the smaller bin sizes of 2 and 4, the estimate of r is close to zero. Overall, these results indicate that, at least by this simple metric, default arrivals are slightly more correlated than suggested by the co-movement across firms of our estimated default intensities.

Table 5: 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 Appendix C. 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	3.62	**3.83**	4.09	4.22	4.43	4.52
4	6.71	**6.58**	7.01	7.32	7.62	7.86
6	10.00	8.97	**9.80**	10.38	10.84	11.34
8	12.75	11.49	12.44	**12.99**	13.76	14.77
10	16.00	13.82	14.69	**15.86**	16.69	17.39

5 Tests for Missing Default Covariates

We have documented violations of various degrees of severity of the joint hypothesis of correctly specified default probabilities and the doubly stochastic property. We now investigate a potential cause of these violations in the form of missing covariates in the PD default-prediction model, a logit-based model that uses only firm-specific covariates. In particular, the underlying Moodys PD 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.

5.1 Testing for Independent Increments

Although all of the above tests depend to some extent on the independent-increments property of Poisson processes, we will 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 $X_1, X_2, \ldots$ in successive bins are independent and identically distributed. We test for independence by estimating an auto-regressive model for $X_1, X_2, \ldots$, under which

$$X_k = A + BX_{k-1} + \epsilon_k, \tag{10}$$

Table 6: Estimates of the auto-regressive model (10) of defaults in successive bins, for a range of bin sizes (t-statistics are shown parenthetically).

Bin Size	No. of Bins	A (t_A)	B (t_B)	R^2
2	118	1.73 (7.66)	0.16 (1.72)	0.03
4	59	2.72 (4.83)	0.34 (2.73)	0.12
6	39	4.20 (3.97)	0.32 (2.01)	0.10
8	29	6.68 (3.83)	0.19 (0.96)	0.03
10	24	6.09 (2.75)	0.39 (1.93)	0.15

for coefficients A and B, and for iid innovations $\epsilon_1, \epsilon_2 \ldots$. Under the joint hypothesis of correctly specified default intensities and the doubly-stochastic property, $A = c$, $B = 0$, and $\epsilon_1, \epsilon_2 \ldots$ are iid de-meaned Poisson random variables. A significantly positive estimate for the auto-regressive coefficient B would be evidence of a failure of the null hypothesis. Serial correlation for small bin sizes could lead, moreover, to fat tails of the distribution of number of defaults in larger bin sizes, and thus could be responsible for the rejections of the Poisson distribution in the larger bins that we reported earlier. Such a result could reflect missing covariates, causing an appearance of default clustering in excess of that implied by the doubly-stochastic property, even if in fact the true multi-firm model of default times is doubly stochastic. For example, if a business-cycle covariate should be included, but is not, and if the 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 6 presents the results of this autocorrelation analysis. The estimated AR(1) coefficient B is always positive, with t statistics indicating a degree of significance that varies with the bin size.

5.2   Macro-economic 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 in default correlation as do unobservable covariates.

Prior work by Lo [1986], Lennox [1999], McDonald and Van de Gucht [1999], and

Duffie and Wang [2003] suggests that macroeconomic performance is an important explanatory variable in default prediction. Without controlling for firm-specific default covariates, for example, industrial production is a highly significant covariate (McDonald and de Gucht [1999]). What matters in our empirical setting is whether macro-economic covariates have significant explanatory power for default probabilities *after* controlling for firm-specific covariates. Among the prior studies, Duffie and Wang included distance to default, the key covariate in Moody's PD model, and found significant additional dependence of default intensities on U.S. personal income growth, for the U.S. machinery and instruments sector for 1971 to 2001.

In this section, we explore the potential role of missing macroeconomic default covariates. In particular, we examine *(i)* whether the inclusion of these macroeconomic variables helps explain default arrivals after controlling for the default covariates used to estimate our default intensities, and if so, *(ii)* whether these variables can potentially explain the violation of the doubly-stochastic assumption. We find that industrial production and GDP growth rates do offer some explanatory power.

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 should therefore be uncorrelated with any variable in the information set available prior to the formation of the k-th bin. Consider the regression,

$$Y_k = \alpha + \beta_1 GDP_k + \beta_2 IP_k + \epsilon_k, \tag{11}$$

where GDP_k and IP_k are the growth rates of US gross domestic product and industrial production observed in the quarter and month, respectively, immediately prior to the beginning of the k-th bin. Under the null hypothesis of correct specification of the default intensities, the coefficients α, β_1, and β_2 are in theory equal to zero. Table 7 reports estimated regression results for a range of bin sizes.

We report the results for the multiple regression as well as for each of the variables separately. For all bin sizes, industrial production enters the regression with sufficient significance to warrant its consideration as an additional explanatory variable in the default intensity model. For bin sizes of both 2 and 10, the coefficient for GDP growth rate is also significant at the 99% confidence level. For each of the bins, the signs of the coefficients in the regressions are negative as one would expect under a mis-specification of missing macroeconomic variables. That is, significantly more than the number of defaults predicted by the PD model occur when GDP and industrial production growth rates are lower than normal. Moreover, the explanatory power of the regression is particularly high for the larger bin sizes, again consistent with the hypothesis that the larger bins are especially affected by a missing macroeconomic covariate. Overall, there is evidence of mis-specification. Given the persistence of macroeconomic variables across time, these missing covariates may be responsible for the presence of the apparent auto-correlation in $X_1, X_2, \ldots$ that we reported in Section 5.1. We verified this by including lagged values of deviations X_{k-1} in the regression. We notice that for all bin sizes, both the intercept and the lagged value

Table 7: Macroeconomic Variables and Default Intensities. For each bin size c, OLS-estimated coefficients are reported for 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	118	0.49	-14.13		5.65
		(2.71)	(-3.03)		
		0.18		-56.03	4.39
		(1.13)		(-2.29)	
		0.55	-11.99	-44.29	8.27
		(3.06)	(-2.44)	(-1.77)	
4	59	0.76	-21.16		5.74
		(1.53)	(-1.76)		
		0.35		-101.53	8.06
		(1.22)		(-2.71)	
		0.90	-17.12	-88.66	11.69
		(1.86)	(-1.48)	(-2.51)	
6	39	1.26	-34.38		7.76
		(1.45)	(-1.58)		
		0.53		-165.04	9.46
		(0.98)		(-2.40)	
		1.43	-27.53	-139.34	14.21
		(1.83)	(-1.32)	(-2.14)	
8	29	1.06	-28.33		2.93
		(0.74)	(-0.76)		
		1.04		-285.63	17.40
		(1.43)		(-2.69)	
		1.71	-21.02	-276.08	18.99
		(1.43)	(-0.63)	(-2.64)	
10	24	2.62	-71.47		19.00
		(1.76)	(-1.97)		
		1.59		-388.33	24.55
		(1.39)		(-2.49)	
		3.83	-64.67	-360.47	39.98
		(4.01)	(-2.36)	(-2.59)	

X_{k-1} are not significant, implying that the use of macro-economic covariates appears to account for some of the variation in excess defaults observed in the data.

It is also useful to examine the role of macro-economic factors when defaults are much higher than expected. Table 8 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) examined previously in Table 3 are correlated with GDP and industrial production growth rates. We report two sets of regressions, the first set based on the prior period's macroeconomic variables and the second set based on the growth rates observed within the bin period.[4]

We report results for those bin sizes, 2 and 4, for which we have a reasonable number of observations. For each of these bins, previous or current period industrial production is significant at typical confidence levels, and has the (negative) sign consistent with the presence of mis-specification by failure to include macroeconomic performance variables in prediction of default. GDP growth rate is not significant in the single variable regressions.

6   Conclusions and Discussion

Defaults cluster in time both because firms' default intensity processes are correlated and also perhaps because, even after conditioning on these intensities, default occurrence is correlated through additional channels such as contagion and 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. The doubly-stochastic assumption forms the current basis of risk management practice, yet to date, no test of its validity has been undertaken. This paper makes the following contributions:

1. We introduce a time-change technique that reduces the process of cumulative defaults to a standard Poisson process under the doubly-stochastic hypothesis. Based on this, we provide newly developed tests of the joint hypothesis that default intensities are correctly measured and that the doubly-stochastic property holds. We are particularly interested in whether defaults are indeed independent after conditioning on intensities.

2. Using various tests, we reject (at traditional confidence levels) the null of correctly measured intensities and the doubly-stochastic property. The Fisher dispersion test and our upper quartile test address the question by each bin size, whereas the test by Prahl (1999) enables a test across all bin windows.

3. There is only mild evidence, however, that defaults are more tightly clustered in time than would be suggested by simultaneous increases in their default

---

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

Table 8: Upper-tail regressions. For each bin size c, OLS-estimated coefficients are shown for 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	40	-0.10 (-0.75)	3.47 (0.84)		1.26
		0.11 (0.64)		-60.21 (-2.50)	14.24
		-0.07 (-0.55)	8.56 (1.99)	-75.57 (-2.79)	20.98
4	17	0.67 (2.52)	-8.30 (-1.48)		9.64
		0.49 (2.24)		-35.11 (-1.16)	4.67
		0.66 (2.37)	-7.13 (-1.15)	-18.57 (-0.57)	10.76

Bin Size	K	Intercept	Current Bin GDP	Current Bin IP	R^2 (%)
2	40	-0.09 (-0.63)	3.13 (0.59)		0.97
		0.03 (0.21)		-23.60 (-1.21)	2.05
		-0.16 (-0.95)	-10.32 (1.19)	-57.99 (-1.58)	8.24
4	17	0.55 (2.42)	-4.26 (-0.68)		2.02
		0.51 (2.61)		-28.05 (-2.14)	10.00
		0.43 (1.96)	3.82 (0.45)	-35.63 (-1.92)	10.90

intensities. Introducing a measure of residual Gaussian copula correlation, after controlling for default intensities, we find that the default clustering in our data can be matched by injecting as little as 2% extra copula correlation. Thus, as a practical matter, simulation models that rely on the doubly stochastic assumption may be useful for certain types of corporate debt risk management applications.

4. We explore further whether default intensities are well-specified. We present evidence that business-cycle covariates offer some explanatory power for default prediction after controlling for standard firm-specific covariates.

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 bank 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 have tested 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 Saunders [2003], and Kashyap and Stein [2004]).

Appendices

A   Moody's Data on Defaults

This appendix provides some details of the creation of the data set used in this paper. Our source of data are two separate databases, one containing default probabilities and the other containing information of defaults. For the empirical work in this paper, we need to account for all the defaults that occur over our sample of firms for which we have PDs. Below, we describe how we link the two databases, and create a clean set of data for our analysis.

In its default database, Moody's records 626 US and Canadian defaults of non-financial firms in the period 1/87 to 10/2000. A few firms default twice over this time period (Grand Union defaulted three times). The defaults in the database are indexed by Moody's Issuer Number (MIN). Although some of these firms are linked to a Cusip or a Bloomberg ticker, many of the firms do not have a link to any external identifier. However, the name of the defaulted firm is provided, as well as some information regarding the nature of default. Moody's database of default probabilities is created using accounting and equity price data, and is limited to firms that had available data in the sample period. Our sample period is January 1987 to October 2000. This data is indexed by the Gvkey.

The defaulted firms that have a Cusip are matched to the PD database using the Gvkey-Cusip link of the combined Compustat-CRSP database. For the remaining firms, we do a manual match using the company name. Some of these firms do not have Gvkeys because they are either subsidiaries, or related to the primary public firm that has defaulted. For example, on 7 April 1987, Texaco Capital, Texaco Capital N.V., Texaco Corporation and Texaco Operations Europe are listed as four separate defaults. Of these, only Texaco Corporation is counted in our sample.

For our empirical investigation, we focus only on defaults among the 1,990 rated firms that are represented in our PD database (firms that have accounting and price information over this period *and* were rated over the sample period of 1987-2000). Of these firms, we count a total number of 241 incidents of defaults over 216,859 firm-months of data. This comprises our final PD and default dataset.

B   Estimation of Default Intensities from PDs

This appendix provides the algorithm for our iterative estimator of default intensities.

1. First, we obtain starting coefficient estimates values from the regression, for $h = 1/12$,

$$s_{t+h}(1) - s_t(1) = \alpha + \beta s_t(1) + e_t, \tag{12}$$

   where α and β are the ordinary-least-squares (OLS) estimators and e_t denotes the residual. From this regression, we get initial estimates of the three parameters as:

$$k = -\frac{\beta}{h} \tag{13}$$

$$\theta = -\frac{\alpha}{\beta} \tag{14}$$

$$\sigma = \frac{V(e)}{\sqrt{\theta h}}, \tag{15}$$

where $V(e_t)$ denotes the sample standard deviation of the residual e_t.

2. Given starting values of $\{k, \theta, \sigma\}$, we obtain an initial estimate of the default intensity λ_t, for each observation time t, using equation (7).

3. Next, we estimate by OLS,

$$\lambda_{t+h} - \lambda_t = a + b\lambda_t + w_t. \tag{16}$$

New parameter estimates are then given by

$$\hat{k} = -\frac{b}{h}, \quad \hat{\theta} = -\frac{a}{b}, \quad \hat{\sigma} = V\left(\frac{w_t}{\sqrt{h\lambda_t}}\right), \tag{17}$$

where, again, $V(\cdot)$ denotes sample standard deviation.[5]

4. Given these updated estimates of the parameters $\{k, \theta, \sigma\}$, we return to Steps 2 and 3, and iterate to numerical convergence.

5. For firms with very few observations (usually < 48), or with minimal variation in PDs, convergence is not guaranteed; in these cases, PDs are converted to intensities using a simpler (constant over 1 year) intensities model.

C   Residual Guassian Copula Correlation

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

1. For name i, cumulative intensity bin size c, and a given bin k, we calculate the total cumulative intensity $C_i^{c,k}$ for name i, 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.

2. For a flat Gaussian copula with correlation parameter r, let X_i be the standard normal for name i in this bin, let $F(X_i)$ be the standard normal CDF, and set $U_i = F(X_i)$ to be the associated uniform. We set $\text{corr}(X_i, X_j) = r$ for all (i, j). Under the conditional-copula model, name i defaults in this bin if $U_i > \exp[-C_i^{c,k}]$, that is, if the uniform exceeds the survival probability.

3. For each of 5,000 independent scenarios, we draw one of the bins at random (equally likely), draw $X(i)$ for each name i, with correlation r, and draw a default for each i or not according to Step 2. Note that name i cannot default if $C_i^{c,k} = 0$.

---

[5]In the current version of our results, we use $V(w_t/\sqrt{\hat{\theta}h})$ in place of the sample standard deviation shown in (17), although our tests indicate that this causes minimal distortion in the estimated intensities.

4. From Step 3, we calculate the empirical distribution of the number of defaults per bin. If the empirical distribution (see Table 3) has the same upper-quartile mean number of defaults per bin as the actual mean upper quartile that we calculated in the simulation, then r is calibrated.

References

[2003] Allen, L. and A. Saunders (2003) "A Survey of Cyclical Effects in Credit Risk Measurement Models," BIS Working Paper 126, Basel Switzerland.

[1986] Billingsley, Patrick (1986). *Probability and Measure, Second Edition*, New York: Wiley.

[2002] Cathcart, L. and L. El Jahel (2002) "Defaultable Bonds and Default Correlation," Working Paper, Imperial College.

[1954] Cochran, W.G. (1954) "Some Methods of Strengthening χ^2 Tests," *Biometrics* v10, 417-451.

[2003] Collin-Dufresne, Pierre, Goldstein, Robert, and Jean Helwege (2003) "Is Credit Event Risk Priced? Modeling Contagion via the Updating of Beliefs," Working Paper, Haas School, University of California, Berkeley.

[2004] Collin-Dufresne, Pierre, Goldstein, Robert, and Julien Huggonier (2004) "A General Formula for Valuing Defaultable Securities," *Econometrica*, v72, 1377-1407.

[1985] Cox, John, Jon Ingersoll, and Steven Ross (1985) "A Theory of the Term Structure of Interest Rates," *Econometrica* v53, 385-407.

[2001] Das, Sanjiv., Laurence Freed, Gary Geng, and Nikunj Kapadia (2001). "Correlated Default Risk," working paper, Santa Clara University.

[2001] Davis, Mark, and Violet Lo (2001) "Infectious Default," *Quantitative Finance*, v1, 382-287.

[2002] deServigny, Arnaud, and Olivier Renault (2002) "Default Correlation: Empirical Evidence," Working Paper, Standard and Poors.

[2003] Duffie, Darrell., and Ke Wang (2003). "Multi-Period Corporate Failure Prediction with Stochastic Covariates," working paper, Stanford University.

[2004] Giesecke, Kay (2004) "Correlated Default with Incomplete Information," *Journal of Banking and Finance*, v28, 1521-1545.

[2003] Gordy, Michael (2003) "A Risk-Factor Model Foundation for Ratings-Based Capital Rules," *Journal of Financial Intermediation*, v12, 199-232.

[2001] Jarrow, Robert., and Fan Yu (2001). "Counterparty Risk and the Pricing of Defaultable Securities," *Journal of Finance* v56, 1765-1800.

[2005] Jarrow, Robert., David Lando, and Fan Yu (2005). "Default Risk and Diversification: Theory and Applications," *Mathematical Finance*, v15(1), 1-26.

[2004] Kashyap, Anil., and Jeremy Stein (2004) "Cyclical Implications of the Basel-II Capital Standards," Working Paper, Graduate School of Business, University of Chicago.

[1999] Kusuoka, Shigeo (1999) "A Remark on Default Risk Models," *Advances in Mathematical Economics* v1, 69-82.

[1994] Lando, David (1994). "Three essays on contingent claims pricing," Ph.D. thesis, Cornell University.

[1998] Lando, David (1998). "On rating transition analysis and correlation," *Credit Derivatives: Applications for Risk Management, Investment and Portfolio Optimization*, Risk Publications, 147-155.

[2002] Lando, D. and T. Skødeberg (2002). Analyzing Rating Transitions and Rating Drift with Continuous Observations. *Journal of Banking and Finance* v26, 423-444.

[1992] Lang, Larry and Rene Stulz (1992), "Contagion and competitive intra-industry effects of bankruptcy announcements, *Journal of Financial Economics* v32, 45-60.

[1999] Lennox, C. (1999) "Identifying Failing Companies: A reevaluation of the Logit, Probit, and DA Approaches," *Journal of Economics and Business*, v51, 347-364.

[1986] Lo, Andrew (1986) "Logit versus Discriminant Analysis: Specification Test and Application to Corporate Bankruptcies," *Journal of Econometrics*, v31, 151-178.

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

[1999] McDonald, Cynthia and Linda M. Van de Gucht (1999). "High Yield Bond Default and Call Risks." *Review of Economics and Statistics*, v81, 409-419.

[1974] Merton, Robert C. (1974). "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *The Journal of Finance*, v29, 449-470.

[2004] Norden, Lars and Martin Weber (2004). "Informational Efficiency of Credit Default Swap and Stock markets: The Impact of Credit Rating Announcements," Working Paper, Department of Banking and Finance, University of Mannheim.

[1999] Prahl, Juergen. (1999) "A Fast Unbinned Test on Event Clustering in Poisson Processes," Working Paper, University of Hamburg, submitted to *Astronomy and Astrophysics.*

[2003] Protter, Philip (2003). *Stochastic Integration and Differential Equations, Second Edition* (New York: Springer).

[2003] Schönbucher, Philipp (2003) "Information Driven Default Contagion," Working Paper, ETH, Zurich.

[2001] Schönbucher, Philipp, and D. Schubert (2001) "Copula Dependent Default Risk in Intensity Models," Working Paper, Bonn University.

[2001] Shumway, Tyler (2001) "Forecasting Bankruptcy More Accurately: A Simple Hazard Model," *Journal of Business* v74, 101-124.

[2000] Sobehart, J., R. 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.

[2000] Sobehart, J.R., Keenan, S.C. and Stein, R.M. (2000), "Benchmarking Quantitative Default Risk Models: A Validation Methodology," *Technical Report*, Moody's Risk Management Services.

[2003] Terentyev, S. (2004) "Asymmetric Counterparty Relations in Default Modeling," Working Paper, Department of Statistics, Stanford University.

[1987] Vasicek, Oldrich (1987) "Probability of Loss on Loan Portfolio," Working Paper, KMV Corporation.

[2003] Yu, Fan (2003) "Default Correlation in Reduced Form Models," University of California, Irvine.

[2004] Yu, Fan (2004) "Accounting Transparency and the Term Structure of Credit Spreads," University of California, Irvine.

[2004] Zhang, Gaiyan (2004) "Intra-Industry Credit Contagion: Evidence from the Credit Default Swap Market and the Stock Market," Working paper, University of California, Irvine.

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