

FDIC Center for Financial Research
Working Paper

No. 2004-04

Common Failings: How Corporate Defaults are Correlated

Sanjiv R. Das
Darrell Duffie
Nikunj Kapadia

September 2004



Federal Deposit Insurance Corporation • Center for Financial Research

Common Failings: How Corporate Defaults are Correlated¹

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.

July 2004

¹This research is supported by a fellowship grant from the Federal Deposit Insurance Corporation (FDIC). We received useful comments from participants at the FDIC CFR conference, and the QFAFAFEW Group San Francisco. We are grateful to Mark Flannery, Robert Jarrow, Edward Kane, Paul Kupiec, Dan Nuxoll, George Pennacchi, Louis Scott, Philip Shively, and Haluk Unal for their suggestions. We are also grateful to Moody's Investors Services and Gifford Fong Associates for data and research support for this paper. The first author is grateful for the support of a Breetwor Fellowship.

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 correlation of their default intensity processes, for example through dependence on common or correlated observable risk factors. Our tests do not require assumptions on the correlation structure 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 at most weak evidence of default clustering in excess of that implied by the doubly-stochastic model and correlation of the firm-specific default covariates.

1 Introduction

Why do corporate defaults cluster in time? Several explanations have been explored in the literature. First, firms may be exposed to common or correlated risk factors whose co-movements cause correlated changes in the conditional probabilities of defaults 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 correlation in default intensities, that is, the first channel on its own, is sufficient to account for the degree of default clustering 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 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-subsiary 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 the correlation of conditional corporate default probabilities (see, for example, Das, Freed, Geng and Kapadia, [2001]), there is relatively little evidence regarding the presence of highly clustered defaults. Lucas [1995] and deServigny and Renault [2002] have estimated historical average one-year default correlations, but do not address the issue of clustering. Collin-Dufresne, Goldstein, and Helwege [2003] 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 premia¹ of other firms, as envisioned by Kusuoka [1999]. Both effects could be at play. Collin-Dufresne, Goldstein, and Helwege do not disentangle these two channels for default-induced widening of yield spreads. Explicitly consider-

¹Collin-Dufresne, Goldstein, and Huggonier [2002] 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 [2000], Giesecke [2002], Kusuoka [1999], Schönbucher and Schubert [2001], Teremtyev [2003], Yu [2003], and Zhou [2001].

ing a failure of the doubly-stochastic hypothesis, Collin-Dufresne, Goldstein, and Helwege [2003], Giesecke [2002], Jarrow and Yu [2001], and Schönbucher [2004] explore learning-from-default interpretations, based on the statistical modeling of frailty, under which default intensities include 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.

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 then exploit the following 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. 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, fixing any scalar $c > 0$, if we define 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), the doubly-stochastic assumption implies 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, \dots, X_k, \dots of defaults in successive time bins of a given accumulated intensity c , with the theoretical Poisson distribution implied by the doubly-stochastic model.
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 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.
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.
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 is positive, could lead to default clustering in excess of that associated with the doubly-stochastic assumption.

An appealing feature of these tests is that they are not dependent on the correlation structure of default intensities, 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. 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. 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 that were observed before bin k begins. It is possible that missing covariates, rather than a failure of the doubly-stochastic property, is responsible for the relatively poor fit of the data to the joint hypothesis that we test.

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 default intensity accumulation. 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 addresses the question of independence of increments of the time-changed process governing default arrival. In Section 6 we test our default intensity data for missing macroeconomic covariates. (This test does not depend on the doubly-stochastic property.) Section 7 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 (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 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 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 default intensities 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 $N = (N_1, \dots, N_n)$ 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 will 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 $U(t) = \int_0^t \sum_{i=1}^n \lambda_i(u) 1_{\{\tau_i > u\}} du$ 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,

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

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

In order to check the independence of $Z_1, Z_2 \dots$, 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 shall test are based on the following characterization.

Poisson property: For any $c > 0$, the random variables

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

are iid Poisson with parameter c .

That is, if we divide our sample period into “bins” that each have an equal cumulative aggregate intensity of c , then we can test the doubly stochastic assumption by testing whether the numbers of defaults in successive bins are independent Poisson random variables with common parameter c .

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, covering 1,990 firms, 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. As the default probabilities and defaults are from separate databases, much of the matching is done manually by matching company names. Given that the default probabilities have been computed by fitting to observed defaults, we can verify the completeness of the matching by comparing the mean default rate implied by the default probabilities to the actual number of defaults. We discuss this in more detail in our analysis below. Appendix A provides further details on the construction of the database.




Figure 1: Cross-sectional sample mean of one-year conditional default probabilities, and the number of firms covered, January 1987 to October 2000.

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 Intensities

In order to test the doubly-stochastic assumption using the new-time-scale Poisson process described in the Proposition above, we estimate default intensities, firm by firm, from the PD data on 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

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

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

$$s_t(T) = E \left[\exp \left(- \int_t^{t+T} \lambda_u du \right) \mid \lambda_t \right]. \quad (2)$$

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

$$s_t(T) = A(T) \exp [-\lambda_t B(T)], \quad (3)$$

where

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

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

$$\gamma = \sqrt{k^2 + 2\sigma^2}. \quad (6)$$

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

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

Our PD data are 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. Figure 2 shows the total of the estimated intensities of all firms, as well as the monthly arrivals of defaults.

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. Fixing time bins containing c units of accumulative default intensity, we then construct calendar times t_0, t_1, t_2, \dots such $t_0 = 0$ and $U(t_i) - U(t_{i-1}) = c$, and let $X_k = \sum_{i=1}^n 1_{\{t_k \leq \tau(i) < t_{k+1}\}}$ denote the number of defaults in the k -th time bin. Figure 3 illustrates the 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 differ vary slightly from the integer bin sizes shown because of daily granularity in 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 PD data are reasonably well estimated, however this is to be expected given the within-sample nature of the estimates. For larger bin sizes, the empirical variances are bigger than their theoretical counterparts under the null of correctly specified doubly-stochastic intensity model of defaults.

Figure 4 presents the observed default frequency distribution, and the associated theoretical Poisson distribution, for bin sizes 2 and 8. For bin size larger than 4, there is a tendency for bi-modality (two peaks), as opposed to the uni-modal theoretical Poisson distribution associated with

²Under 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 $3 + c^{-1}$.

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

the hypothesis of doubly-stochastic defaults.

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, \dots, X_K are independent Poisson distributed variables with mean parameter c is Fisher's dispersion test (Cochran [1954]). Under this null,

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

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

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

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 6 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 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 distribution of X_1, \dots, X_K . By Monte Carlo simulation, we generated 10,000 data sets, each consisting of K *iid* Poisson random 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 p -value would be approximately 0.5.

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

The sample p -values are presented in Table 3, and suggest, for larger bin sizes, fatter upper-quartile tails than those of the theoretical Poisson distribution. (That is, our one-sided tests implies rejection for larger bins of the null joint hypothesis, 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 strong rejection of a Poisson-distributed upper-quartile distribution at standard confidence levels.

4.3 Prah’s Test of Clustered Defaults

Fisher’s dispersion and our tailored upper-tail test do not exploit well the information available across all bin sizes. In this section, we apply a test for “bursty” default arrivals due to Prah [1999]. Prah’s test is sensitive to cluster-like deviations from the 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.

Table 4: Selected moments of the distribution of cumulative amount of intensities between successive default times. Under the joint hypothesis of doubly-stochastic defaults and correctly measured default intensities, the distribution is exponential.

Moment	Empirical	Exponential
Mean	1.07	1.07
Variance	1.19	1.16
Skewness	2.13	2.00
Kurtosis	7.46	6.00

Prahl's test statistic is based on the fact that, in the new time scale under which default arrivals are those of a Poisson process (with rate parameter 1), the inter-arrival times Z_1, Z_2, \dots are *iid* exponential of mean 1. (Because of data granularity, our mean is slightly larger than 1.) In contrast to the previous tests based on a Poisson number of defaults, this test examines instead the inter-arrival default times, which are exponential under the null. The sample moments of these time-rescaled inter-arrival times are provided in Table 4.

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

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

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

$$\begin{aligned} \alpha &\simeq 0.1839 \\ \beta &\simeq 0.2431. \end{aligned}$$

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 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.

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 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, 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 summary, while the default inter-arrival times (under the intensity-based time change) do not fit the model-implied exponential distribution, Prah's test does not indicate (bursty) default clustering in excess of what would be implied by the doubly-stochastic property and co-movement of the default intensities.

5 Testing for Independent Increments

Although all of the above tests depend in part 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, \dots in successive bins are independent and identically distributed. We test for independence by estimating an auto-regressive model for X_1, X_2, \dots , under which

$$X_k = A + BX_{k-1} + \epsilon_k, \quad (10)$$

for coefficients A and B , and for *iid* innovations $\epsilon_1, \epsilon_2, \dots$. Under the joint hypothesis of correctly specified default intensities and the doubly-stochastic property, $A = c$, $B = 0$, and $\epsilon_1, \epsilon_2, \dots$ are *iid* de-meaned Poisson random variables. A large and positive auto-regressive coefficient B would be evidence of a failure of the null hypothesis. This form of serial correlation,

Table 5: 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 small bin sizes, could generate fat tails of the distribution of number of defaults per bin for larger bin sizes, and thus could be responsible for the apparent rejections of the Poisson distribution in the larger bins that we reported earlier. This could perhaps be evidence of mis-specification of the underlying PD model, for example through missing covariates for default prediction. Such a result could reflect serial correlation of 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.

Table 5 presents the results of this autocorrelation analysis. The AR(1) coefficient B is always positive, and sometimes significantly larger than zero at traditional confidence levels. Indeed, the auto-regressive coefficient B tends to be “more significant” for small bin sizes, consistent with an interpretation of the earlier failure of the test for Poisson distributed upper tails in large bins as potentially due to a failure of the independence assumption for small bins, and perhaps due to mis-specification of the underlying intensity model through a failure to include default covariates, such as macroeconomic variables, that have some persistence over time.

6 Tests for Missing Default Covariates

Our lack of support for the joint hypothesis of correctly specified default probabilities and the doubly stochastic property might be related to missing covariates in the PD default-prediction model, a logit-based model that uses only firm-specific covariates. In particular, this default prediction model may be missing covariates that are common to many firms, and would therefore reveal additional default time correlation under the doubly-stochastic model.

Prior work by Shumway [2001], Lennox [1999], Lo [1986], and Duffie and Wang [2003] indeed suggests that macroeconomic performance is an important explanatory variable in default prediction. Among these 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 two macroeconomic variables, United States G.D.P. growth rate (GDP) and personal income growth rate (PI). In particular, we examine (i) whether the inclusion of these macroeconomic variables helps predict defaults in addition to the default intensities, and if so, (ii) whether these variables can potentially explain the apparent failure of the doubly-stochastic assumption.

We first examine whether the default intensities, based on Moody's default probability estimates, indeed indicate mis-specification from lack of a macroeconomic covariate. 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,

$$Y_k = \alpha + \beta_1 PI_k + \beta_2 GDP_k + \epsilon_k, \quad (11)$$

where PI_k and GDP_k are the growth rates of U.S. personal income and U.S. growth in gross domestic product observed in the quarter immediately prior to the beginning of the k -th bin. Under the null hypothesis of correct specification, *whether or not the doubly-stochastic assumption is satisfied*, the coefficients β_1 and β_2 are in theory equal to zero. Table 6 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 bin sizes of both 2 and 10, the coefficient for GDP growth rate is significant at the 99% level. For each of the bins, the signs of the coefficients in the single equation 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 personal income growth rates are lower than normal. Overall, there appears to be mild evidence of mis-specification. Given the persistence of macroeconomic variables across time, these missing covariates may be partly responsible for the presence of the apparent auto-correlation in X_1, X_2, \dots that we reported in Section 5. One may therefore wish to consider whether any excess clustering of defaults (beyond that implied by the doubly-stochastic property) is related to this potential mis-specification of the default intensity processes.

Table 7 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 personal income and GDP 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.³

As for these upper-tail-size regressions, the estimated coefficients for PI and GDP based on the prior period's growth rates are not significant at typical confidence levels. The coefficient for current-period PI for bin-size 4, however, is significant at typical confidence levels, and has a sign consistent with the presence of mis-specification by failure to include macroeconomic performance variables in prediction of default. Overall, the extent of significance, as well as the fit (as judged from the R^2) are weakly suggestive of macroeconomic underpinnings of conditional dependence.

³The within-period growth rates are computed by compounding over the daily growth rates that are consistent with the reported quarterly growth rates.

Table 6: 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 personal income growth rate and the GDP growth rate. 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	Personal Income	GDP	R^2 (%)
2	118	0.10 (0.58)	-11.15 (-0.65)		0.37
		0.49 (2.71)		-14.13 (-3.03)	5.65
		0.46 (2.45)	27.37 (1.33)	-19.74 (-3.35)	6.98
4	59	0.25 (0.56)	-25.39 (-0.63)		0.93
		0.76 (1.53)		-21.16 (-1.76)	5.74
		0.74 (1.43)	18.55 (0.41)	-24.85 (-1.75)	6.07
6	39	0.53 (0.69)	-56.34 (-0.73)		2.14
		1.26 (1.45)		-34.38 (-1.58)	7.76
		1.24 (1.37)	6.37 (0.08)	-35.53 (-1.49)	7.78
8	29	1.06 (0.74)	-28.35 (-0.76)		2.93
		0.13 (0.12)		-2.65 (-0.03)	0.00
		0.97 (0.67)	65.44 (0.59)	-42.16 (-0.91)	4.28
10	24	1.15 (0.78)	-127.60 (-0.92)		6.10
		2.62 (1.76)		-71.44 (-1.97)	18.99
		2.57 (1.61)	44.67 (0.37)	-81.00 (-2.53)	19.39

Table 7: 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 personal income (PI) and GDP 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 PI	Previous Qtr GDP	R^2 (%)
2	40	-0.05 (-0.71)	5.23 (0.49)		0.35
		-0.10 (-0.75)		3.47 (0.84)	1.26
		-0.09 (-0.71)	-7.43 (-0.48)	5.54 (0.89)	1.52
4	17	0.64 (2.42)	-25.07 (-1.53)		10.78
		0.67 (2.52)		-8.30 (-1.48)	9.64
		0.67 (2.42)	-16.99 (-0.67)	-3.52 (-0.39)	11.39
Bin Size	K	Intercept	Current Bin PI	Current Bin GDP	R^2 (%)
2	40	-0.00 (-0.03)	-0.94 (-0.09)		0.01
		-0.09 (-0.63)		3.13 (0.59)	0.97
		-0.08 (-0.58)	-20.53 (-0.65)	8.66 (0.71)	2.90
4	17	0.69 (2.86)	-32.96 (-2.85)		15.71
		0.55 (2.42)		-4.26 (-0.68)	2.02
		0.55 (2.56)	-72.94 (-2.70)	17.45 (1.60)	26.50

7 Concluding Comments

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. The latter channels are not admitted in a doubly-stochastic setting. By a time change that reduces the process of cumulative defaults to a standard Poisson process, we provide test statistics of the joint hypothesis that default intensities are correctly measured and the doubly-stochastic property. We are particularly interested in whether defaults are indeed independent given intensities. We believe this to be the first such empirical test. For several types of tests, we reject (at traditional confidence levels) the null of correctly measured intensities and the doubly-stochastic property, at traditional confidence levels. We present some evidence, however, of potential mis-specification of these default probability estimates, in that they do not include business-cycle covariates that may offer some predictive power for default above and beyond the role of firm-specific covariates. Moreover, there is at best weak evidence that defaults are more tightly clustered in time than would be suggested by common increases in their default intensities.

The economic impact of a failure of the doubly-stochastic property for the risk management of credit portfolios is of critical interest to investors and bank regulators. For example, the level of economic capital necessary to support levered portfolios at high confidence levels is heavily dependent on the degree to which this often-assumed property actually applies. This is especially the case in light of the upcoming changes in bank capital regulations under the proposed Basel II accord on regulatory capital (see Gordy [2003], Allen and Saunders [2003], and Kayshap and Stein [2004]). By focusing attention on a crucial practical aspect of doubly-stochastic models, and by providing a statistical test of default clustering in credit-sensitive asset portfolios, we hope to have provided information that is useful in the design of models for risk management.

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 datasets, and the set of defaults that results from our procedures.

In its default database, Moody's records 628 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). Moody's records defaults only for firms that it has rated at some point in the firm's history. 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. After both these matches, many firms remain without a Gvkey. 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.

The number of defaults that are available for our empirical work is further reduced as many firms were not rated by Moody's according to our PD database over the period 1987-2000. The final dataset, corresponding to the default of firms that are present in the PD database over our sample period, contains 241 incidents of defaults among a total of 1,990 firms and over 216,859 firm-months of data.

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 $h = 1/12$,

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

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:

$$k = -\frac{\beta}{h} \quad (13)$$

$$\theta = -\frac{\alpha}{\beta} \quad (14)$$

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

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,

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

New parameter estimates are then given by

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

where, again, $V(\cdot)$ denotes sample standard deviation.⁴

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

⁴In 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.

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 Jahl (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.
- [2002] Collin-Dufresne, Pierre, Goldstein, Robert, and Julien Huggonier (2002) "A General Formula for Valuing Defaultable Securities," forthcoming, *Econometrica*.
- [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.
- [2000] Davis, Mark, and Violet Lo (2000) "Infectious Default," Working Paper, Imperial College.
- [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.

- [2002] Giesecke, Kay (2002) "Correlated Default with Incomplete Information," Working Paper, Humboldt University, Berlin.
- [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* 56, 1765-1800.
- [1999] Jarrow, Robert., David Lando, and Fan Yu (1999). "Default Risk and Diversification: Theory and Applications," working paper, Cornell University.
- [2004] Kayshap, 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.
- [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.

- [1974] Merton, Robert C. (1974). "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *The Journal of Finance*, v29, 449-470.
- [1999] Prah, J. (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).
- [2004] Schönbucher, Philipp (2004) "Frailty Models, Contagion, and Information Effects," 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," Working Paper, U.C. Irvine.
- [2004] Yu, Fan (2004) "Accounting Transparency and the Term Structure of Credit Spreads," Working Paper, U.C. Irvine.

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




Figure 2: Aggregate (across firms) default intensities and firm defaults from 1987-2000.




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.





Figure 4: Comparison of the empirical and Poisson distributions of defaults for bin sizes 2 and 8.




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.