

Correlated Default Risk

SANJIV R. DAS, LAURENCE FREED, GARY GENG,
AND NIKUNJ KAPADIA

SANJIV R. DAS
is a professor at Santa Clara University in Santa Clara, CA.
srdas@scu.edu

LAURENCE FREED
is a managing director at Bear Stearns Asset Management in New York, NY.
lfreed@bear.com

GARY GENG
is a risk strategist at Amaranth Group, Inc. in Greenwich, CT.
ggen@amaranthllc.com

NIKUNJ KAPADIA
is an associate professor at University of Massachusetts in Amherst, MA.
nkapadia@som.umass.edu

Recently, an unusually high number of firms in the economy defaulted, with the default rate for Moody's-rated speculative-grade issuers reaching as high as 10.2% in 2001. In their annual review, Moody's summarized these credit events as follows, "Record defaults—unmatched in number and dollar volume since the Great Depression—have culminated in the bankruptcies of well-known firms whose rapid collapse caught investors by surprise."¹ What factors cause the economy-wide default rate to change over time, and why does it vary as much as it does? In this article, we investigate the likelihood of joint default across firms in the economy by examining the co-variation of individual firms' default probabilities.² This, in turn, provides insight into how, and why, the economy-wide default rate varies over time.

We provide a comprehensive empirical investigation of how default probabilities co-vary using a database of issuer-level default probabilities for the period 1987–2000. This database provides a unique opportunity to understand how default risk behaves both in the cross-section of firms and in the time-series for almost all U.S. public non-financial firms. More importantly from the standpoint of using the results, the dataset lends *quantitative* expression to this behavior. For instance, our analysis allows us to understand the extent to which the record defaults of 2001 were, indeed, a surprise.

Defaults of firms in the economy will cluster if there are common factors that affect individual firms' default risk. The structural model of Merton [1974]³ identifies two of these factors as the firm's debt ratio and volatility. Co-variation of individual firms' debt ratios or volatilities will result in their default probabilities being correlated. The economy-wide default rate will vary widely if, first, debt ratios or volatilities across firms are correlated, and second, if there is wide variation in debt ratios, volatilities, or their correlations over time. By relating the co-variation in default probabilities to variation in debt, volatility, and their correlations, we also provide an economic understanding of why the economy-wide default rate varies over time. In short, we examine the contribution of the correlation between default probabilities, i.e., the first part of the standard reduced-form structure of doubly stochastic processes, to joint default risk. Recent evidence in Das, Duffie, Kapadia and Saita [2005] shows that after conditioning on default intensities, the residual correlation that may be ascribed to conditional default is of the order of 1 to 5% only. That is, the majority of joint default risk emanates from covariation in default probabilities, the focal point of the investigation in this article.

Specifically, we observe that although both debt levels and firm volatilities change over time, firm volatilities—and their correlations—vary much more than debt levels. In

particular, firm volatilities show steep increases and declines in relation with the business cycle or major economic events. Moreover, the Merton [1974] model suggests that the default probability of a firm is more sensitive to a change in volatility than the debt level. These two observations suggest that the economy-wide default rate will be strongly related to economy-wide volatility, and thus to the state of the economy. This is precisely what we observe. We summarize our findings below.

First, default probabilities of individual firms of a given rating class vary substantially over time. The mean default probability across all firms in the economy more than doubles between December 1993 and December 1999, increasing from a little over 1% to close to 3%. Much of this time-variation is determined by changes in economy-wide volatility.

Second, default correlations of individual firms, like correlations between individual firm volatilities, are not stable over time. The median correlation between a pair of firms is close to zero in the low volatility period of the mid-1990s, but increases to about 25% for Baa/Ba rated firms and is over 30% for higher rated firms in the high volatility period of the late 1990s. The latter correlations are higher than the corresponding asset (and equity) return correlations over that period. The time variation in default correlations follows the same pattern as correlations of individual firm volatilities. Correlations between individual firm volatilities also increase when economy wide volatility is high. In contrast, asset return correlations are relatively stable over time.

Third, we document cross-sectional differences across rating classes. Volatility of firms of the highest rating classes increase the most in times of economic stress and are also the most highly correlated. Correspondingly, these firms show the steepest increases in default probabilities and default correlations.

Besides providing an insight into why defaults across the economy vary, our results have direct risk management and pricing implications. An understanding of the co-variation of default risk is fundamental to portfolio management in the corporate debt markets and required for the pricing of securities such as Collateralized Debt Obligations and basket default swaps whose payoffs are affected by the number of defaults at a portfolio level. It is therefore of some interest to understand how our primary finding—that both default probabilities and their correlations vary with the state of the economy—can be applied in practice. We propose and test a parsimonious statistical

model that accounts for both these observations. The model, in turn, allows us to illustrate the economic impact of our findings.

To account for our observation that default probabilities vary with economic events, we allow the economy-wide default risk to be regime-dependent. We find strong support for a two-regime model, with a high default regime and a low default regime, the former having a mean default level more than twice that of the latter. Moreover, we demonstrate that each regime shows a different correlation structure: default correlations are higher in the high default regime as compared with those in the low default regime. Therefore, systematic variation in joint default risk may be modeled within a simple reduced-form framework,⁴ and such a model then allows us to quantify default risk at a portfolio level.

Using estimates from our model, we illustrate the economic impact of changing economic regimes. We provide two examples. First, we compute the n 'th default probabilities over a basket of firms, and demonstrate that these probabilities vary widely between regimes. For example, the probability of observing a default in a basket of 10 medium-rated (Baa/Ba) firms increases more than three-fold from 4% in the low-default regime to 13.5% in the high-default regime. Second, to get an understanding of the impact of changing economic conditions on defaults across the economy, we compute the distribution of defaults across a portfolio of low-grade firms. This analysis indicates that the (out-of-sample) probability of observing the record defaults in low-grade firms in 2001 may have been as high as 20%.

This article is related to and has implications for other recent work in the literature. First, the empirical results complement the theoretical results of Zhou [2001]. Within the context of a structural model of default, Zhou considers the implications of correlated asset returns on correlations between defaults. Modeling asset correlations has been recommended by the Basel Committee, and is often implemented in practitioner models such as the one by RiskMetrics. Our empirical findings suggest that modeling asset volatilities is even more important for the analysis of joint default risk. Second, our findings provide an economic basis for understanding how default risk is priced. There has been considerable recent work in understanding the dynamics of credit spreads. Campbell and Taksler [2003] find that changes in credit spreads are related to changes in volatility, and our results provide additional evidence of the importance of volatility

in determining default risk. Xiao [2003] observes that credit spreads of high quality bonds are more highly correlated than credit spreads of lower rated bonds, consistent with our finding that default correlations are higher for higher grade firms. Lucas [1995] and De Servigny and Renault [2002] use joint migration to default across rating classes to measure joint default risk. Our use of default probabilities allows us to measure joint default risk within a rating class even when there are no rating transitions.

The empirical evidence of this article focuses on the impact of the common factors that drives default risk. In addition, defaults may also cluster because of contagion-like effects when the default of one firm affects the probability of default of another firm, or when the default of a firm provides information that causes investors to update their estimates of default probabilities for other firms (see, Collin-Dufresne, Goldstein & Helwege [2003], Davis & Lo [2001], Driessen [2002], Giesecke [2002], Jarrow & Yu [2001], Schönbucher [2004], and Yu [2003]). Das, Duffie, Kapadia and Saita [2005] provide evidence that suggests that defaults cluster primarily because of the common factors that affect individual firm's default probabilities.

The rest of this article is as follows. We start by examining the underlying determinants of default probabilities. This provides a basis for understanding the dynamics of both default probabilities and default correlations. Based on our empirical observations, we then provide a statistical model to simultaneously model time-variation in default probabilities and default correlations. An analysis section illustrates the economic importance of our findings. Finally, we provide comments on future work.

DEFAULT PROBABILITIES AND CORRELATIONS

Determinants of Default Probabilities

In order to provide a framework for our empirical investigation, consider the structural model of Merton [1974].⁵ In the model, equity-holders can put the firm to the debt-holders at the face value of debt. As noted in Vassalou and Xing [2004], the probability of default (PD) is equal to $1 - \Phi(DTD)$, where the function $\Phi(\cdot)$ is the standard normal cumulative distribution function, and DTD is the "distance-to-default" defined as,

$$DTD = \frac{\ln(V/D) + (\mu_V - 0.5\sigma_V^2)T}{\sigma_V\sqrt{T}}.$$

Here, V is the firm value, D is the face value of zero-coupon debt of maturity T , μ_V is the growth rate of the firm under the physical measure, and σ_V is the asset return volatility. Although not directly observable, the firm asset value and asset return volatility can be computed from the observed values of the stock price (S) and stock return volatility (σ_S), inverting via the Black-Scholes [1973] option pricing model.

In both academic (Vasallou & Xing [2004]) and industry applications (Moody's—KMV *EDF*, Moody's *RiskCalc*), the model has been used to impute one-year probability of defaults. As μ_V is difficult to estimate and sensitive to errors, the distance-to-default is often computed in practice as,⁶

$$DTD = \frac{1 - D/V}{\sigma_V\sqrt{T}}. \quad (1)$$


Thus, the Merton model identifies two of the primary determinants of the default probability for an individual firm as the fraction of debt in the capital structure and the firm volatility. The probability of default of a firm will increase with an increase in its debt ratio or an increase in its volatility.

We can understand how default probabilities vary across firm and over time by observing how debt levels and firm volatilities vary across firms and over time.⁷ Exhibit 1 plots the median debt (D/V) and annualized firm volatility (σ_V) across rated firms in the economy over the period 1987–2000. We break our sample into three quality categories, where firms of Moody's rating single-A or higher are classified as high-grade, Ba and Baa are classified as medium-grade, and single-B and C as low grade. The plot indicates that both debt and volatility have varied over this period. For high-grade (low-grade) firms, the debt of the median firm ranges from 10.2% to 16.8% (32.3% to 48.4%), and volatility from 18% to 36.5% (29% to 45%). Debt levels across firms in the economy appear to trend over time, while changes in volatility are more severe and tied to economic events. Debt levels first increase over 1987 to 1991, then trend downward in the aftermath of the 1991–92 recession, and again trend upwards in the latter half of the 1990s. Volatility shows large increases immediately after the 1987 crash, during

EXHIBIT 1

Debt and Volatility Over 1987–2000

The plot graphs the median debt fraction (D/V) and the firm return volatility (σ_V) of firms in each rating category for each of the 166 months over the period January 1987 to October 2000. The firm value and the firm return volatility are estimated on a monthly frequency using Merton [1974]. The debt is measured as the sum of current debt plus half the long-term debt. Firms with Moody's rating single-A or higher are classified as high-grade, Ba and Baa are classified as medium-grade, and single-B and C as low grade.



the 1991–92 recession, and the latter half of the 1990s stock market bubble that foreshadowed the bear market and subsequent recession of 2001–02. Excepting the unlikely event that an increase in firm volatility is offset by a specific decline in debt level, default probabilities of firms are not constant over time.

What is the relative impact of changes in debt and volatility on the probability of default of an individual

firm? Defining Δ as the difference operator, observe from equation (1) that,

$$\Delta(DTD) \approx \frac{-1}{\sigma_V \sqrt{T}} \Delta(D/V) - \frac{1-D/V}{\sigma_V^2 T} \Delta(\sigma_V \sqrt{T}),$$

indicating that, at low debt levels, the DTD is more sensitive to changes in volatility than to changes in debt levels.

Over our sample period, the median high-grade (low-grade) firm had an average volatility and debt of 24% (34%) and 11% (40%), respectively. At these volatility and debt levels, consider the impact of an absolute change of 1% in either debt or volatility on the *DTD*. For high grade firms, the impact of the change in volatility is 3.7 times the impact of the change in debt value, and for low grade firms, it is 1.8 times. Thus, for the average firm in the economy, the *DTD* is more sensitive to changes in volatility than changes in debt levels. In addition, over a short horizon, volatility is often much more variable than debt levels, especially for high and medium rated firms. As seen in Exhibit 1, volatility across the economy sharply increases in the aftermath of the 1987 crash as well as in the latter half of the last decade.

Determinants of Default Probability Correlations

The correlation between *PDs* of two firms will depend on the correlation between the underlying determinants of the default probabilities: correlation between individual firm debt levels, firm returns, and firm volatilities. For example, Zhou [2001] derives the joint probability of default for two firms when returns are correlated.

Exhibit 2 reports the median correlation between pairs of firms in our data. We classify our sample by rating as well as by sub-period. The rating groups are as previously defined, and the sub-periods are 1/87–6/90, 7/90–12/93, 1/94–6/97, and 7/97–10/00. We make the following observations. First, the median correlation between firms for each of these variables is positive. For high grade and medium grade firms, the highest correlation is between the volatilities of firms. For example, in the latest sub-period, the median correlation between volatilities for high-grade (medium-grade) firms is 0.71 (0.35), while the median correlation between debt levels is 0.25 (0.26). Second, differences between the rating classes appears to be driven mostly by differences in correlation between volatilities. Across all sub-periods, the relative difference in median correlations across all the three rating groups is the highest for the volatility correlations. Third, over time for each rating class, correlations between firm returns are the most stable, while the correlations between volatilities change the most. Interestingly, the correlations between volatilities are the highest in the Periods I and IV when the level of volatility itself is high across the economy (see, Exhibit 1).

Further Analysis

The results of the previous sections indicates that joint default risk will vary both over time and in the cross-section of firms. However, these results do not allow us to quantify joint default risk. To do so, we extend our dataset beyond the estimates from the Merton [1974] model. As Vassalou and Xing [2004] note, the strict application of the Merton [1974] model generates default probabilities that do not match actual aggregate default events in the economy. Instead, the standard approach in the industry (e.g., *EDF*, *RiskCalc*) is to estimate a model for default probabilities econometrically using a large dataset of default events and the *DTD* as an explanatory variable. The resultant estimates of default probabilities are thus calibrated to match the aggregate level of historical defaults. For this reason, Vassalou and Xing [2004] call the estimate from the Merton model a “default likelihood indicator,” *DLI*. We shall follow their terminology to define the *DLI* as,

$$DLI = 1 - \frac{1}{2}(DTD), \quad (2)$$

and reserve the term *PDs* to refer to probabilities of default that have been calibrated to match historical default levels.

CORRELATIONS OF DEFAULT PROBABILITIES

Default Probability Data

We obtain data on individual firms’ probabilities of default (“*PDs*”) from Moody’s Investors Service through its subsidiary Moody’s Risk Management Services (“*MRMS*”). With their extensive database of defaults, *MRMS* fits a model for short-term default risk using the distance to default from the Merton [1974] model and additional financial statement information. The output from their *RiskCalc* model is an estimate of a one-year default probability for an individual public firm at a monthly frequency. Details of the model and its econometric fit are in Sobehart and Stein [2000]. The model uses as input firm-specific information: a) company financial statement information (including leverage, profitability, and liquidity measures), b) the distance-to-default, and c) Moody’s ratings, if available. The use of information other than the distance-to-default is consistent with the Merton model when asset values are not perfectly observable (Duffie & Lando [2001]).

EXHIBIT 2

Correlations of Firms' Return, Debt and Volatility

The Exhibit reports correlations of (i) returns on firm value, (ii) debt fraction (D/V) and (iii) volatility (σ_V) across firms in each rating class. Firm value and return volatilities are computed using the Merton model with data from Compustat and CRSP. Correlations are computed for every pair of firms in the class. The median of all pair-wise correlations is reported in the Exhibit. Firms with Moody's rating single-A or higher are classified as high-grade, Ba and Baa are classified as medium-grade, and single-B and C as low-grade. The four sub-periods are: 1/87–6/90, 7/90–12/93, 1/94–6/97, and 7/97–10/00.

Group	Period I			Period II		
	Return	Debt	Volatility	Return	Debt	Volatility
High Grade	0.13	0.12	0.85	0.08	0.06	0.27
Medium Grade	0.10	0.10	0.59	0.06	0.12	0.15
Lower Grade	0.05	0.10	0.19	0.03	0.20	0.13
Group	Period III			Period IV		
	Return	Debt	Volatility	Return	Debt	Volatility
High Grade	0.05	0.07	0.10	0.07	0.25	0.71
Medium Grade	0.05	0.02	0.07	0.07	0.26	0.35
Lower Grade	0.04	0.05	0.06	0.08	0.27	0.06

The use of Moody's *PDs* has important advantages in analyzing joint default risk. First, Moodys calibrates their *PDs* to match the level of realized default; thus inferences about default correlations are directly related to actual economy-wide joint default. If defaults are independent conditional on *PDs*, then the correlations of *PDs* are also the correlations of default. This assumption of conditional independence (see, Jarrow, Lando & Yu [2005]) is a popular one in default models. Second, the Moody's *PD* uses the DTD as a co-variate for determining the probability of default. Duffie, Saita and Wang [2004],

for a subset of firms in the economy, demonstrate that changes in the DTD has the greatest impact on determining changes in the probability of default. Sobehart and Stein [2000] show that these *PDs* have predictive power, and perform as well, or better, than alternative measures. Finally, we are able to find a consistent pattern of correlations relative to the Moody's data by creating our own data set of default likelihood indicators.

The data consists of a panel of monthly probabilities of default for North American non-financial public firms over January 1987 to October 2000. The total number of

unique firms that enter our sample is 7,363. On any given date, the precise number of firms varies depending on new firm inceptions, mergers, or defaults. For some of our analysis, we sub-divide our sample by period and by rating classification. The four sub-periods include all firms that have continuous observations within any of these sub-periods. The number of sub-periods is motivated by the trade-off between having sufficient time series observations for estimation of correlation matrices, yet ensuring that the period is relatively homogeneous so that differences in correlation matrices across sub-periods may be observed. As noted earlier, the sub-periods are of almost equal length, and comprise the periods, 1/87–6/90, 7/90–12/93, 1/94–6/97, and 7/97–10/00. As before, we group firms by credit rating into three groups: high grade, medium grade and low grade. A firm is classified into a rating class according to its average rating within the sub-period.

Exhibit 3 describes the data. The total number of unique firms range from 3,202 to 5,170 over the sub-periods. The average number of firms in the high, medium and low grade groups are 215, 405 and 184, respectively. The vast majority of the firms in our sample are unrated, ranging from 2724 in Period I to 4142 in Period IV. The fraction of firms classified as high-grade declines from over 6% in Periods I and II to about 4% in Period IV. This decline is offset mainly by an increase in the fraction of firms classified as low grade. However, the fraction of firms in the two largest categories—medium-grade and non-rated categories—is stable over time. As these two categories together comprise about 90% of the total firms, the mix of firms across the economy is, for the most part, stable over time.

As might be expected, the mean default probability increases monotonically as the average rating declines, across each sub-period. For instance, the mean *PD* in Period IV is 0.23%, 1.17% and 5.65% for high, medium and low grade firms, respectively. The mean *PD* of unrated firms is in the range of 1.63% to 2.45%, suggesting that, if rated, these firms would fall in the lower rating classes.

For comparison, we also form an industry grouping by SIC broad industry code. More than half the firms in our sample belong to the manufacturing sector, Sector 4. The sector with the least number of firms is Sector 1 (agriculture, forestry and fishing) with as few as 14 firms in Period II. Although there is variation in *PDs* across sectors, this variation is small compared with that across

credit classes. Firms in sectors 1 and 5 have the lowest average *PD*, while firms in sector 10 have the highest. The high default risk of sector 10 arises from firms of SIC group 99 (firms of industries that cannot be classified into any of the other SIC groups).

Time Variation in Economy-Wide Default Probabilities

Two facts are immediately evident from Exhibit 3. First, there is considerable time-variation in *PDs* across sub-periods. Second, the pattern of time-variation is consistent across all firms. Except for firms in Sector 10, the mean default probability for groups of firms is the lowest in Period III and highest in Period IV. If individual firm default probabilities change over time and these changes are correlated, then this implies that the economy-wide default probability varies over time.

Exhibit 4 plots the monthly time series of the economy-wide probability of default, constructed by averaging the default probabilities of individual firms, $MEANPD(t) = \frac{1}{N(t)} \sum_{i=1}^{N(t)} PD_i(t)$, where $N(t)$ is the number of firms with a *PD* in month t . It is evident that the economy-wide probability of default varies substantially at times. From our previous analysis, we expect the economy-wide mean default probability to be related to economy-wide factors for volatility and debt. Comparing Exhibit 4 with Exhibit 1, we, in particular, observe that the time-variation in the mean default probability is closely related to the variation in firm volatilities. Each period of sharp increase in the level of the economy-wide default probability corresponds to a period of sharp increases in firm volatilities.

We verify the relative importance of debt and volatility at the aggregate level by a regression of the economy-wide mean default probability on economy-wide factors for volatility and debt. Define these to be the equal-weighted averages, $MKTVOL(t) = \frac{1}{\bar{N}(t)} \sum_{i=1}^{\bar{N}(t)} \sigma_{V_i}(t)$ and $MKTDEBT(t) = \frac{1}{\bar{N}(t)} \sum_{i=1}^{\bar{N}(t)} D(t)/V(t)$, where $\bar{N}(t)$ are the number of firms for which we have available data in month t . Using a first-order linear approximation, we estimate the following OLS regression,

$$\Delta MEANPD(t) = \alpha + \beta_1 \Delta MKTVOL + \beta_2 \Delta MKDEBT + \beta_3 (\Delta MKTVOL \times \Delta MKTDEBT) + \epsilon.$$

Exhibit 5 reports the results of the regression. We also report results for a subset of the variables to allow for comparison.

EXHIBIT 3

Descriptive Statistics of Default Probabilities


The Exhibit reports the mean default probabilities for U.S. non-financial public firms, sorted by rating and sector. The sample period of January 1987 to October 2000 is divided into 4 sub-periods: 1/87–6/90, 7/90–12/93, 1/94–6/97, and 7/97–10/00. The mean default probability for each group is then calculated by averaging monthly observations of Moody's default probabilities over each sub-period and over all firms within that group. Moody's rating single-A or higher are classified as high-grade, Ba and Baa are classified as medium-grade, and single-B and C as low grade. The sector groupings are by the broad industry classification according to the SIC code. Number of firms are shown in brackets.

Group	Period I (%)	Period II (%)	Period III (%)	Period IV (%)
High Grade	0.09 {213}	0.09 {223}	0.07 {210}	0.23 {215}
Medium Grade	0.75 {344}	0.77 {337}	0.49 {403}	1.17 {535}
Low Grade	4.56 {130}	5.01 {125}	3.29 {203}	5.65 {278}
Not Rated	1.66 {2724}	1.89 {2517}	1.63 {3183}	2.45 {4142}
Sector 1	1.59 {15}	1.99 {14}	1.21 {21}	2.03 {20}
Sector 2	2.04 {245}	1.80 {215}	1.36 {243}	2.78 {276}
Sector 3	2.43 {39}	2.39 {42}	1.92 {51}	2.47 {73}
Sector 4	1.50 {1788}	1.68 {1687}	1.38 {2083}	2.09 {2633}
Sector 5	1.17 {357}	1.48 {353}	1.35 {402}	2.40 {491}
Sector 6	1.75 {179}	2.13 {169}	1.87 {204}	2.87 {270}
Sector 7	1.74 {236}	1.98 {216}	1.79 {307}	2.92 {396}
Sector 9	1.80 {509}	2.14 {472}	1.99 {645}	2.76 {964}
Sector 10	1.99 {43}	3.53 {34}	3.56 {43}	3.83 {47}

EXHIBIT 4

Default Probabilities

Average default probability across all issuers in the sample on a monthly basis.



We can use the estimates of the coefficients to gauge the relative economic importance of volatility and debt. Over the 166 months of our period, the average $\Delta MEANPD$ is 1.05 basis points, while the average $\Delta MKTVOL$ and $\Delta MKTDEBT$ is 0.1828% and 0.0133%, respectively. Given the estimated coefficients, the fraction of the average $\Delta MEANPD$ over the 166 months of our period explained by $\Delta MKTVOL$ and $\Delta MKTDEBT$ is 74.9% and 9.36%, respectively. As expected from our previous analysis, the market-wide factor of volatility has a larger economic impact on the time-variation in the economy-wide default probability.⁸ As market-wide volatility is sensitive to economic conditions, the number of defaults in the economy is thus tightly linked to the state of the economy.

Empirical Evidence on Default Intensity Correlations

We next analyze how default correlations (defined below) vary over time and in the cross-section. To do so, we first convert the one-year default probability into a corresponding default intensity,

$$\lambda_i = -\ln(1 - PD_i), \quad (3)$$

where λ_i is the constant default intensity that corresponds to the 1-year default probability PD_i .⁹ This eliminates the maturity dependence from our data, and also allows our results to be interpreted within the framework of reduced form models (see, Jarrow & Turnbull [1995], Madan & Unal [1998], Duffie & Singleton [1999] and Duffie & Lando [2001]).

We examine the correlation between default intensities using two models. First, we compute the correlations between changes in default intensities. Let

$$\lambda_i(t) - \lambda_i(t-1) = \epsilon_i(t), \quad (4)$$

then the correlation between default intensities of firms i and j , ρ_{ij} , is the correlation between ϵ_i and ϵ_j over this period.

In our second model, we model the default intensity as a discrete time AR(1) process,

EXHIBIT 5

Default Probabilities, Market Volatility and Market Debt

The Exhibit reports the results of an OLS regression of changes in the mean *PD* on changes in market volatility and debt, using monthly observations over the period January 1987 to October 2000. We report the coefficients, the t-statistic (in parenthesis) and the adjusted R^2 . Standard errors are computed using the Newey-West procedure with 6 lags. The regression equation is,

$$\Delta \text{MEANPD} = \beta_0 + \beta_1 \Delta \text{MKTVOL} + \beta_2 \Delta \text{MKTDEBT} + \beta_3 (\Delta \text{MKTVOL} \times \Delta \text{MKTDEBT}) + \epsilon.$$

β_0 $\times 10^{-4}$	β_1	β_2	β_3	Adj. R^2 (%)
0.31 (0.36)	0.0420 (3.34)**			10.98
0.98 (1.04)		0.0723 (2.70)**		9.67
1.99 (0.27)	0.0431 (3.31)**	0.0740 (2.95)**	0.4738 (0.19)	21.21

$$\lambda_i(t) = \alpha_i + \beta_i \lambda_i(t-1) + \tilde{\epsilon}_i(t). \quad (5)$$

The correlation between firms i and j is the correlation between $\tilde{\epsilon}_i$ and $\tilde{\epsilon}_j$. The model of equation (5) is a mean-reverting intensity processes as used in the reduced form literature (Duffee [1999]). When we estimate the model across the firms in our database, the median β_i ranges from 0.90 to 0.94 across the rating classes and sub-periods.

We report both the usual Pearson correlation as well as the rank correlation. There are two reasons for reporting the rank correlation. First, the rank correlation imposes less stringent conditions on the data as the results are valid even if default intensities, like ratings, only order firms on the dimension of default risk. Second, the use of the rank correlation ensures that the results and conclusions of this section are applicable to other measures of default risk that are not necessarily calibrated to the actual level of defaults in the economy, but simply rank firms by their default probabilities. For example, these results also hold if the distance-to-default from the Merton model is used as an ordinal measure of default risk.

We compute the correlations for every pair of firms in our sample for sub-samples of our data, and the median of the pair-wise correlations is reported in Exhibit 6.

Panel A reports the statistics for the quality groups, and Panel B for the industry sectors. We report the following findings.

First, the median default correlation is non-negative across all credit classes and in every period. This result holds for both models of the default intensity process, equation (4) or equation (5), as well as whether the correlation coefficient is computed as the Pearson or as the rank correlation. For the model of equation (4), the magnitude of the Pearson coefficient ranges from 0.02 to 0.37. When computed from the residuals of equation (5), the correlation coefficients are of similar magnitude. That the median correlation is positive across all sub-samples is consistent with our earlier observation that default probabilities are affected by an economy-wide volatility factor.

Second, default intensities of high grade firms are more highly correlated than firms of lower rating classes. This conclusion is also robust across both models and measures of correlation coefficient. The difference between rating classes is particularly high in Periods I and IV. For example, in Period I, the median Pearson correlation computed from equation (5) is 0.37 for high-grade firms, compared with 0.23 and 0.16 for medium and low-grade firms. Non-rated firms have correlation coefficients

EXHIBIT 6

Correlations of Defaults Intensities

The Exhibit reports the median correlation ρ_{ij} computed using default intensities from Moody's PDs. The correlation is defined as the correlation between the shocks in each of the regressions,

$$\lambda_i(t) - \lambda_i(t-1) = \epsilon_i(t), \quad \text{MODEL 1}$$

$$\lambda_i(t) = \alpha_i + \beta_i \lambda_i(t-1) + \tilde{\epsilon}_i(t), \quad \text{MODEL 2}$$

where $\lambda_i(t)$ is the default intensity of firm i in month t . Correlations are estimated pairwise for each pair of firms in the sample, and the median is reported. Panel A reports the results by credit class. The first line reports the Pearson correlation coefficient. The second line reports the rank correlation coefficient. Panel B reports results by SIC code; to conserve space, only the Pearson correlation coefficient is reported.

Panel A

Group	Period I		Period II		Period III		Period IV	
	Model 1	Model 2	Model 1	Model 2	Model 1	Model 2	Model 1	Model 2
High Grade	0.36	0.37	0.10	0.10	0.02	0.01	0.37	0.38
	0.33	0.30	0.12	0.11	0.04	0.04	0.34	0.31
Medium Grade	0.22	0.23	0.10	0.10	0.03	0.02	0.24	0.25
	0.25	0.23	0.14	0.12	0.06	0.06	0.27	0.24
Low Grade	0.16	0.16	0.06	0.07	0.02	0.02	0.17	0.17
	0.19	0.17	0.13	0.13	0.05	0.05	0.20	0.18
Not Rated	0.16	0.16	0.05	0.06	0.02	0.02	0.16	0.17
	0.15	0.13	0.07	0.07	0.02	0.02	0.16	0.13

Panel B

Group	Period I		Period II		Period III		Period IV	
	Model 1	Model 2	Model 1	Model 2	Model 1	Model 2	Model 1	Model 2
Sector 1	0.54	0.19	0.14	0.06	-0.04	0.00	0.40	0.00
Sector 2	0.16	0.09	0.04	0.00	0.00	0.01	0.50	0.10
Sector 3	0.29	0.16	0.25	0.05	0.02	0.02	0.23	0.07
Sector 4	0.29	0.17	0.12	0.05	0.00	0.00	0.21	0.06
Sector 5	0.42	0.13	0.17	0.06	0.01	0.00	0.15	0.05
Sector 6	0.32	0.16	0.14	0.06	0.00	0.00	0.20	0.05
Sector 7	0.40	0.19	0.14	0.07	0.02	0.01	0.13	0.05
Sector 9	0.17	0.11	0.06	0.04	0.00	0.00	0.13	0.04
Sector 10	0.15	0.09	0.04	0.01	0.08	0.00	0.10	0.03

similar to those of medium and low-grade firms. The observation that default probabilities of high-grade firms are more highly correlated, and especially so in Periods I and IV, follows the pattern previously observed for asset return volatilities in Exhibit 2. Using data on credit spreads, Xiao [2003] observes that changes in credit spreads of high grade bonds are also more highly correlated than credit spreads of lower rated bonds.

Third, the median correlation within each group varies over the sub-periods. The striking pattern that can be observed in Exhibit 6 is that the time-variation in correlations is consistent across sub-groups. Across all rating classes, the correlations are higher in Periods I and IV, compared to Periods II and III. For each of the rating classes, the lowest Pearson correlation for Model 1 is close to zero in Period III, but increases to a range of 16–36%

EXHIBIT 7

Principal Components

For each credit class, the Exhibit reports the fraction of the variance of the residual (in %) from Equation 4 (Model 1) and Equation 5 (Model 2) that is explained by the first, and first two principal components, respectively.

Group	Period I		Period II		Period III		Period IV	
	Model 1	Model 2	Model 1	Model 2	Model 1	Model 2	Model 1	Model 2
High Grade	59.84	55.01	24.91	24.70	15.21	15.28	60.18	55.88
	95.08	95.20	36.40	36.44	26.83	27.08	95.39	95.41
Medium Grade	16.38	16.34	25.46	27.27	16.69	16.70	16.51	16.45
	28.49	29.45	40.36	41.83	29.56	29.92	29.37	29.97
Low Grade	22.10	22.64	18.44	19.67	10.88	11.10	22.75	23.43
	34.21	34.98	31.80	32.80	18.74	19.53	35.08	35.98

in Period IV. Once again, the pattern of time-variation in correlations of default intensities is closest matched in Exhibit 2 by the pattern of variation in the correlations between firm volatilities. This similarity also extends to the cross-sectional differences in the magnitude of time-variation of correlations. High-grade firms have the greatest variability in correlations of both asset return volatilities and default intensities, and low-grade firms show the least.

Panel B of Exhibit 6 reports the results when firms are classified into industry sectors. Median correlations are, once again, predominantly positive. As with rating classes, correlations are time-varying, with correlations in Periods I and IV higher than in Periods II and III across almost all sectors. With the sole exception of sector 10 (Public Administration), the lowest median correlation is in Period III. Thus, the time variation of correlations appears to be systematic across the firms, irrespective of how firms are classified.

We provide two alternative sets of results that, although of independent interest, also serve as robustness tests. First, as an alternative to considering pair-wise correlations, we analyze the principal components of the residuals from Equation (4) and Equation (5). Exhibit 7 reports the fraction of the variance of the residuals explained by the first, and the top two principal components, respectively. The results indicate that the primary cross-sectional and time-series conclusions of Exhibit 6 are robust. In Periods I and IV, the first two components explain more than twice the variance for high grade firms

than for lower grades. Overall, the conclusions from Exhibit 7 are consistent with those of Exhibit 6.

We also report correlations across rating classes for the Merton *DLI* of equation (2). The use of the *DLI* helps us to verify that the results are not dependent on the specific econometric model used to calibrate the default probabilities. Exhibit 8 reports the results. Although the magnitude of the median correlation is lower as compared with Model 1 in Exhibit 6 (as estimates of default probabilities include additional factors besides the *DLI*), the correlations nevertheless exhibit the same patterns in the cross-section and time-series.

An important conclusion that emerges from the results of this section is that *both* the time-variation in default intensities as well as time-variation in the correlation between default intensities are of substantial magnitude. More importantly, as increases in default probabilities appear to coincide with increases in default correlations, both these effects will result in substantial variation in the economy-wide default risk.

MODELING TIME-VARIATION IN DEFAULT PROBABILITIES AND DEFAULT CORRELATIONS

In this section, we propose a statistical model that explicitly accounts for systematic time-variation in both default probabilities and their correlations. Empirical evidence of the previous sections indicates that economy-

EXHIBIT 8

Default Likelihood Indicator

For each credit class, the Exhibit reports rank correlations of changes in the 'default likelihood indicator', $DLI = 1 - \Phi(DTD)$. $\Phi(\cdot)$ is the normal cumulative distribution function, and DTD is the estimate of the "distance-to-default" from equation 1. Rank correlations are computed for every pair of firms, and the median is reported. The number of firms are reported in brackets.

Group	Period I	Period II	Period III	Period IV
High Grade	0.31 {197}	0.11 {180}	0.05 {152}	0.18 {203}
Medium Grade	0.17 {294}	0.09 {314}	0.02 {339}	0.10 {474}
Low Grade	0.09 {96}	0.07 {98}	0.02 {161}	0.08 {172}
Non Rated	0.08 {2552}	0.04 {3063}	0.00 {3207}	0.05 {3322}

wide levels of default risk, default correlations and volatility are related to each other. Increases in volatility lead to increases in the economy-wide default level as well as increases in the correlations of default intensities. We directly model this link between economy-wide default risk and the magnitude of correlations through a Hamilton [1989] type regime-shifting model with different correlations across regimes. A comparison of sub-period IV (high default risk, high correlations) versus sub-period III (low default risk, low correlations) suggests that a two-period regime-switching model is a likely candidate for good fit. Similar time-varying correlation models have been proposed in Erb, Harvey and Viskanta [1994], Longin and Solnik [1995] and Ang and Chen [2002] for modeling time-varying and asymmetric equity correlations.

Determining Regimes for Default Intensities

We estimate a two-regime model for default risk in the economy. Let the average (across all issuers) default intensity, denoted by $\lambda(t) = [\frac{1}{N(t)} \sum_{i=1}^{N(t)} \lambda_i(t)]$, follow an AR(1) model that depends on the prevailing regime, $k_t \in \{1, 2\}$, at time t (we will suppress the subscript on k_t below):

$$\lambda(t) - \theta^k = \beta^k [\lambda(t-1) - \theta^k] + v^k \tilde{z}_t$$

$$\text{Regimes } k_t = \begin{cases} 1 & \text{high default regime} \\ 2 & \text{low default regime} \end{cases}$$

where $\tilde{z}_t$ has a standard normal distribution. k_t follows a Markov chain with a transition matrix, $\begin{pmatrix} q_1 & 1-q_1 \\ 1-q_2 & q_2 \end{pmatrix}$, where q_1 and q_2 are the transition probabilities, for $s > t$, $q_1(s, t) = Pr(k_s = 1 | k_t = 1)$ and $q_2(s, t) = Pr(k_s = 2 | k_t = 2)$. The mean reversion rate, $1 - \beta^k$, the mean default rate, θ^k , and the volatility v^k are regime dependent. The discrete transition density is

$$f[\lambda(s) | \lambda(t)] = \exp \left(\frac{-\left(\lambda(s) - \lambda(t) - (1 - \beta^k)(\theta^k - \lambda(t))\right)^2}{2(v^k)^2} \right) \times \frac{1}{\sqrt{2\pi(v^k)^2}} \quad (6)$$

We write the Markov probabilities at time t in logit form as,

$$q_{kt}(s, t) = \frac{\exp(a^k + b^k x_t)}{1 + \exp(a^k + b^k x_t)}, \quad k_t = 1, 2 \quad (7)$$

EXHIBIT 9

Estimation of the Regime Switching Model

The Exhibit reports estimates from a two-regime model for default probabilities. Panel A is for the case of a constant transition matrix, and Panel B for the case where transition probabilities vary with the short rates. Panel C presents the results when the state variable is set to average firm return volatility.

Parameters	Panel A Constant Transition Probabilities		Panel B <i>r</i> Dependent Probabilities		Panel C <i>MKTVOL</i> Dependent Probabilities	
	Estimates	t-stats	Estimates	t-stats	Estimates	t-stats
β^1	0.78	17.37	0.78	17.21	0.78	18.12
β^2	0.92	57.48	0.92	51.94	0.92	58.10
θ^1	1.77%	18.43	1.76%	17.03	1.77%	19.12
θ^2	0.74%	15.61	0.74%	16.33	0.74%	15.77
v^1	0.07%	6.57	0.07%	7.61	0.07%	7.56
v^2	0.04%	15.28	0.04%	16.47	0.04%	15.95
a^1	1.96	3.39	3.46	0.47	1.41	0.30
a^2	3.45	6.59	5.14	2.56	7.65	2.28
b^1			-0.26	-0.39	2.32	0.17
b^2			-0.29	-0.91	-14.21	-1.31

where x_t is a state variable that impacts the transition probability, and a^k and b^k are transition function parameters to be estimated. Here, we allow for the possibility that the transition probabilities may vary over time with the short rate ($x_t = r_t$), or with average firm volatility ($x_t = MKTVOL(t)$), or be constant, i.e., $b^k = 0$. As we expect regimes, if any, to be correlated with economic events, the short-rate and volatility are natural candidates for determining transition probabilities. The standard Hamilton approach is used to fit the following objective function:

$$\begin{aligned} & \{\beta^1, \beta^2, \theta^1, \theta^2, v^1, v^2, a^1, a^2, b^1, b^2\}^* \\ & = \arg \max \sum_{t=1}^T \log[f(\lambda_t)] \end{aligned} \quad (8)$$

Maximum likelihood estimates for the model are presented in Exhibit 9.¹⁰ Panel A provides the estimates


under the assumption of constant transition probabilities, $b^k = 0$. The mean level of default in regime 1 ($\theta^1 = 1.77\%$) is more than twice that of regime 2 ($\theta^2 = 0.74\%$). The volatility, v^k , is also higher in regime 1 versus regime 2 (0.07% vs. 0.04%). The probability parameter (a^k) results in high values of q_k , suggesting strong persistence in each regime. The probability of remaining in regime 1 is lower (0.8765) than the probability of remaining in regime 2 (0.9692). Raising the transition matrix to the power of infinity provides the long-run stable probabilities of each regime, which are roughly 20% in the regime 1 and 80% in regime 2. All estimates are statistically significant in Panel A. Panel B and C reports the estimates where the transition probabilities are allowed to be state-dependent. The coefficient b^k is insignificant in both panels, and therefore, the results are economically similar to those in Panel A.

The parameter estimates across the two regimes identify regime 1 (2) as a regime with high (low) mean

EXHIBIT 10

Time Series of Regime Probabilities

This Exhibit presents the time series plot of probability of being in the high default intensity regime. The probability of being in the low intensity regime is one minus the high intensity probabilities, since there are only two regimes in the estimated model. The transition probabilities are generated from the model where they are a function of the short rate. The results with constant transition probabilities or when transition probabilities are a function of volatility are very similar.



default rates and volatility. In Exhibit 10, the probability of being in regime 1 is presented for the time period spanned by the data set. The probabilities track the patterns originally observed in Exhibit 4, indicating, in particular, the existence of high/low default rate regimes.

Regime-Dependent Parameters for Rating Classes

In the previous sub-section, we determined the two regimes within our sample period in terms of the process for the average level of default in the economy. Next, we look at how the parameters of the intensity process vary across these two regimes for the average firm in each rating class. We identify a specific time period as being of regime 1 (High) or 2 (Low), if the probability of being in that regime is more than 0.5 (see, Exhibit 10). Of the 166 month period from January 1987 to October 2000, 30 months are identified as being in regime 1.

Using the defined regimes, we estimate the parameters for each rating class in each regime. We model the rating level default intensity

$$\lambda_i^k(t) = \frac{1}{N_i(t)} \sum_{j=1}^{N_i(t)} \lambda_{ij}^k(t)$$

as follows, where we use i to index the rating class and j to index a firm within a rating class.

$$\begin{aligned} \lambda_i^k(t) = & \lambda_i^k(t-1) + (1 - \beta_i^k) [\theta_i^k - \lambda_i^k(t-1)] \\ & + v_i^k \epsilon_i(t), \quad k \in \{1, 2\}. \end{aligned} \quad (9)$$

Assuming a standard normal distribution for $\epsilon_i(t)$, the parameters are estimated by maximum likelihood,

$$\max_{\Omega} \sum_{t=1}^T \ln F[\lambda_i^k(t)], \quad (10)$$

where $\Omega = \{\beta_i^k, \theta_i^k, v_i^k\}, k \in \{1, 2\}$.

$$F[\lambda_i^k(t)] = \mathbf{1}_1 \times f[\lambda_i^1(t)] + \mathbf{1}_2 \times f[\lambda_i^2(t)], \quad (11)$$

EXHIBIT 11

Regime Dependent Parameters within each Rating Category

For each rating category, i , the Exhibit reports the maximum likelihood estimates of the model,

$$\lambda_i^k(t) = \lambda_i^k(t-1) + (1 - \beta_i^k)[\theta_i^k - \lambda_i^k(t-1)] + v_i^k \epsilon_i(t), k \in \{1, 2\}.$$

across two regimes. The regimes are based on the estimation in Panel A of Exhibit 9. The t-statistics are reported in parenthesis.

Rating Class			Parameters			
	β_i^1	β_i^2	θ_i^1 (%)	θ_i^2 (%)	v_i^1 (%)	v_i^2 (%)
High	0.97 (82.00)	0.99 (49.32)	0.25 (5.67)	0.11 (3.77)	0.03 (7.18)	0.01 (14.80)
Medium	0.86 (18.64)	0.98 (49.37)	1.48 (9.00)	0.41 (2.54)	0.09 (9.37)	0.05 (17.34)
Low	0.85 (8.92)	0.96 (7.77)	4.81 (8.93)	2.38 (7.77)	0.12 (8.21)	0.12 (16.07)

and $f[\lambda_i^k]$ is the conditional density for the stochastic process in equation 9 and $\mathbf{1}_k$ are indicator functions based on the regimes identified in the previous subsection.

The results in Exhibit 11 show clearly that the mean default rate is different in the two regimes for every class. Mean default rates (θ_i^k) increase by 127%, 261% and 102%, for the high, medium and low grade firms, respectively. Volatilities (v_i^k) also increase in the high regime, and as credit quality declines. The difference in the volatility between the two regimes differs by credit class, with the lowest credit class showing little difference across regimes, in contrast to a three-fold increase for the highest rating class. For each rating class, the rate of mean reversion, $(1 - \beta_i^k)$, is also different across the two regimes, and is higher in regime 1. Mean reversion increases as credit quality decreases.

Overall, the regime-dependent parameters for the average firm in each class show similar patterns as those identified earlier for the overall default level of the economy.

Issuer Correlations Across Regimes

For firms in each rating class, we compute the covariance matrix of Model 2 (equation 5) residuals in each of

the two regimes, and test for differences in covariances across regimes using the following test.¹¹ Let S^k , $k = \{1, 2\}$ be the covariance matrix for each regime k . The null hypothesis is that the covariance matrices are equal across the two regimes, i.e., $S^1 = S^2$. To construct the test, define

$$C^k = \frac{S^k}{n^k - 1},$$

where n^k , $k = \{1, 2\}$, are the number of observations used to compute each covariance matrix. Compute the pooled covariance matrix as follows:

$$C = \frac{\sum_k (n^k - 1) C^k}{n - m}, k = \{1, 2\}$$

where $n = \sum_k n^k$, and $m = 2$ (testing for the equality of only 2 covariance matrices). The null hypothesis of equality of the covariance matrices is tested using a modified log-likelihood ratio statistic L ,

EXHIBIT 12

Statistical Differences in Covariance Matrices

This Exhibit presents the chi-square statistics for the differences in covariance matrices in the high and low default regimes. The test was performed on a randomly chosen group of 30 firms from each rating class. The sample is limited to 30 firms as the number of months in the shorter regime—regime 1—are thirty. Similar results are obtained for alternative samples.

Rating Class	Test-Statistics	Degrees of Freedom	p-value
High Grade	985.69	465	0.00
Medium Grade	657.65	465	0.00
Low Grade	904.75	465	0.00

$$L = M \times h,$$

$$M = (n - m) \ln |C| - \sum_k (n^k - 1) \ln |C^k|,$$

$$h = 1 - \frac{2p^2 + 3p - 1}{6(p + 1)(m - 1)} \left(\sum_k \frac{1}{n^k - 1} - \frac{1}{n - m} \right),$$

where p is the dimension of the covariance matrix, and L is distributed $\chi^2[p(p + 1)(m - 1)/2]$. If L exceeds the critical value, then we reject the hypothesis that the covariance matrices are equal.

Results are presented in Exhibit 12 on a random sample of 30 firms in each rating category. For each of the three rating categories, we can easily reject the hypothesis that covariance matrices are equal across the two regimes. Thus, the results suggest that correlations are, indeed statistically different between regimes.

In addition, we also considered the fraction of the variance explained by the first principal component in each regime across each rating class. In the high default regime, the fraction explained by the first principal component is 61%, 55% and 69% for high, medium and low grade firms, respectively. In contrast, the fraction explained by the first principal component in the low default regime is 32%, 34% and 35% for each rating class, respectively. Overall, the evidence indicates the correlations across the two regimes are different, with the correlation, on average, being higher in the high-default regime as compared with the low-default regime.

The evidence in this section indicates that a statistical regime-shifting model is able to capture both the time-variation in default probabilities and correlations that were observed in the previous sections. The regimes are based on the economy-wide level of default probability, and thus accounts for the two observations that the changes in joint default risk are systematic, and that changes in both the level of the default probability and default correlations are related to common factor(s). Because the implementation of the model does not require any additional information beyond the time-series of default probabilities, it can be easily implemented.

In the next section, we provide applications that use the analysis of this and the previous sections to quantify the differences in regimes in terms of the likelihood of (joint) default across a portfolio of firms.

Analysis

In this section, we provide two numerical examples to assess the economic importance of our findings. We do so using standard factor copula techniques that are widely used by practitioners. First, we compute n 'th default probabilities to illustrate the differences across regimes. As the probability of the n 'th default from a basket of issuers is sensitive to both the levels of the default intensity as well as the correlations between the intensities, these probabilities are a good way to illustrate the economic importance of the regimes detected earlier. Second,

we examine the economic impact of changing economic conditions. We quantify the tails of the distribution of defaults over a large portfolio of firms, and ask how unusual it is to observe the economy-wide level of defaults in 2001.

N'th to Default Probability

Consider the computation of the probability of exactly n defaults from a credit basket of $m \geq n$ bonds.¹² These probabilities are sensitive to both the level of default and the degree of correlation between the issuers, and allow us to quantify the effects of both in a simple one-period model.

Let each bond have a distinct issuer i , whose default risk is characterized by a default intensity, $\lambda_i = \theta_i + \nu_i^* x_i$, where θ_i is the mean intensity, and ν_i^* is the unconditional standard deviation. The variable x_i consists of a random shock which is a function of a latent factor deviation Y , common across all issuers, with constant cross-sectional correlation ρ , and an idiosyncratic component ϵ_i , i.e.,

$$\begin{aligned} x_i &= \rho Y + \sqrt{1-\rho^2} \epsilon_i, \\ \epsilon_i &\sim N(0,1), \\ Y &\sim (\mu_Y, \sigma_Y^2). \end{aligned}$$

The variables Y and ϵ_i are uncorrelated, as are the ϵ_i s. This specification injects the required correlation amongst the intensities for each issuer through the common component Y .

Assuming unit time, the survival probability is equal to $s_i = e^{-\lambda_i}$, and therefore, the expected survival probability, conditional on Y , is

$$\begin{aligned} \bar{s}_i(Y) &= E[\exp(-\lambda_i) | Y] \\ &= E\{\exp[-(\theta_i + \nu_i^* x_i)]\} \\ &= E\left\{\exp\left[-\left(\theta_i + \nu_i^* [\rho Y + \sqrt{1-\rho^2} \epsilon_i]\right)\right]\right\} \\ &= \exp\left\{-\theta_i - \nu_i^* \rho Y + \frac{(\nu_i^*)^2}{2}(1-\rho^2)\right\}. \end{aligned} \quad (12)$$

Likewise, the expected conditional default probability is $\bar{p}_i(Y) = 1 - \bar{s}_i(Y)$.¹³

In order to examine the probability distribution of the number of defaults, we use probability generating functions, implementing the ideas previously developed

in articles by Finger [1999], Mina and Stern [2003], and Gregory and Laurent [2004]. We denote the probability of n defaults from m bonds in unit time as $p(n)$. From m issuers, n can take values in the set $\{0, 1, 2, \dots, m\}$. Conditional on Y , the probability of n defaults is denoted $p(n | Y)$. We define the probability generating function of $p(n)$ as $g(t | Y)$ conditional on Y :

$$\begin{aligned} g(t | Y) &= \sum_{n=0}^m p(n | Y) t^n, \quad 0 \leq t \leq 1 \\ &= p(0 | Y) + p(1 | Y)t + p(2 | Y)t^2 + \dots \\ &\quad + p(m | Y)t^m, \end{aligned} \quad (13)$$

where t is the parameter of the probability generating function (pgf). For each individual issuer i , the pgf is as follows:

$$\begin{aligned} g_i(t | Y) &= p_i(0 | Y) + p_i(1 | Y)t \\ &= 1 - \bar{p}_i(Y) + \bar{p}_i(Y)t \\ &= \bar{s}_i(Y) + (1 - \bar{s}_i(Y))t. \end{aligned} \quad (14)$$

Conditional on Y , the pgf of each issuer is independent of the other issuers, and we may write the joint pgf as:

$$\begin{aligned} g(t | Y) &= \prod_{i=1}^m g_i(t | Y), \\ &= \prod_{i=1}^m [\bar{s}_i(Y) + (1 - \bar{s}_i(Y))t], \end{aligned} \quad (15)$$

which is calculated using equation (12). We let $t \in \{0, 1/m, 2/m, \dots, 1\}$, this represents $(m+1)$ equations, one for each t . This calculation provides the LHS value for the system of equations in (13), which then contains $(m+1)$ probabilities $p(n | Y)$, which may be solved for using a simple matrix inversion. By integrating out Y , we obtain

$$p(n) = \int_Y p(n | Y) dY$$

These are precisely the values we are interested in, i.e., the probability density function for the number of defaults n .

For illustration, we set the number of issuers $m = 10$. We assume that the factor deviation Y is standard normal. We use parameters values β_p , θ_p , and ν_i from Exhibit 11. In

this one-period model, we obtain the unconditional standard deviation of intensities from the Exhibit, and denote them v_i^* , where v_i^* is set to $\sqrt{v_i^2 / 2(1 - \beta_i)}$, approximating unconditioning. The correlations corresponding to the high and low default regimes are the median correlation from Period III and IV, respectively, in Exhibit 6. Using the procedure outlined above, we compute the values of $p(n)$, $n = 0 \dots m$, across regimes and rating classes. Our goal is to examine the difference in these distributions across regimes for each rating class.

The results are presented in Exhibit 13. The difference in the distributions across the two regimes is apparent. For example, the probability of observing a default in a portfolio of 10 high-grade firms is 2.47% in the high default regime, more than twice the probability in the low-default regime. Even more striking is the difference in the tails of the distributions for larger number of defaults—the probabilities of default in the high-default regime are a large multiple of the values in the low regime (although limitations on the number of decimals that we can report in Exhibit 13 does not always make it apparent).

The procedure outlined above for the implementation of the n -th to default model is extendable to a much larger basket of bonds. As an illustration, we extended the approach to $m = 100$ bonds. The integration of the probabilities in this approach is undertaken using a Fourier transform approach instead, which speeds up the computations. Using the same parameters as in Exhibit 13, we generated the distribution of probabilities for n out of m defaults, i.e., $p(n)$. For the high default risk regime, the probability density functions for all three rating classes are presented in Exhibit 14. The plots for the low default risk regime are also plotted in the same Exhibit. The tails of the distribution are larger in the high risk regime. In the low rating class, there is a 5% probability of 11 or more defaults in the high default regime, whereas in the low default regime, there is a 5% probability of only 7 or more defaults.

Impact of Changing Economic Conditions

In the next illustration, rather than using average parameter estimates, we use default intensity data across

EXHIBIT 13

Probabilities of Exactly n Defaults from $m = 10$ Bonds

The two exhibits below show the parameters used and the probabilities of n defaults.

Panel A: Input parameters

Rating class	High default regime			Low default regime		
	$(\theta_i\%)$	(v_i^*)	Corr	$(\theta_i\%)$	(v_i^*)	Corr
High	0.25	0.1225	0.38	0.11	0.0707	0.01
Medium	1.48	0.1701	0.25	0.41	0.2500	0.02
Low	4.81	0.2191	0.17	2.38	0.4243	0.02

Panel B: Probabilities of n defaults

No. Defaults n	High Default Regime			Low Default Regime		
	Rating High	Rating Medium	Rating Low	Rating High	Rating Medium	Rating Low
	$p(n)$	$p(n)$	$p(n)$	$p(n)$	$p(n)$	$p(n)$
0	0.975327	0.862451	0.618184	0.989063	0.959859	0.788274
1	0.024387	0.128565	0.304590	0.010883	0.039405	0.189786
2	0.000284	0.008632	0.067539	0.000054	0.000728	0.020562
3	0.000002	0.000344	0.008875	0.000000	0.000008	0.001320
4	0.000000	0.000008	0.000765	0.000000	0.000000	0.000056
≥ 5	0.000000	0.000001	0.000047	0.000000	0.000000	0.000002

EXHIBIT 14

Plots of the Probability Density Functions of the Number of Defaults n out of m Bonds for the High and Low Default Risk Regimes

In this illustration $m = 100$.






EXHIBIT 15

Difference in the Probability $p(n)$ of the Number of Defaults n , across High and Low Default Risk Regimes

The plots below are for each rating class (High, Medium and Low). They depict the difference in the PDF over the number of defaults, obtained by subtracting the low regime PDF from the high regime one. The Exhibit shows that the tails of the distribution are fatter for the high risk regime, as would be expected.



the firms in our dataset to examine the economic impact of changing economic conditions across a large portfolio of firms. To do so, we first combine data from sub-periods I and IV to comprise a “high” default risk regime, and data from sub-periods II and III to form a “low” default risk regime. We then form portfolios of firms that have data across all months in our sample. The number of firms in each rating category in our analysis are 114, 238, and 106 for the high, medium and low rating classes respectively.¹⁴

Within each regime and rating class, using the raw intensity data, we compute the vector of mean intensities for all firms and the covariance matrix. Assuming that intensities are lognormal,¹⁵ and a horizon of one year, we can compute the distribution of aggregate intensity (denoted A) for each rating class. The probability of n defaults (assuming independence across defaults conditioned on A) is denoted $p(n|A)$. To obtain the probability of n defaults, we then integrate, i.e.,

$$p(n) = \int_A p(n|A) f(A) dA, \forall n,$$

where $f(A)$ is the probability density function for A , and $p(n|A)$ is Poisson with parameter A . Computing $p(n)$ for all n results in the probability function for the number of defaults.


In Exhibit 15, we present the difference in probability functions between the high and low regimes. The PDF of the low regime is subtracted from the PDF for the high regime. This is done for each of the three rating classes. We observe that the tail of the distributions is much fatter in the high default regime.

To gain a sense of the economic impact of the differences across these two regimes, consider the low rating class in both regimes, as a proxy for the High-Yield debt market. We plot in Exhibit 16 the cumulative default probability of the number of defaults of the 106 firms in that category in each regime. To fix ideas, consider the fact that the default rate in the High-Yield Sector in 2001 was about 10%—at a historical peak. This would be about 11 firms or more in our sample defaulting in about one year (recall that our sample period ends in October 2000). In Exhibit 16 we have plotted a vertical line at this number of firms. We can see that the probability of reaching or surpassing

EXHIBIT 16

Difference in the Cumulative Probability of the Number of Defaults, across High and Low Correlation Regimes for the Low Rating Class

The total number of firms in this class was 106. The green vertical line represents the cut off for 10% of the total number of firms.



this level of default is less than 10% in the low regime, but is more than twice in the high regime—there is a 1-in-5 chance of seeing the default levels of 2001. In summary, when one models joint changes in default probabilities and default correlations across regimes, the high number of defaults observed in 2001 is not a low probability event.

CONCLUSION

In this article, we document how the correlation of default probabilities varies both in the cross-section of firms and over time. Our primary conclusion is that joint credit risk varies because both default probabilities and default correlations vary with economic conditions. Moreover, market-wide volatility plays an important role in determining this time-variation in joint default risk. Clustering of defaults occurs during times of high volatility, and does so because both default probabilities and correlation between defaults increase.

We provide below a summary of our findings and their implications:

1. Although the fact that there is time-series variation in default probabilities is perhaps not surprising, the magnitude of this variation certainly is. Across each broad rating category, the default probability increases by more than 100% between times of low default risk and times of high default risk.
2. When default probabilities rise, so do their correlations. Correlations rise from close to zero to levels of 17–38% that are much higher than the corresponding correlations between asset returns. The concurrent increase in both default probabilities and their correlations results in variation in joint default risk over time and the clustering of defaults (fat-tailed behavior of the distribution of defaults) with important implications for the management of portfolio credit risk.
3. Both default probabilities and default correlations are related to firm asset volatility. This provides an economic explanation for the relation between credit spreads and economy-wide volatility (Campbell and Taksler [2003]). Firm volatility is also important for

determining differences in joint default risk across the cross-section of firms. Structural models of joint default risk should explicitly model correlations of volatilities, along with the correlations between firm returns.

4. Both the default intensity process of individual firms and the correlation between the intensity processes between pairs of firms may be modeled within a statistical framework, using regimes based on an aggregate economy-wide default level. Such a reduced form approach is relatively simple to implement given estimates of default probabilities. We find considerable differences across regimes for the values of products that depend on the level and correlation of default, such as n -th to default basket contracts.

More remains to be learned about why and how defaults cluster. In fact, joint default risk may be even higher than what we observe in our dataset. First, given the industry practice of estimating default probabilities using the Merton [1974] distance to default leads us to focus primarily on debt levels and firm volatility as two main drivers of default risk. Yet, recent research (Duffie, Saita & Wang [2004]) indicates that macroeconomic variables such as personal income growth and term structure level and slope are also significant in explaining the dynamic movement of default probabilities beyond just the distance to default. Second, our focus has been on the common factors that jointly affect individual firm default probabilities. Even after conditioning on these probabilities, defaults may further cluster because of contagion-like effects, when the default of one firm affects the probability of default of another firm. Recent evidence in Das, Duffie, Kapadia and Saita [2005] indicates that the latter is also significant, but contributes much less to default clustering than the co-movement of default probabilities. This article has developed a basic understanding of joint default risk, and motivates a further exploration of the clustering of defaults.

ENDNOTES

We are extremely thankful for many constructive suggestions from Gurdip Bakshi, N. Chidambaran, Darrell Duffie, Rong Fan, Gifford Fong, John Knight, N. R. Prabhala, Jun Pan, Dmitry Pugachevsky, Shuyan Qi, Ken Singleton, Rangarajan Sundaram, Suresh Sundaresan and Haluk Unal. We received useful feedback from participants at various seminars:

at the AIMR talks in Tokyo, Singapore and Sydney, the AIMR Research Foundation Workshop in Toronto, Risk conferences in Boston and New York, BGI in San Francisco and Citicorp in New York, the Credit Conference at Carnegie-Mellon University, Institutional Investors' Fixed Income Forum, the Mathematical Sciences Research Institute workshop on Event Risk, the Federal Deposit Insurance Corporation and the Office of the Comptroller of the Currency, and discussants and participants at the 2003 meetings of American Finance Association and European Finance Association. The first author gratefully acknowledges support from the Price Waterhouse Cooper's Risk Institute, the Dean Witter Foundation, and a Research Grant from Santa Clara University. We are also grateful to Gifford Fong Associates, and Moody's Investors Services for data and research support for this article.

¹Default & Recovery Rates of Corporate Bond Issuers, Moody's Special Comment, February 2002.

²Precisely, we examine default probability correlations or the correlation between default intensities that are equivalent to our data on one-year default probabilities. We defer more detailed discussion to Section 3 below.

³For theoretical modeling of credit risk, see the structural models of Merton [1974], Longstaff and Schwartz [1995], Leland and Toft [1996], [1995], Collin-Dufresne and Goldstein [2001], and reduced form models of Duffie and Singleton [1999], Jarrow and Turnbull [1995], Das and Tufano [1996], Jarrow, Lando and Turnbull [1997], Madan and Unal [1998], Das and Sundaram [2000], Duffie and Lando [2001], among others.

⁴Allen and Saunders [2002]. Altman, Brady, Resti and Sironi [2002] find that the state of the economy also drives the level of recovery on default.

⁵Other structural models include Longstaff and Schwartz [1995], Leland and Toft [1996], and Collin-Dufresne and Goldstein [2001]. Duffie, Saita and Wang [2004] show empirically the relevance of modeling distance to default.

⁶For example, see footnote 8 in Sobehart and Stein's [2000] description of the implementation of Moody's *RiskCalc* model. In the Merton [1974] model, as the stock is priced using the risk-neutral distribution, the expected return of the firm is not required.

⁷To do so, we estimate debt ratio and firm volatilities across all firms in the Compustat-CRSP combined database. We first extract data for the period 1987–2000 on stock prices and debt per share. Using historical equity return volatilities for one year, we then invert the Merton [1974] model to obtain firm value V and firm volatility σ_V . We solve a system of two equations: one, which represents equity as a call option, and the other, the relationship of equity volatility to firm volatility. This pair of equations contain two unknowns, V and σ_V which are easily solved for using standard numerical procedures. In our computations to estimate a one-year probability of default,

we set $T = 1$, and D equal to the total of the face values of short term debt, and one-half of long term debt. We filter out firms whose volatilities are less than 0.1% or greater than 500%. Firms with debt ratio greater than 0.95 are considered to be in default, and also eliminated from the sample. Data on the one-year risk free interest rate is obtained from the Federal Reserve.

⁸This conclusion is consistent with Campbell and Taksler's [2003] recent finding that corporate bond spreads are correlated with volatility.

⁹The transformation of PDs to default intensities is based on the assumption that intensity process is constant. As a check, for a number of firms, we also derive the instantaneous intensity process by modeling it as an affine model of the CIR form. The correlation of the resultant intensity process under the two separate transformations is close to 1. See also the discussion in Das, Duffie, Kapadia and Saita [2005]. Thus, given the objectives of this article, we choose to work with the constant intensity process.

¹⁰The specifications for intensities in this and the next subsection do not preclude negative intensities. A logarithmic specification (that ensures positive intensities) did not converge for the regime-switching estimation model. However, the parameter estimates for the volatility v^k in both regimes in comparison to the mean values of intensities θ^k shows that the probability of negative intensities is extremely negligible.

¹¹There are several useful references to this class of test. The simplest source is from <http://www.gseis.ucla.edu/courses/ed231a1/notes3/covar.html>. See also Takemura, Akimichi and Kuriki, Satoshi—"Maximum Covariance Test for Equality of Two Covariance Matrices"—University of Tokyo, and The Institute of Statistical Mathematics. See also the SAS code in <http://ewe3.sas.com/techsup/download/stat/CovMxTest.html>

¹²Laurent and Gregory [2003] provide a very nice overview of this class of products. See Duffie and Garleanu [2001] for the pricing of CDOs.

¹³For m bonds, assuming Y is normally distributed, the expected number of defaults may be shown (after some lengthy though simple calculations) to be equal to

$$\begin{aligned} E(n) &= \int_Y \prod_{i=1}^m \bar{s}_i(Y) dY \\ &= m - \exp \left[\left(\sum_{i=1}^m (v_i^*)^2 \right) (1 - \rho^2) - \left(\sum_{i=1}^m \theta_i \right) \right. \\ &\quad \left. - \left(\sum_{i=1}^m v_i^* \right) \rho \mu Y + \frac{1}{2} \left(\sum_{i=1}^m v_i^* \right)^2 \rho^2 \sigma_Y^2 \right] \end{aligned}$$

¹⁴The number of firms is less than that totally available in each subperiod (see, Exhibit 3) because we have combined subperiods and only used firms with full data across sub periods.

¹⁵Since the analysis is based on raw data, the choice of lognormal joint distribution for the intensities ensures positive values. Any other joint distribution with positive support over intensity values is also possible.

REFERENCES

- Allen, L., and A. Saunders. "A Survey of Cyclical Effects in Credit Risk Measurement Models." Working paper, New York University, 2002.
- Altman, Edward. I., Brooks Brady, Andrea Resti, and Andrea Sironi. "The Link Between Default and Recovery Rates: Implications for Credit Risk Models and Procyclicality." *Journal of Business*, 2002.
- Ang, A., and J. Chen. "Asymmetric Correlation of Equity Portfolios." *Journal of Financial Economics*, Vol. 63, No. 3 (2002), pp. 443–494.
- Black, F., and M. Scholes. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81 (1973), pp. 637–654.
- Campbell, John., and Glen Taksler. "Equity Volatility and Corporate Bond Yields." *Journal of Finance*, 58 (2003), pp. 2321–2349.
- Collin-Dufresne, P., R. Goldstein, and J. Helwege. "Is Credit Event Risk Priced? Modeling Contagion via the Updating of Beliefs." Working paper, Haas School, University of California, Berkeley, 2003.
- Collin-Dufresne, P., and B. Goldstein. "Do Credit Spreads Reflect Stationary Leverage Ratios?" *The Journal of Finance*, Vol. 56, No. 5, 2001.
- Das, S., D. Duffie, N. Kapadia and L. Saita. "Common Failings: How Corporate Defaults Cluster." *Journal of Finance*, 2005.
- Das, S., G. Fong., and G. Geng. "The Impact of Correlated Default Risk on Credit Portfolios." *Journal of Fixed Income*, December 2001, Vol. 11, No. 3 (2001), pp. 9–19.
- Das, Sanjiv., and Gary Geng. "Correlated Default Processes: A Criterion-Based Copula Approach." *Journal of Investment Management*, Vol. 2, No. 2 (2004), pp. 44–70.

- Das, S., and P. Tufano. "Pricing Credit Sensitive Debt when Interest Rates, Credit Ratings and Credit Spreads are Stochastic." *The Journal of Financial Engineering*, Vol. 5, No. 2 (1996), pp. 161–198.
- Das, S., and R. Sundaram. "A Discrete-Time Approach to Arbitrage-Free Pricing of Credit Derivatives." *Management Science*, Vol. 46, No. 1 (2000), pp. 46–62.
- David, Mark., and Violet Lo. "Infectious Defaults." *Quantitative Finance*, Vol. 1, No. 4 (2001), pp. 382–387.
- de Servigny, Arnaud., and Olivier Renault. "Default Correlation: Empirical Evidence." Working paper, Standard & Poors 2002.
- Driessen, Joost. "Is Default Event Risk Priced in Corporate Bonds?" *Review of Financial Studies*, 18 (2005), pp. 165–195.
- Duffie, D. and N. Garleanu, "Risk and Valuation of Collateralized Debt Obligations." *Financial Analysts Journal*, Vol. 57, No. 1 (January–February 2001), pp. 41–59.
- Duffie, Darrell., and David Lando. "Term Structure of Credit Spreads with Incomplete Accounting Information." *Econometrica*, 69 (2001), pp. 633–664.
- Duffie, D. and K. Singleton "Modeling Term Structures of Defaultable Bonds." *Review of Financial Studies*, Vol. 12, No. 4 (1999), (Special), pp. 687–720.
- Duffie, Darrell., Leandro Saita and Ke Wang. "Multi-Period Corporate Default Prediction with Stochastic Covariates." Graduate School of Business, Stanford University 2004.
- Duffee, G. "Estimating the price of Default Risk." *Review of Financial Studies*, 12 (1999), pp. 197–226.
- Erb, Claude., Campbell Harvey, and Tadas Viskanta. "Forecasting International Equity Correlations." *Financial Analysts Journal*, (November/December 1994), pp. 32–45.
- Finger, Chris. "Conditional Approaches for CreditMetrics Portfolio Distributions." *CreditMetrics Monitor*, (April 2004), pp. 14–33.
- Giesecke, Kay. "Correlated Default with Incomplete Information." Working paper, Humboldt University, Berlin 2002.
- Gregory, Jon and Jean-Paul Laurent. "In the Core of Correlation." Unpublished Manuscript 2004.
- Hamilton, J.D. "A New Approach to the Economic Analysis of Nonstationary Time Series and The Business Cycle." *Econometrica*, 57 (1989), pp. 357–384.
- Jarrow, R. and S. Turnbull. "Pricing Derivatives on Financial Securities Subject to Credit Risk." *Journal of Finance*, Vol. 50, No. 1 (March 1995), pp. 53–86.
- Jarrow, R., D. Lando, and S. Turnbull. "A Markov Model for the Term Structure of Credit Spreads." *Review of Financial Studies*, 10 (1997), pp. 481–523.
- Jarrow, R., D. Lando, and F. Yu., "Default Risk and Diversification: Theory and Empirical Implications." *Mathematical Finance*, Vol. 15, No.1 (2005), pp. 1–26.
- Jarrow, R., and Fan Yu. "Counterparty Risk and The Pricing of Defaultable Securities." *Journal of Finance*, Vol. 56, No. 5 (2001), pp. 1765–1799.
- Laurent, Jean-Paul, and Jon Gregory. "Basket Default Swaps, CDOs and Factor Copulas." Working paper, ISFA Actuarial School and University of Lyons, 2003.
- Leland, Hayne., and Klaus Toft. "Optimal Capital Structure, Endogenous Bankruptcy, and The Term Structure of Credit Spreads." *Journal of Finance*, Vol. 51, No.3 (1996), pp. 987–1019.
- Longin, F., and B. Solnik. "Is the Correlation in International Equity Returns Constant: 1960–1990?" *Journal of International Money and Finance*, 14 (1995), pp. 3–26.
- Longstaff, F., and E. Schwartz, "A Simple Approach to Valuing Risky Fixed and Floating Rate Debt." *Journal of Finance*, Vol. 51, No. 3 (September 1995), pp. 789–819.
- Lopez, Jose, A. "The Empirical Relationship between Average Asset Correlation, Firm Probability of Default and Asset Size." Working paper, Federal Reserve Bank of San Francisco 2002.
- Lucas, D. J. "Default Correlation and Credit Analysis." *The Journal of Fixed Income*, (March 1995), pp. 76–87.
- Madan, Dilip and H. Unal. "Pricing the Risks of Default." *Review of Derivatives Research*, Vol. 2, No. 2/3 (1998), pp. 121–160, 449–470.
- Merton, R. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates." *Journal of Finance*, 29 (1974), pp. 449–470.

Mina, Jorge, and Eugene Stern. "Examples and Applications of Closed-Form CDO Pricing." *RiskMetrics Journal*, Vol. 4, No. 1 (2003), pp. 5–34.

Schönbucher, Philipp. "Frailty Models, Contagion, and Information Effects." Working paper, ETH, Zurich, 2004.

Sobehart, J., and R. Stein. "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.

Vassalou, Maria, and Yuhang Xing. "Default Risk in Equity Returns." *The Journal of Finance*, Vol. 59, No. 2 (2004), pp. 831–868.

Xiao, Jerry. "Constructing the Credit Curves: Theoretical and Practical Considerations." Working paper, RiskMetrics, 2003.

Yu, F. "Default Correlations in Reduced Form Models." *Journal of Investment Management*, 2003.

Zhou, C. "An Analysis of Default Correlations and Multiple Defaults." *Review of Financial Studies*, 14 (2001), pp. 555–576.

To order reprints of this article, please contact Dewey Palmieri at dpalmieri@iijournals.com or 212-224-3675.

©Euromoney Institutional Investor PLC. This material must be used for the customer's internal business use only and a maximum of ten (10) hard copy print-outs may be made. No further copying or transmission of this material is allowed without the express permission of Euromoney Institutional Investor PLC.

The most recent two editions of this title are only ever available at <http://www.euromoneyplc.com>