Unresolved Issues in Modeling Credit-Risky Assets

STUART M. TURNBULL

STUART M. TURNBULL is Bauer Chaired Professor in the Bauer College of Business at the University of Houston in Houston, TX.
sturnbull@uh.edu

To quantify the risk of a portfolio of credit-risky instruments for regulatory purposes, it is necessary to describe the probability distribution of possible portfolio values at some specified horizon. This involves two distinct steps. First, it is necessary to determine the value of the portfolio at the horizon in a given state of nature, implying that we must have a valuation model that can price the different types of credit-sensitive instruments in the portfolio. The second step is to describe the probability distribution of portfolio values under the natural probability measure.

Two different approaches can be used to price credit-sensitive instruments: either structural-based models or reduced-form models. In either case, two natural questions arise. First, can the model describe the current term structure of spreads, and, second, if the model is used for dynamic hedging, can it accurately describe the evolution of spreads?[1]

Regulators also raise questions about the testing of a credit pricing model, given the limited default data. All models must be calibrated. For loans and bond-related instruments, bond spreads—if available—are the natural choice, while credit default swap prices are often used for credit derivatives. In either case, information about recovery in the event of default is usually a necessary input. The criteria used for calibration depend on whether the model is simply required to match a chosen set of prices or whether the time series properties of the model are important.

Apart from the issue of pricing individual instruments, it is necessary to describe the probability distribution of a portfolio as it evolves over time. In the original CreditMetrics framework, it is assumed that the credit quality of an obligor could have one of n possible credit ratings including the state of default; the term structure of credit spreads is assumed to be deterministic. At the horizon it is necessary to enumerate every possible combination of credit qualities. The Merton [1974] model can be used, and a multivariate normal distribution for asset returns is assumed. Small changes in correlation are found to have a major impact on the measured risk of the portfolio and its economic capital.

In applying this framework it is necessary to specify a ratings transition matrix over the horizon that reflects the current state of the economy. The modeling of default dependence in this framework is achieved by combining the Merton [1974] structural model with a rating transition matrix.

The original CreditMetrics method of modeling default dependence with a multivariate normal probability distribution is generalized by Li [2000], who introduces a normal copula function to link the individual marginal distributions. Since then copula functions have become the standard way to model default dependence in risk management and for pricing structural products.

In the New Basel Capital Accord, the

foundation internal ratings-based (IRB) approach allows banks to develop and use their own internal ratings. The loss given default (LGD) and exposure at default (EAD) are fixed and based on values set by the regulators. Under the advanced IRB approach, banks can use their own loss given default and exposure at default values, subject to regulatory approval. The flexibility to determine these three parameters so that they reflect the nature of a bank's portfolio gives a bank incentives to move to the advanced IRB approach, although a bank's specification of these parameters is subject to supervisory review.

This raises the question of what constitutes a reasonable approach to estimating expected loss. The issue of model testing is of central importance to financial regulators and to financial institutions adopting the advanced internal ratings-based methodology, so it is important to recognize what we actually know about some of the issues in credit and what issues remain unresolved.

We examine some of the major unresolved areas in credit pricing and risk management, focusing on recent academic findings. We start by examining our current knowledge about modeling two of the important components of expected loss: the loss given default, and the probability of default. To model the loss given default requires that we describe the determinants of the recovery rate. There are a variety of approaches to the modeling of recovery rates and some recent advances in academic default prediction models that can be extended to multiperiod horizons.

Next we discuss current empirical knowledge of the behavior of credit spreads. We assess the ability of structural models to examine the cross-sectional and time series variation in credit spreads and the ability of reduced-form models to examine the time series variation in spreads.

Then we consider some of the problems that can arise in the pricing of single-name credit default swap (CDS) instruments. We first consider how calibration can affect pricing and hedging. We start by considering a digital credit default swap, which has the great advantage that the payoff if a credit event occurs is independent of the recovery rate on the reference entity. Digital credit swaps are only a small part of the credit swap market, so we use regular CDS to calibrate our model for pricing (although here the calibrated price is dependent on the recovery rate).

Second, we consider the specification of the credit event that we are modeling. For regular CDS, the protection buyer owns a cheapest to deliver option. If a credit event occurs, the protection buyer has some choice over the asset to deliver, assuming physical delivery. The set of deliverable obligations must satisfy certain restrictions with respect to maturity and be pari passu, usually with respect to senior unsecured debt. This is especially important if the credit event is a restructuring. With a number of different restructuring clauses, how do we estimate the difference in pricing?

The third problem involves hedging. How do we hedge a simple credit default swap? This is a fundamental question as to the methodological approaches to pricing that assume market completeness.

The fourth issue is counterparty risk and how it affects the pricing of single-name credit default swaps due to default dependence. It automatically arises if we consider a portfolio of credit default swaps. How do we model this type of dependence?

Finally, we address some issues that arise in the pricing of multiname products and in the modeling of the distribution of portfolio values. Modeling default dependence is central to modeling the probability distribution of a portfolio of credit-risky assets, pricing multiname products, and pricing counterparty risk. Yet we have little empirical information about default dependence. Also, how does default of an obligor or news about changes in its creditworthiness affect the creditworthiness of other obligors? The actual modeling of default dependence may be done with structural models and copula functions or by a reduced-form methodology. A final issue is the testing of models.

I. MODELING EXPECTED LOSS

Both the loss given default and the probability of default are key inputs in estimating the expected loss at a specified horizon.[2] We cannot directly test a model of expected loss, but we can test the models used to generate estimates of the loss given default and the probability of default. We start by examining what we know about the determinants of recovery rates, and then we review recent academic models estimating the probability of default.

Recovery Rates

In the event of default, claim holders receive some fraction of their claims, which is referred to as the *recovery rate*. Recovery rate information is a critical input in estimating the expected loss and a critical input for pricing the majority of credit-sensitive instruments.[3]

Most of our information about the determinants of recovery rates comes from actual recovery data that can immediately be used to calculate expected loss. This is not the case for pricing. At present there is little information about how to model recovery rates under the appropriate pricing measure.

There are many issues that need to be addressed in the modeling of recovery rates. First is the definition of default. The definition of default according to the Bank for International Settlements covers a wide range of corporate events:

- Obligor determined as unlikely to pay its debt obligations in full.
- Distressed restructuring involving the forgiveness or postponement of principal, interest, or fees.
- Missed payments of credit obligations.
- Obligor has filed for bankruptcy or similar protection from creditors.

The variety in this list of events implies that, everything else equal, the recovery will depend on the type of event that triggers default. It is this diversity of events and the resulting disparity in recovery that led to the introduction of Modified Restructuring, Modified-Modified Restructuring, and No Restructuring default swap contracts.

The second issue is the measurement of recovery. The recovery rate is often taken as the price of a bond within a one-month period of default. The event of default may affect a bond's liquidity and thus its price, implying that any measure of recovery is affected by the bond's liquidity. The issue of liquidity is usually ignored in any discussion of recovery.

Guha [2002] looks at bond prices just after default. Prices of bonds of the same seniority converge to almost the same price, independent of maturity. Guha argues that the amount recovered is best modeled as a fraction of the face value. Another measure would be the recovery at emergence from default.[4]

Given that restructuring of claims often occurs during bankruptcy, and the absolute priority rule is frequently violated, it is far from clear how to measure recovery. The term of the bankruptcy process is highly uncertain. All these considerations greatly complicate estimation of the present value of the payments to claim holders. In practice, various ad hoc discount rates are used, implying that present value estimates probably have a higher standard deviation than estimates generated by the first method.[5]

Schuermann [2004] and Altman, Resti, and Sironi [2005] list some conclusions about recoveries:

1. Depend on seniority and collateral.
2. Vary over the credit cycle.
3. Vary across industries.
4. Often have a bimodal distribution.

Moody's [2003] documents a monotonic relation between average (and median) recovery rates with priority in the capital structure over the 21-year period 1982-2002.[6] Gupton, Gates, and Carty [2000] and Acharya et al. [2003] find that syndicated loan recoveries are substantially greater than recoveries on senior unsecured debt.

That recovery rates vary with the credit cycle is to be expected. Frye [2000] and Schuermann [2004] document that the distribution of recovery rates is shifted toward the left during a recession, compared to the distribution during an expansion. Frye [2003], Moody's [2003], and Altman, Resti, and Sironi [2005] find a negative relation between aggregate default rates and recovery rates.

Average recovery rates vary across industries. In Moody's [2003] for the period 1982-2002, public utilities had the highest recovery rates, followed by financial institutions, with telecommunications last.

There is often wide cross-sectional variation in recoveries, even after keeping seniority, security, and industry fixed. Altman, Resti, and Sironi [2005] consider aggregate recovery rates. They find that the logarithm of the weighted-average aggregate default rate, changes in the weighted-average aggregate default rate, and the total amount of high-yield bonds outstanding for a particular year are important explanatory variables.[7] Macroeconomic variables are generally not statistically significant.

Acharya, Bharath, and Srinivasan [2003] regress recovery rate, measured by the bond price within one month of default, against a number of variables classified into one of four groups: contract-specific characteristics, firm-specific characteristics, industry-specific characteristics, and macroeconomic conditions. Variables found to be statistically significant at the 1% or 5% level are:[8]

- Contract-specific characteristics: Logarithm of issue size; seniority.[9]
- Firm-specific characteristics: Profit margin (defined as EBITDA/sales one year prior to default); debt concentration (defined as Herfindahl index measured using the amount of the debt issues of the defaulted company).[10]
- Industry-specific characteristics: Industry distress (defined as a dummy variable that equals one if the industry is in distress and zero otherwise).[11]

- Macroeconomic conditions: In the presence of industry-specific characteristics macroeconomic variables are insignificant.

If a firm has been struggling for a long time to avoid default, its assets may have been depleted, and hence the recovery rate will be reduced. Researchers do not address whether firm history prior to default influences the recovery rate.

Modeling the recovery rate to depend on firm-specific and macroeconomic variables is useful for risk management purposes, but for such models to be useful for pricing, the probability measures must be changed to those used for pricing. Apart from the work of Bakshi, Madan, and Zhang [2001], there is little guidance as to the appropriate modeling of recovery under the pricing measure.

Unresolved issues that remain include:

1. Whether the history of a firm's condition prior to default influences the recovery rate.
2. We have no information about the out-of-sample performance of models. If models are going to be used for regulatory capital reasons or for pricing, we need to test their accuracy.
3. There is little guidance as to the modeling of recovery rates under the appropriate pricing probability measure, given that recovery rates depend on systematic factors.

New Developments in Default Prediction Models

Hillegeist, Keating, and Cram [2004] use a discrete duration model to examine the relative performance of a Merton [1974] model compared to the Altman [1968] Z-score and Ohlson [1980] O-score models. They find that the Merton-based model outperforms the accounting-based models. Interestingly, its performance is relatively insensitive to the precise specification of the critical liability value used in the distance to default variable. This variable does not entirely explain the variation in failure probabilities across firms. Chava and Jarrow [2003], Bharath and Shumway [2004], and Campbell, Hilscher, and Szilagyi [2004] reach a similar conclusion.

Janosi, Jarrow, and Yildirim [2002] specify an intensity function depending on the spot default-free interest rate and the cumulative unanticipated rate of return on the market index. They estimate the stochastic processes describing the procedure. The probability of default over a finite horizon will depend on the parameters of these variables.

Chava and Jarrow [2003] use a number of standard accounting variables and a number of market variables: relative size, excess return on equity, and stock's volatility. Relative size is defined as the logarithm of the value of the firm's equity divided by the total market value of equity. Excess return is defined as the monthly return of the firm's equity minus the value-weighted market index return. The firm's equity volatility is computed as the standard deviation over the last 60 days of observable daily market prices. Industry dummy variables are also found to be statistically significant. The Chava and Jarrow model provides a small improvement in out-of-sample performance over the Shumway [2001] model.

Duffie and Wang [2003] use a reduced-form methodology with five covariates: distance to default, personal income growth, sector performance, firm-level earnings, and firm size. They estimate the underlying stochastic processes for these state variables. Sector performance, measured by sector average across firms of the ratio of earnings to assets, is found to be statistically insignificant when the first two covariates are used and is hence dropped from the analysis. Duffie and Wang do not test the performance of their model against other models.

The models of Janosi, Jarrow, and Yildirim [2002] and Duffie and Wang [2003] can all be applied to estimating multiperiod default probabilities, which is an alternative methodology to estimating state-dependent probability transition matrices.[12]

Primary among unresolved issues is that all prediction models use accounting and macroeconomic variables as basic inputs. While a firm's reported accounting variables must satisfy accounting standards, there are many unidentified assumptions applied in the production of these numbers. Many factors such as managerial ability, firm culture, and ability to respond to changes in market and production conditions are also imperfectly measured, if measured at all, by the types of covariates employed. And there is no empirical work that attempts to model contagion arising from information effects if investors are imperfectly informed about some common factors that affect the true probability of default.

Expected Loss

Several studies directly address the dependence of recovery and default rates by assuming that both depend on a single common factor. Frye [2000] describes the determinants of the probability of default using the single-factor model suggested by Finger [1999] and Gordy [2000]. The rate of return on the asset value of the j-th firm, denoted by $A_j(t)$, is assumed to depend in a linear way on a common factor:

$$A_j(t) = pX(t) + (1 - p^2)^{1/2}e_j(t)$$

where p is the correlation between the asset return and a common factor; X(t) is the common factor; and $e_j(t)$ is an idiosyncratic factor. It is assumed that X(t) and $e_j(t)$ are independent, normally distributed random variables. Default occurs if the asset return is below some critical value.

The rate of recovery is described by

$$R_j(t) = \mu_j + \sigma q X(t) + \sigma(1 - q^2)^{1/2}Z_j(t)$$

where $Z_j(t)$ is a idiosyncratic factor. It is assumed that X(t) and $Z_j(t)$ are independent, normally distributed random variables, which implies that the recovery rate is also normally distributed. The standard deviation of the recovery rate is σ, and q is a measure of the sensitivity of the recovery rate to the macro factor X(t). The term μ_j depends on the seniority of the debt.

Using a two-stage estimation procedure, Frye [2000] finds evidence of a negative correlation between default and recovery rates. In a similar-one factor framework, Schönbucher [2003a, pp. 147-150] assumes that recovery rates are described by a logit normal transformation, while Pykhtin [2003] assumes a lognormal distribution. He derives closed-form expressions for the expected loss and the economic capital.[13]

These three specifications are tested in Dullmann and Trapp [2004], assuming homogeneity among obligors.[14] They find that the normal distribution and the logit normal transformation provide a better fit than the lognormal distribution. They also find empirical evidence of a systematic factor and negative correlation between recovery rates and default rates.

An unresolved issue is that the Dullmann and Trapp [2004] results do not favor the logit normal transformation over the normal distribution. Given that recovery rates are defined over a finite interval, usually between zero and one, their finding suggests model misspecification.

We know from the work of Acharya, Bharath, and Srivinisan [2003] and Altman, Resti, and Sironi [2005] that recovery rates depend on more than a single systematic factor. We need to model the probability of default and recovery rates in a multifactor setting.

II. WHAT DO WE KNOW ABOUT CREDIT SPREADS?

In many applications of the reduced-form methodology, the intensity function is calibrated to match the term structure of the credit spread. This assumes that the whole spread is generated by credit risk and consequently raises the question as to what we know about credit spreads. Many structural models form the basis for default prediction and for modeling default dependence. One of the ways to test these models is to examine their ability to describe the properties of credit spreads.

We first briefly describe some of the known empirical properties of credit spreads, and then discuss how structural and reduced-form models can describe the different properties of credit spreads.

Empirical Evidence

Elton et al. [2001] argue that three factors affect corporate bond spreads: expected loss, a tax premium, and a risk premium. There is a tax premium because corporate bonds are taxed at the state level, while interest income on government bonds is not. There is a risk premium because credit spreads depend on systematic factors. Elton et al. [2001] use a simple model that assumes risk-neutrality, stationary default probabilities, and constant recovery rates to estimate the spread due to default. Estimated spreads are found to be substantially lower than observed spreads.

Expected loss accounts for only a small part of observed credit spread, while state taxes explain a substantial portion of the difference. Factors in a three-factor Fama-French model go a long way toward explaining the differences.

Collin-Dufresne, Goldstein, and Martin [2001] (CGM) argue that a number of variables should affect the change in credit spreads:

- Changes in the spot rate (negative relationship).[15]
- Changes in the slope of the Treasury yield curve (negative relationship).[16]
- Changes in leverage (positive relationship).
- Changes in volatility of the assets of the firm (positive relationship).
- Changes in the probability or magnitude of a downward jump in the value of the firm (positive relationship).

- Changes in the business climate (negative relationship).[17]

From their empirical analysis they find:

- A negative correlation between the spot interest rate and credit spreads.
- Insignificant convexity and slope that are often of the wrong sign.
- Significant volatility.
- Economically and statistically significant S&P 500 index return.
- Significant changes in the S&P smirk.

Financial and economic variables explain only 25% of observed credit spread changes. A principal components analysis on the residuals shows that they are mostly driven by some unspecified single factor. An expanded regression, which included 1) some liquidity measures, 2) some non-linear spot interest rate terms, and 3) some lagged variables could only explain about 34% of the spread changes. A principal components analysis still indicated that residuals were highly correlated.

To check that their finding of a systematic factor is not spurious, CGM use both the Longstaff-Schwartz [1995] and the Collin-Dufresne and Goldstein [2001] models to generate a data set of credit spreads. They regress changes in spreads against changes in leverage and spot interest rates, and find, contrary to their empirical findings, that change in leverage is a major explanatory factor, with an R-square of about 90%. CGM conclude that it is difficult to reconcile their empirical findings with the structural models we have. While models may be able to fit the current term structure of spreads, they cannot explain the time series variation.

The view that credit risk explains only a small part of credit spreads is challenged by Longstaff, Mithal, and Neis [2003]. Using credit default swap and corporate bond data to jointly estimate the expected loss under the pricing measure, they find that the proportion of the instantaneous spread explained by credit risk is 62% for AAA/AA-rated bonds, 63% for A-rated bonds, 79% for BBB-rated bonds, and 89% for BB-rated bonds.[18] They also show that the proportion of the spread not explained by credit risk appears to be related to different measures of liquidity. There is only weak support for the hypothesis that the non-credit risk component of spreads is due to taxation.

Macroeconomic factors also affect credit spreads. Jaffee [1975] finds that different macroeconomic variables can explain a substantial amount of the cyclical variation in corporate bond yields. Altman [1983], looking at aggregate business failure in the U.S. economy, finds four macroeconomic variables that have explanatory power: real economic growth, stock market performance, money supply growth, and business formation.

Wilson [1997] uses macroeconomic variables to explain aggregate corporate default rates across different countries. Yang [2003] shows that macroeconomic variables such as inflation and real activity can explain a significant amount of the variation in corporate bond prices.

There is still considerable uncertainty as to the major determinants of credit spreads. To summarize the empirical evidence, we find:

1. There is a negative correlation between credit spreads and the short-term default-free interest rate.
2. Credit risk explains only part of the credit spread.
3. Credit spreads depend on systematic factors.
4. Credit spreads depend on firm-specific factors.
5. Credit spreads depend on liquidity and possibly taxation.

All these factors should be considered in modeling the evolution of credit spreads.

Structural Models

In the application of structural models, the two different approaches used for calibration reach quite different conclusions. In the first approach, models are calibrated to the term structure of default probabilities under the natural probability measure. The second approach uses firm value, leverage, payout ratio, and estimation of exogenous parameters, such as the default interest rate process.

For the first approach, we review the work of Huang and Huang [2003], who calibrate a wide array of structural models, including jump diffusion models, and generate a term structure of credit spreads for different credit classes. The models are calibrated to match 1) the average probability of default under the natural probability measure over different horizons; 2) the average loss as a fraction of the face value of debt; 3) the average leverage ratio; and 4) the equity premium. These models for investment-grade firms can explain less than 30% of the average credit spread. For high-yield, the models can explain between 60% and 80% of the average credit spread.[19]

A similar conclusion is reached by Turnbull [2003] using a reduced-form model that is calibrated to match the term structure of default probabilities under the natural measure over a five-year horizon. For each obligor, the term structure of default probabilities is generated using KMV's Credit Monitor. The probability measure is changed to the risk-neutral measure (with the default premium assumed to be unity), and a term structure of credit spreads is generated, assuming some average recovery rate. It is found that credit spreads were one-third to one-half of those observed in the market.

In the second approach, Eom, Helwege, and Huang [2004] test five different structural models. They use firm-specific parameters such as firm value, leverage, and payout ratio based on historical corporate data. They find the Merton [1974] model generates spreads that are too narrow. The Geske [1977] compound option model also generates spreads that are too narrow, but it is an improvement over Merton. The Leland and Toft [1996] model overestimates spreads, even for short-maturity bonds. The Longstaff and Schwartz [1995] model generates spreads that are too high on average, including excessive spreads for risky bonds, and underestimates the spreads for low-risk bonds. The Collin-Dufresne and Goldstein [2001] model overestimates spreads on average.

Eom, Helwege, and Huang conclude that structural models do not systematically underpredict credit spreads, but accuracy is a problem. They make no attempt to infer the term structure of default probabilities under the natural measure or to examine the time series properties of their models.

In any application of a structural model that uses accounting information, it is implicitly assumed that such information is accurate. An interesting extension of the structural modeling approach is given in the work of Duffie and Lando [2001], who show that with incomplete accounting information it becomes necessary to model the event of default as an inaccessible stopping time, along the lines of the reduced-form models introduced by Jarrow and Turnbull [1992, 1995]. Yu [2005] provides some empirical evidence relating accounting transparency and credit spreads.

The structural models are overall far from accurate. Unresolved questions are:

1. Calibration of a structural model.
2. Reasons for the inaccuracy of structural models.
3. Modeling of information asymmetry and the implications for instrument dynamics.

Reduced-Form Models

In modeling credit spreads, Duffee [1999] assumes that the risk-neutral intensity function depends on a firm-specific factor and two default-free interest rate factors, all described by square root processes. Credit spreads are found to be weakly negatively correlated with the short-term interest rate. The firm-specific factor and its price of risk are significant. While the influence of a factor varies over credit ratings, the firm-specific factors are found to be correlated across firms, implying a possible misspecification of the intensity function. There is no estimation of the default risk premium.

Janosi, Jarrow, and Yildirim [2000] assume that the intensity function depends on the short-term default-free interest rate and the cumulative unanticipated return on the equity index (see Jarrow and Turnbull [2000] for details). They also include a liquidity term modeled as a convenience yield. Like Duffee, they find a negative relation between spreads and the short-term interest rate. They find the coefficient on the equity market return to be statistically insignificant, contradicting the finding in Collin-Dufresne, Goldstein, and Martin. The different models tested have relatively low R-square values, implying that they can explain only a small part of the variation in credit spreads.

Driessen [2005] models the intensity function as dependent on two short-term default-free interest rate factors, two latent factors common to all firms, and a firm-specific factor. These factors, described by square root processes, are found to be statistically significant. The prices of risk for the two common factors are statistically significant, while the price of risk for the firm-specific factor is close to zero for the median firm. Lower-rated firms are found to be more sensitive to the common factors than higher-rated firms.

The default risk premium, μ, is defined as the ratio of the intensity function under the pricing measure Q to the intensity function under the natural measure P. It is assumed to be constant and identical across firms.[20]

The default risk premium is estimated by equating estimates of the historical probability of default to the probability of default estimated from the model; it cannot be estimated with high statistical precision, even after adjusting for differential taxation and liquidity. The estimate is 2.91 with standard error 1.11 based on S&P data for the period 1981-2000 and 2.15 with standard error 1.03 based on Moody's data for the period 1970-2000.

Unresolved issues with respect to reduced-form models include:

1. How to relate the state variables of the intensity function to firm-specific and macroeconomic variables.
2. The out-of-sample performance of models.
3. In the current formulations, the intensity function is assumed to be a linear function of state variables. Even if the state variables are assumed to follow square root processes, given the empirical estimation, there is no guarantee that the intensity function is positive, which is what we need.

III. ISSUES IN PRICING SINGLE-NAME CREDIT INSTRUMENTS

Many issues that affect the pricing of single-name products must be addressed before it is meaningful to discuss multiname products. We consider some of the issues that arise in the pricing of credit default swaps, perhaps the simplest form of credit derivative; the issues are applicable to all single-name credit-risky derivative instruments.

For any model there is always imprecision in parameter values. Often simple point estimates are used as inputs to price an instrument, implying that the estimated price is an implicit function of the point estimates. For many credit derivatives, the recovery rate is an input parameter. The usual assumption is to treat the recovery rate as a constant, although it is well known that recovery rates are stochastic and vary over the credit cycle.

The first issue we consider arises because of input parameter imprecision and its effect on calibration. The second issue is application of the standard model to pricing credit default swaps (CDS), when there is more than one type of credit event. We must clearly identify the stochastic processes for the different types of credit events and recognize that cash and default swap markets may depend on different factors. (Recall that defining what constitutes a credit event arose in our discussion of recovery rates.)

The third issue we discuss is the question of hedging. We argue that the standard static hedge is only an approximation. Arbitrage-free models assume markets are dynamically complete, and hence hedging and pricing are intrinsically related. Again we argue that because of different definitions about what constitutes a credit event, uncertain recovery rates, and a lack of hedging instruments, markets are also dynamically incomplete.

The last issue we address is counterparty risk. The credit default swap market is an interbank market, so counterparty risk is always present. This affects pricing and hedging in the market.

Calibration and Pricing of Credit Digital Swaps

A protection buyer promises to make payments, S_D, to the protection seller at dates $T_1, \ldots, T_n$, provided there is no credit event. The present value of these payments is given by:[21]

$$S_D E_t[PV_{PB}(t)] \equiv S_D \sum_{j=1}^{n} \times E_t[B(t, T_j) \perp (\Gamma_{RE} > T_j) \mid F_t]$$

where $B(t, T_j)$ is the discount factor; $t < T_1$; Γ_{RE} is the time of default by the reference entity; and F_t is the information set at time t. If there is a credit event, the protection buyer must pay accrued interest (a complication we ignore).[22]

At the time of a credit event, the protection seller makes a payment of L to the protection buyer. Without loss of generality, we set $L = 1$. The present value of the protection leg is given by

$$E_t[PV_{PS}(t)] \equiv E[B(t, \Gamma_{RE}) \perp (T_n \geq \Gamma_{RE}) \mid F_t]$$

where $T = T_n$. The price, S_D, is set so that the present value of the premium leg equals the present value of the protection leg:

$$S_D E_t[PV_{PB}(t)] = E_t[PV_{PS}(t)]$$

Note that the price does not depend on the recovery rate for the reference entity. How do we actually price this instrument?

The credit digital swap market is a small market, and the usual approach to pricing is to deduce the intensity function from the prices of regular default swaps. The present value of premium payments is given by $S_{CDS}E_t[PV_{PB}(t)]$, and the present value of the protection leg is given by $E_t[(1 - \delta)\, PV_{PS}(t)]$ where δ is the recovery rate on the reference entity. This is determined by the cheapest to deliver bond (we ignore the optionality associated with the cheapest to deliver option).

The price is determined by again equating the present value of the premium payments and the present value of the protection payment:

$$S_{CDS} E_t[PV_{PB}(t)] = E_t[(1 - \delta) PV_{PS}(t)]$$

Using this last expression, we can relate the premium of a regular swap to the premium on the digital swap:

$$\frac{S_{CDS}}{E_t(1-\delta)} + \frac{\text{cov}_t[\delta, PV_{PS}(t)]}{E_t(1-\delta)E_t[PV_{PB}(t)]} = S_D$$

Frye [2000], Acharya, Bharath, and Srinivasan [2003], and Altman et al. [2003] find a negative correlation between the recovery rate and the frequency of default. This is under the natural probability measure. Assuming a negative correlation holds under the pricing measure implies that:

$$S_{CDS} \geq S_D E_t(1-\delta)$$

Taking the value of S_{CDS} as given, the spread for the digital swap is inversely related to our estimate of the expected loss. There is normally a lot of uncertainty about the expected loss, and small changes in the expected loss can produce large changes in the estimated digital premium. This problem is compounded with differences in the definitions of what constitutes a credit event in the cash and swap markets. For instruments where we need the recovery rate and do not want to assume it is a known constant, we have very little information under the pricing measure Q. The method of calibration can cause false dependence.[23]

Defining the Credit Event: Restructuring

Defining a credit event affects default prediction models, the measurement of recovery rates, and the pricing and hedging of credit-risky bonds and credit derivatives. The last two years have seen a major attempt to clarify the definition of a credit event and the resulting recovery rate. Our discussion draws on material in O'Kane, Pedersen, and Turnbull [2003].

In the current standard default swap contract linked to a corporate (non-sovereign) reference credit, three credit events can trigger the payment of protection: bankruptcy, failure to pay, and restructuring. Both bankruptcy and failure to pay are called *hard* credit events because following either one, pari passu bonds and loans of the issuer should trade at or very close to the same price. We also use the term *default* to refer to a hard credit event.

Restructuring is different, and is characterized as a *soft* credit event. Restructuring does not make a company's debt immediately due and payable, so the assets of a company will generally continue to trade, with long-dated assets trading below short-dated assets.

In 2003, the mechanism for settling a default swap following all three credit events was the same. In the event of a credit event, protection buyers would settle by delivering the face value amount of deliverable obligations to the protection seller in return for the face value amount paid in cash. The market standard was to specify a basket of delivery obligations consisting of bonds or loans at least pari passu with the reference obligation. There was usually a limit of 30 years on the maximum maturity. Protection buyers would choose from this basket of deliverables; they are effectively long a cheapest to deliver option.

The 2003 International Swaps Dealers Association definitions state that a debt obligation is considered restructured if there is:

1. An interest rate reduction.
2. A reduction in principal or premium.
3. A postponement or deferral (maturity extension).
4. A change in the priority ranking of payments.
5. A change in currency or composition of payment of principal or interest.

In addition, more than three unaffiliated holders must hold the restructured obligation, and at least two-thirds of the holders must have consented to the restructuring. Note that a restructuring is viewed as a credit event if there is deterioration in the creditworthiness or financial condition of the reference entity.

Following the release of the ISDA 2003 definitions, there are now four types of restructuring clauses: Old Restructuring (Old-R), Modified Restructuring (Mod-R), Modified-Modified Restructuring (Mod-Mod-R), and No Restructuring (No-R).[24] Note that Mod-Mod-R allows a greater range of deliverables than Mod-R. The standard contract in the U.S. is Mod-R, and in Europe, Mod-Mod-R.

For each type of contract, a credit event will affect the recovery rate. To price the contracts, it is necessary to model the probability of an occurrence of a credit event and the resulting recovery rate—see O'Kane, Pedersen, and Turnbull [2003].

A major challenge is the calibration of pricing models. A number of questions must be addressed, including the probability of a restructuring and the impact of a restructuring on the pricing of bonds and hence the cheapest to deliver bond. The current prices of bonds for a particular obligor should reflect the possibility of 1) default, 2) a restructuring event, and 3) default after restructuring. It is impossible to infer all these parameters from market prices.

Hedging a Credit Default Swap

Pricing theory assumes a unique equivalent martingale measure. This can be achieved either in equilibrium or by arbitrage. If we rely on arbitrage, then we must be able to hedge a credit default swap. But how?

A static hedge argument is traditionally used to price a default swap, although the hedge is only an approximation. If static hedging fails, this leaves dynamic hedging. We argue that dynamic hedging also fails.

The premium payments in a default swap contract are defined in terms of a default swap spread, S, which is paid periodically on the protected notional until maturity or a credit event. It is possible to show that the default swap spread can, to a first approximation, be replicated by a par floater bond spread (the spread to LIBOR at which the reference entity can issue a floating-rate note of the same maturity at a price of par) or the asset swap spread of an asset of the same maturity, provided it trades close to par.

Suppose an investor buys a par floater issued by the reference entity with maturity T. The investor can hedge the credit risk of the par floater by purchasing protection to the same maturity date. Suppose this par floater (or asset swap on a par asset) pays a coupon of LIBOR plus F. Default of the par floater triggers the default swap, as both contracts are written on the same reference entity. With this portfolio the investor is effectively holding a default-free investment, ignoring counterparty risk.

The purchase of the asset for par may be funded on the balance sheet or on repo—in which case we make the assumption that the repo rate can be locked into the bond's maturity. The resulting funding cost of the asset is LIBOR plus B, assumed to be paid on the same dates as the default swap spread S. O'Kane and Turnbull [2003] show the default swap spread, S, is given by

$$S = F - B$$

The hedge is not exact, as it relies on several assumptions that do not hold in practice. The result is only an approximation, resulting in small but observable differences.

If a static hedge is limited, what about dynamic hedging? To hedge a credit swap, we form a replicating portfolio. Suppose that we are long protection; the value of the credit swap is the difference between the value of the protection minus the cost of the protection. To replicate this portfolio, we must replicate the changes in value without default and the change in value when default occurs. This requires two instruments written on the reference entity and the default-free instrument.

Being long protection is equivalent to being short the reference entity. It is not easy to short corporate names, and the recovery rates may differ unless we short the cheapest to deliver bond. We could use credit swaps of different maturities to form the replicating portfolio, but the secondary credit swap market is not a liquid market.

Our inability to replicate a credit swap raises questions about pricing credit derivatives and the minimum price to charge for such instruments.

Counterparty Risk

Counterparty risk arises when there is a risk that the writer of a contract might default and thus be unable to perform. This affects the pricing of the contract, as now it is necessary to consider the default risk of the counterparty and the default dependence between the reference entity and the counterparty. Jarrow and Turnbull [1992, 1995] were the first to consider the effect on pricing. If the reference entity and the counterparty are independent, a relatively simple expression can be derived.

A credit default swap is a bilateral contract, with different exposures to the counterparty for the protection buyer and protection seller.[25] There are two facets to counterparty risk: default risk, and mark-to-market risk. If the protection buyer defaults, the protection seller suffers 1) at most one payment of the swap premium and 2) the cost of replacing the protection buyer—a mark-to-market risk.

If the protection seller defaults, there are two cases to consider. First, if the reference entity defaults first and the protection seller defaults before the settlement of the contract, the protection buyer has no credit protection. Second, if the protection seller defaults first, it is necessary to replace the protection seller. This can be particularly costly if the credit quality of the reference entity has deteriorated.

To determine the effects of counterparty risk, it is necessary to specify the default dependence between the reference entity and the protection buyer and seller. This raises the question of modeling default dependence. To compute the mark-to-market risk, it is necessary to price a replacement contract at the time of default by the counterparty. Given that default by the counterparty can occur at any time over the life of the contract, this raises a computation problem, especially for structured products.[26]

In summary, imprecise parameter values can have unforeseen consequences when we calibrate prices of instruments, even those that do not explicitly depend on the particular parameter. Different definitions about what constitutes a credit event affect the recovery rate and pricing. While it is possible to model this diversity, the challenge is to calibrate such a model. Markets are not complete, and eventually we must face the issue of pricing in incomplete markets and data availability.

Unresolved issues in this respect include:

1. How to incorporate event definition diversity into a practical model.
2. How to incorporate pricing in an incomplete market, given data availability.
3. How to incorporate a standard approach to address the effect of counterparty risk on pricing.

IV. PRICING MULTINAME PRODUCTS

Pricing multiname products is complicated because we must model how the credit spreads of the different obligors evolve over time and the effects of the default of an obligor on the remaining obligors. An obligor in a particular sector may experience some form of difficulty that affects the market's assessment of its creditworthiness. Other obligors in the same sector may also see their credit spreads widen simply because they are in the same sector—guilt by association.

The academic literature proposes a number of different approaches: the structural approach, the reduced-form methodology, and the use of copulas. Practitioners tend to use a combination of these approaches. We review these different approaches and raise questions about testing of models and the effect of data limitations.

We start with a discussion of what we actually know about default dependence.

What Do We Know Theoretically About Default Correlation?

Zhou [2001] extends the typical barrier framework to consider two obligors. By assuming either that the critical value is a constant or exponential, a closed-form expression can be derived for the joint probability of default.[27]

Zhou derives a number of conclusions, assuming that the expected rate of return on the asset equals the growth rate of the critical value.

1. The default correlation and asset correlation are of the same sign. This is an intuitive result. If two assets are positively correlated, and one firm defaults, then the value of the other firm's assets has probably declined and the probability of default is increased.
2. Default correlations are generally low over short horizons; then they first increase and then decline. Over short horizons, defaults are rare and primarily idiosyncratic. This explains why default correlations are generally low over short horizons. Given the assumption that the asset's expected rate of return equals the growth rate of the barrier, the probability of default converges to one as the time horizon lengthens. Consequently, the default correlation converges to zero.
3. Firms of high credit quality tend to have low default correlation over a typical horizon. For high-credit quality firms, the initial probability of default is low, and hence the default correlation is low. Default by one firm signals that the value of another firm may have declined, but over short intervals the probability of default is still very low.
4. The time of peak default correlation depends on the credit quality of the underlying firms. High-credit quality firms take longer to peak. (The intuition here is similar to that in conclusion (1).)
5. Default correlation is dynamic, as the credit quality of firms changes over time.

The last conclusion has important practical implication: Default correlations can change even though the underlying processes for asset values are stationary. This highlights the need for care in developing internally consistent models for default dependence.

By using a simple model, and assuming that the distributions for two firms are bivariate lognormal, Zhou derives a number of interesting conclusions that do not depend on modeling the credit cycle of the economy.

What Do We Know Empirically About Default Correlation?

Servigny and Renault [2002] (S&R) analyze all defaults and credit rating changes recorded by Standard & Poor's for a 21-year period from 1981 through 2001, using U.S. annual data. For each year they estimate probabilities of default and rating transitions. Average default correlations are 0.1% per year for investment-grade firms and 1.7% for non-investment grade firms.[28] There is wide dispersion across different industries.

Servigny and Renault use a one-factor model first described by Merton [1974] (and extended in CreditMetrics) to calibrate the critical asset values to match the estimated probabilities. Assuming multivariate normality, they derive an expression for the joint probability of default for two obligors. This bivariate distribution depends on the correlation of asset returns. With an estimate of the joint probability, one can infer the implied asset correlation and calculate the default correlation.[29]

This approach allows S&R to address the relation between equity default correlation (the default correlation using equity returns) and asset default correlation. They use 60 months of equity returns to estimate the equity correlation, and regress the default correlation from equity as the independent variable against the asset default correlation. Finding a low intercept of 0.0027 and a slope coefficient of 0.7794 leads S&R to conclude that 1) equity default correlation is greater than the asset default correlation, and 2) equity default correlation is a poor proxy for the asset default correlation, given a low R-square value of 0.2981.

To test the bivariate normal distributional assumption, S&R also use a t-copula. Keeping the same marginal distributions, they estimate the degrees of freedom. In 88% of their cases, the joint probability of default has an implied degree of freedom greater than 50, implying minimal differences between using either a normal or t–copula.

S&R also examine the effect of time horizon on the estimated default correlation. They find in general that the estimated default correlation increases as the horizon is lengthened from one to three to five years. The one exception is the empirical estimate of the default correlation between the automotive and financial sectors. The finding that default correlation increases is consistent with the theoretical work of Zhou [2001], where default correlation can increase even with constant asset correlation.

S&R also examine whether asset correlation increases with horizon. Given estimated joint probabilities of default over different horizons, they calculate inferred asset correlation, and argue that asset correlations increase with horizon. The implied asset correlation also changes over the business cycle.

Das et al. [2002] use U.S. data supplied by Moody's, covering the period January 1987 through October 2000. For each firm, Moody's RiskCalc model is used to generate monthly estimates of the probability of default over a one-year horizon. A default intensity for each obligor, $\lambda(t)$, is calculated, given an estimate of the probability of default.[30]

It is shown that:

- Default intensities of individual firms within a given rating class vary substantially over time, and the variation depends on the initial credit class and industry sector.
- There appears to be a systematic component to default intensities.
- Correlations between default intensities vary through the business cycle.
- Firms of the highest credit quality have the highest default correlations, contradicting Zhou [2001].

A simple regime model is used to model this dependence. A more detailed model is described in Das and Geng [2002].

Results of this type of study must be treated with caution. The raw data for the study are generated by a black box, a proprietary model developed by Moody's, so the findings reflect properties of the model used by Moody's and not necessarily the raw data. Each raw observation is an estimate of the true value.

With a "user beware" caution about both the Das et al. [2002] and the Servigny and Renault [2002] studies, we can summarize their conclusions as follows:

1. Default and asset correlations are positively correlated.
2. Default correlations vary cross-sectionally, even after controlling for industry and initial credit rating.
3. Default correlations are time-varying, even if we have stationary asset distributions.
4. Default correlations increase with time horizon.

Effect of a Default on Other Obligors

It is generally assumed that the effect of a default by an obligor will have a negative effect on the remaining obligors. The economy may be so poor that the failure of one firm sends out a dire signal about conditions within an industry. Presumably investors will become more pessimistic about the creditworthiness of related firms. Then again, investors may be all too aware of poor economic conditions within a sector, and default by one firm may convey little or no new information about the remaining obligors.

The state of the economy and different competitive conditions in different sectors affect both the probabilities of default and the effects of a default on the creditworthiness of the remaining obligors. It is well known that defaults tend to increase as the economy goes into recession, and different sectors will have different responses to a given state of the economy. In a competitive industry, the default of any one obligor may have short-run positive effects for other firms in the industry before firms adjust their production or new firms enter the industry. In other sectors with different industrial structures, the effects of a default may be beneficial, as there will be less competition and remaining firms can pick up the slack. Consequently, if we want to model default dependence, we should address how the creditworthiness of an obligor depends on the state of the economy and the nature of different sectors.

Davis and Lo [1999, 2001] assume in the case of default by one obligor that the default intensities of the other obligors are increased by a factor called the *risk enhancement factor* that is assumed to be greater than one. This is a "guilt by association" sort of model. It is also assumed that after an exponentially distributed time default intensities return to their original level.

Davis and Lo call their model *macroeconomic shock*, a misnomer, given their assumption that the probability of direct default and the probability of a jump, given a default, are constants independent of the state of the economy. Perhaps a more appropriate name is *macro shock*. While their model is interesting to practitioners, Davis and Lo do not discuss how to implement their model using market data.

In a similar vein, Jarrow and Yu [2001] consider large and small firms. If a small firm defaults, this has no effect on the creditworthiness of the remaining firms, while if a large firm defaults, the creditworthiness of the remaining firms is assumed to deteriorate. It is difficult to generalize this approach to multiple sectors, and the authors do not discuss implementation.

Gagliardini and Gourieroux [2003] and Schönbucher [2003b] introduce the idea of *unobservable heterogeneity* or *frailty* into the modeling of information-driven contagion. Frailty, which represents the effects of unobservable variables that affect the hazard function, is represented as a random variable that multiplies the hazard function. In this approach, default contagion arises from information effects, given investors are imperfectly informed about parameters that affect the hazard function.[31]

Beliefs about the frailty parameter are updated when a credit event occurs, and default intensities of remaining obligors jump as a result. This affects the level of default dependence among obligors.

While such an approach provides a simple way to model how a sector is affected by a credit event and how credit spreads will recover, given no events, to date there is no published empirical work using this specification in the area of credit.

The major unresolved issue in this regard is that at present we have no model that 1) recognizes the industrial structure of a particular sector and 2) can be calibrated and applied to real data.

Structural Approach

The structural approach forms the basis for the current regulatory approach to modeling credit risk. According to the Merton [1974] model, for two obligors the probability of defaulting over the horizon T is given by:

$$\Pr(\Gamma_1 \leq T \text{ and } \Gamma_2 \leq T) = N_2[N^{-1}(p_1), N^{-1}(p_2), \rho] \quad (1)$$

where N(.) is the normal distribution function; ρ is the correlation between the rates of return on the firm/asset values; $N_2(., ., .)$ is the bivariate normal distribution function; and p_j, j = 1, 2, denotes the probability of obligor j defaulting over the horizon. Note that the Merton model does not consider any forms of market frictions such as corporate or personal taxation, so there is no difference between firm value and asset value.

The last equation, generalized to an arbitrary number of assets, underlies the methodology used in CreditMetrics. The correlation matrix is estimated using a factor model of the form:

$$R_j = \sum_{k=1}^{F} \beta_{jk} I_k + \varepsilon_j$$

where R_j is the rate of return on equity for the j-th name, F is the number of factors, I_k is the rate of return on some traded index, and ε_j is the idiosyncratic risk. Various forms of this factor model are employed to accommodate global portfolios. For KMV the correlation matrix is for asset returns, while for CreditMetrics it is for equity returns.[32]

In the Merton [1974] model, default can occur only when the zero-coupon debt matures. An alternative approach to circumvent this limitation is to assume that if the value of the firm's assets falls below some critical value, $K_j(t)$, default occurs. The time to default for firm j is defined as:

$$\Gamma_j = \inf\{t; V_j(t) \leq K_j(t) \mid V_j(0) > K_j(0)\}$$

For this approach to be operational, it is necessary to specify the critical asset value $K_j(t)$.

Usually it is assumed that $K_j(t)$ is deterministic. When the credit quality of the obligor is deteriorating or the economy is going into recession, it is usually harder to find funding, implying that the critical value will in general be state-dependent. The standard assumptions are: $K_j(t)$ is a constant; $K_j(t)$ is exponential; or $K_j(t)$ is a deterministic function calibrated to match some market observables, such as the term structure of credit spreads or prices of credit default swaps.[33] Extending this approach to more than two assets—see Zhou [2001]—is complicated, and closed-form solutions do not seem to exist.

One objection to the Merton/CreditMetrics approach is the assumption that asset values are lognormally distributed. In practice, equity returns are often used to estimate the correlation matrix. In this case, it is assumed that equity returns are lognormally distributed, an assumption that is inconsistent with the Merton model. It is well known that equity return distributions are fat-tailed. This has led to the use of different types of copula functions to model the multivariate distribution of stopping times.

Copula Approach

The copula approach takes the marginal distributions as given and stitches them together with a copula function in order to specify the joint distribution. This stitching is achieved using Sklar's theorem, which describes the conditions for the application of a copula function.

Examples of Copulas. Two frequently employed copula functions are the bivariate normal copula and the t-copula. The normal copula function is defined by

$$C(u_1, u_2) = N_2[N^{-1}(u_1), N^{-1}(u_2), \gamma]$$

where γ is the correlation coefficient.

The joint distribution of default times can be written as:

$$\Pr[\Gamma_1 \le t_1 \text{ and } \Gamma_2 \le t_2] = N_2[N^{-1}[F_1(t_1)], N^{-1}[F_2(t_2)], \gamma] \tag{2}$$

Comparing expressions (1) and (2), we see that $F_j(t_j) = p_j$ and $\gamma = \rho$. This allows us to identify the correlation parameter.

The t-copula function is defined by

$$C(u_1, ..., u_n) = t_{v,\Sigma}[t_v^{-1}(u_1), ..., t_v^{-1}(u_n)]$$

where $t_{v,\Sigma}$ is a multivariate standard t distribution with v degrees of freedom, t_v is a univariate standard t-distribution, and Σ is the correlation matrix. The joint distribution of default times can be written as

$$\Pr[\Gamma_1 \le t_1 \text{ and } ...., \Gamma_n \le t_n] = t_{v,\Sigma}[t_v^{-1}[F_1(t_1)], ..., t_v^{-1}[F_n(t_n)]]$$

To apply this model, we need to specify the degree of freedom, and we need to estimate the correlation matrix for either approach. How do we estimate these parameters?

Standard Application of the Copula Approach. For standard application of the copula approach, see Mashal and Naldi [2001] and Schönbucher [2003a]. In practice, the structural and reduced-form approaches are combined with the copula methodology. The Merton model is used to specify the distribution of stopping times implying that it is necessary to specify the correlation matrix of asset returns. Asset returns cannot be observed so equity returns are used as a proxy.

It is sometimes argued that different degrees of leverage make it inappropriate to use equity returns to model default dependence, a problem that is particularly severe for low-quality obligors. The Merton model refutes this type of argument.[34]

Under a normal copula, which is the industry norm, the basic assumptions are:

1. Univariate default times are calibrated to market data, using a simple reduced-form model calibrated to credit default swap data.
2. Univariate equity returns follow a normal distribution.
3. Dependence is modeled as a normal copula.
4. Dependence structure among default times is identical to the dependence structure among equity returns.

Mashal, Naldi, and Zeevi [2003] investigate two questions: 1) comparing use of equity and asset returns, and 2) the assumption of a normal copula. Using equity and asset returns generated by the KMV model, they argue that as there is little difference in the results, it is immaterial whether one uses equity or asset returns. This is an encouraging result, although it is not surprising.

Asset returns cannot be directly observed. KMV starts with equity returns to generate a time series of asset returns using the Merton [1974] model. Consequently, its asset-generating process implicitly assumes that asset and equity returns are instantaneously perfectly correlated.[35]

Mashal, Naldi, and Zeevi compare a normal copula to a t-copula and provide empirical evidence that the t-copula provides a more appropriate description of equity returns. This has a significant impact on the pricing of multiname credit instruments.

There are a number of limitations to the current application of the copula methodology for modeling default dependence. The hypothesis is that the description of the dependence among equity returns can be used to describe dependence among stopping times, given the copula function.

First, we have no evidence justifying this type of hypothesis. It might be true that a particular copula provides a more accurate description of equity returns than a normal copula, but there is no direct way to test the second part of the hypothesis that this copula is appropriate for describing the structure of stopping times. The current approach of specifying for each obligor some marginal distribution for the stopping time, specifying some form of copula to model the joint distribution of stopping times, and then using asset or equity returns to estimate required parameters for the copula is problematic.

One possible route would be to modify the Merton [1974] framework, and rather than assume a lognormal distribution to describe the value of the firm use a more general stochastic process that allows for fat tails and jumps in the value of the firm. Then we would introduce some form of barrier model so that default can occur at any point. Hence the marginal distribution of the stopping time is implicitly defined. The specification of the copula function would be dictated by the specification of the joint dependence of asset values.

Alternatively, one could specify some general process describing the stopping time for an obligor and some arbitrary copula function to describe the joint distribution of stopping times along the lines described by Schönbucher and Schubert [2001]. The challenge that this approach faces is justifying the underlying assumptions and calibration empirically.

Second, for the most part the choice of copula is arbitrary and seems to be dictated by ease of use. For practitioners, ease of use and calibration are important issues. However, for modeling default dependence even "simple" copulas become complicated once obligors start to default—see Schönbucher and Schubert [2001].

Fan [2003] argues that a mixed copula should be used for modeling equity return or foreign exchange dependence. Given the hypothesis, this implies that a mixed copula should be used for modeling default dependence, and this will impact the pricing of multiname credit instruments.

Third, in the standard application the marginal distribution for the stopping time is assumed to be exponential. In general this is inconsistent with the underlying structural model used to justify the use of equity correlations.

Finally, once an obligor defaults, the effect on the remaining obligors is dictated by the initial and arbitrary choice of the copula. This approach is independent of the state of the economy and independent of the industrial structure of the different sectors of the economy (see Jarrow and Yu [2001] and Yu [2003]).[36]

An unresolved issue with regard to the copula approach is that to be of practical use, any methodology describing the joint distribution of stopping times must be calibrated to current data. This places a severe restriction on the types of models that can be used. We require an internally consistent model that can be calibrated.

Reduced-Form Approach

Reduced-form models provide a natural way to describe the evolution of credit spreads. If the intensity function is modeled as a Cox process, this will in general imply that defaults are correlated. Typically, a set of state variables is specified, and for each obligor a time series analysis of credit spreads is used to estimate the coefficients. For many obligors, time series and cross-sectional analysis can be combined.

Given that for each obligor the default process depends on the same set of state variables, defaults are correlated—see Jarrow and van Deventer [2005]. Some academicians and practitioners have argued that reduced- form models are unable to generate acceptable levels of default correlations, but this view is challenged by Yu [2003].

Yu [2003] first considers the models developed by Duffee [1999] and Driessen [2005]. The intensity function is assumed to be identical under the natural pricing measures. This is a strong and questionable assumption, as Driessen presents empirical evidence that the default premium is greater than one. Yu adjusts for liquidity and taxation, and estimates the adjustment coefficients by calibrating to Moody's one-year transition matrix.[37]

The Driessen model generates substantially higher default correlations than the correlations produced by the Duffee model, which is not surprising, given the two additional economywide factors. The adjustment for liquidity and taxation is seen to have a major impact on the estimates of the default correlation inferred from credit spreads. For example, for a horizon of five years the unadjusted default correlation for A-rated firms is 0.93%, and the adjusted value is 11.78%. Default correlations increase monotonically with maturity. For maturities of over a year, default correlations are reduced as credit quality declines.

Yu also develops a number of different models that address how different industrial structures affect default correlations. Current applications of the copula methodology are unable to deal with this type of issue. Again the issue of calibration is not addressed.

Suppose a practitioner is trying to price a tranche of a collateralized debt obligation involving over a hundred names. To apply the reduced-form methodology, we need to estimate the intensity function for each obligor. It is an unfortunate reality that for many obligors we simply do not have enough bond or spread data to support the necessary econometric analysis. Until such data exist, or we have sufficient data from other markets such as the credit default swap market, the reduced-form methodology can be applied to only a small part of the fixed-income market. Using data from other markets also generates additional problems if these markets are not liquid, as there will be different liquidity premiums in the different markets (see Jarrow and Turnbull [2000] and Longstaff, Mithal, and Neis [2003]).

Testing a Model

To test an option pricing model, often some form of replicating portfolio is formed and the properties of the residual error, the difference between the price of the option and the replicating portfolio, are examined. This approach assumes we have an active secondary market for trading the option. For both single- and multiname credit derivatives, this is not the case, which prevents direct testing of pricing models. For investment-grade obligors, default is a rare event, so it is not possible to directly test a model of default dependence. Even for high-yield obligors we have limited data, which restricts testing of models.

How do we test a model of default dependence, given data limitations? This is one of the main complaints from regulators—our inability to test models that are used to generate risk measures for credit portfolios.

For testing copula models describing default dependence, the usual approach is to test the ability of the model to describe the joint distribution of equity returns. It is then assumed that the model is capable of describing the distribution of stopping times.

Clearly this is not a test of the model. We do have bond data, and this allows us to test models describing the evolution of credit spreads for different obligors.

In the reduced-form methodology, individual intensity functions are calibrated using a time series of credit spreads. The combining of time series and cross-sectional data and developing hypotheses about the joint co-movements of spreads remains to be done, but this approach does provide a way to test models, and the copula methodology does not.

In summary, we have no reliable empirical knowledge about default correlations. Given that default is a rare event, it is unlikely we can generate reliable statistical estimates. There do not seem to be any models we can use to model contagion in credit markets, although the frailty approach offers hope. There is no general economic framework that links the probability of default of single firms with a particular copula, apart from the original CreditMetrics framework. The reduced-form methodology must be adopted to avoid reliance on credit spreads.

Unresolved issues with regard to testing include:

1. Models must be able to a) describe contagion in credit markets, b) be calibrated, and c) be empirically tested.
2. The reduced-form methodology must be extended so that it can be more widely employed.

V. SUMMARY

We have explored a number of the important areas that directly affect risk management and the pricing of credit-risky instruments. There are many unresolved issues. The work of Acharya, Bharath, and Srinivasan [2003] and Altman, Resti, and Sironi [2005], who identify some of the factors that affect recovery rates, is a major step forward. It needs to be generalized to incorporate past performance of a firm, and out-of-sample testing is required. Similarly, we must generalize work on the modeling of the probability of default in a multiperiod setting to incorporate contagion arising from information effects if investors are imperfectly informed about common factors that affect the true probability of default.

Testing in these two areas of research—recovery rate and default probabilities—will always be inherently difficult, given the scarcity of data. A combination of these two streams of work will be an important contribution to the area of risk management, as it will facilitate the modeling of expected loss through the credit cycle, conditional on the current state of the economy.

We need to incorporate our empirical findings into pricing models for individual instruments and for modeling default dependence. At present default dependence is modeled using ad hoc copula functions with ad hoc parameter identification. More work is needed to relate the use of copulas to the specifics of firms and to the state of the economy. The reduced-form approach provides a natural means to model default dependence, but more work is required to justify the transition from the natural probability measure to the pricing measure and to identify the default risk premium. Once this is done, then by calibrating the model to actual default data and by changing to the pricing probability measure, this approach provides a natural way to model default dependence for pricing multiname products and addressing the pricing effects of counterparty risk.

Both structural and reduced-form approaches must address the issue that spreads may also be affected by other factors, apart from the risk of default.

ENDNOTES

The comments of Tobias Adrian, Robert Jarrow, Joshua Rosenberg, Til Schuermann, and seminar participants at the New York Federal Reserve Bank are gratefully acknowledged.

[1]In some markets dynamic hedging is rarely used, so the ability of the model to describe the evolution of spreads is of little relevance. It is necessary only to generate a term structure of spreads.

[2]The loss given default is calculated as (one minus the recovery rate).

[3]Recovery rates for bonds and loan are usually expressed as a fraction of the par value.

[4]An alternative approach to infer recovery rates from traded bond prices is described in Bakshi, Madan, and Zhang [2001], who find empirical support for recovery as a fraction of the discounted par value.

[5]See Acharya, Bharath, and Srinivasan [2003, Table 1].

[6]Recovery rate is the bond price within one month of default. If the coupon on the debt is high relative to the prevailing term structure, it is possible for the recovery rate to be greater than unity.

[7]The total amount of high-yield bonds is interpreted as a supply variable. Substituting the bond default amount produces similar results.

[8]The variables considered by category are: contract-specific characteristics: seniority, maturity, and current assets; firm- specific characteristics: profit margin, tangibility, logarithm of total assets, Q-ratio, firm return, firm volatility, and debt concentration; industry-specific characteristics: industry distress, median industry Q-ratio, median industry concentration, industry liquidity, and median industry leverage; and macroeconomic conditions: Fama-French factors, S&P 500 stock return, annual gross domestic product, product growth rate, aggregate weighted-average default rate, and total face value amount of defaulted debt.

[9]Large issues may earn higher recoveries than small issues, as large stakeholders in the bankruptcy proceedings may be able to exert greater bargaining power.

[10]Firms with more issues and a more dispersed creditor base, that is, lower debt concentration, may experience greater coordination problems and in turn lower recovery rates.

[11]Recovery rates for an industry in distress are expected to be reduced as a consequence. An industry is in distress if its median stock return in the year of the default is less than or equal to −30%. A second definition depends on median industry sales growth. If median industry growth is negative over a one- or two-year period prior to default, the industry is classified as in distress. The two definitions produce similar results.

[12]Kamakura uses the Janosi, Jarrow, and Yildirim [2002] model to generate multiperiod estimates of default probabilities. Moody's KMV model also generates a term structure of default probabilities.

[13]A logit normal transformation implies that recovery rates lie within the range zero to one. For the normal distribution, the recovery rate lies between minus infinity to plus infinity, and for the lognormal distribution between zero and plus infinity.

[14]This is a strong and unrealistic assumption that may affect the empirical results. It is employed because of the mathematical simplifications it produces and justified on the basis that it represents a good approximation for large international well-diversified banks.

[15]The usual argument is that a rise in the spot rate increases the drift term and reduces the probability of default (under the pricing measure) and hence reduces the spread. Longstaff and Schwartz [1995] and Duffee [1999] find a negative relationship. The argument used to justify a negative correlation is a partial result, as an increase in the spot rate may diminish the value of the firm and hence increase the probability of default.

[16]Litterman and Scheinkman [1991] find the spot rate and the slope of the yield curve important explanatory variables. An increase in the slope increases the expected spot rate and hence reduces the credit spread.

[17]Monthly S&P 500 returns are used (not sector returns).

[18]The spread is measured relative to the Treasury curve. The percentages increase when the Refinancing Corporation and swap curves are used as the riskless curve.

[19]The models include Longstaff and Schwartz [1995], Anderson, Sundaresan, and Tychon [1996], Leland and Toft [1996], Mella-Barral and Perraudin [1997], Collin-Dufresne and Goldstein [2001], and a jump diffusion model.

[20]It is assumed that $\lambda^{Q}_{j,t} = \mu\lambda^{P}_{j,t}$ for all t and j.

[21]An introduction to the reduced-form methodology is given in Duffie and Singleton [2003] and Schönbucher [2003a].

[22]See O'Kane and Turnbull [2003] for a detailed discussion.

[23]See Bakshi, Madan, and Zhang [2001], who infer estimates of the recovery rate under the measure Q.

[24]Mod-R is now the standard contract in the U.S. market.

It limits the maturity of deliverable obligations to the maximum of the remaining maturity of the swap contract and the minimum of 30 months or the longest remaining maturity of a restructured obligation. It applies only when the protection buyer has triggered the credit event.

Mod-Mod-R is the new European standard. It limits the maturity of deliverable obligations to the maximum of the remaining maturity of the CDS contract and 60 months for restructured obligations and 30 months for non-restructured obligations. It allows the delivery of conditionally transferable obligations rather than only fully transferable obligations. This should widen the range of loans that can be delivered. It too applies only when the protection buyer has triggered the credit event.

No-R clauses eliminate restructuring as a credit event, leaving just the hard (default) credit events. Some insurance companies and commercial banks favor this type of contract, and it remains to be seen if it will become liquid.

[25]For interest rate and foreign exchange swaps, see Duffie and Huang [1996] and Jarrow and Turnbull [1997].

[26]See Mashal and Naldi [2005] on the pricing of contracts in the presence of counterparty risk and Turnbull [2004] for the effects of counterparty risk on profit and loss.

[27]In the typical barrier model, default is said to occur the first time the value of the firm falls below some critical value called the barrier. In order to derive closed-form solutions, it is assumed that the barrier is either a constant or exponential—see Schönbucher [2003a, Chapter 9] for details.

[28]This contradicts the work of Zhou [2001] and Das et al. [2002].

[29]The estimated probabilities of joint default will be estimated with large standard errors. This will affect the estimation of the implied correlation. Estimation error would have a great effect when the joint default probabilities are low. Unfortunately S&R give no error statistics such as standard deviations, it is impossible to judge the significance of their results.

[30]The probability of default over the next period, given survival up to that point, is given by the expression $1 - \exp[-\lambda(t)]$.

[31]Frailty is a common term in the statistical analysis of biological and health data—see Hougaard [2000, Chapter 7]. In economics it is referred to as unobservable heterogeneity—see Kiefer [1988].

[32]In Merton diffusion models, equity and asset correlations are locally equal.

[33]Leland and Toft [1996] provide a theoretical justification for assuming a constant. In the case of an exponential $K_j(t)$, it is assumed that $K_j(t) = K_j\exp(\lambda_j t)$ where λ_j is a constant.

[34]In the Merton model the conclusion that asset and equity correlations are locally equal does not depend on leverage. See Schönbucher [2003a, Chapter 10].

[35]The result by Servigny and Renault [2002] must also be treated with caution. Work by Standard & Poor's provides estimates of joint probabilities and with the small number of observations, the standard errors have been high. The inversion procedure Servigny and Renault use is sensitive to measurement errors; they are measuring tail probabilities, and small differences may have a significant impact on the estimate of the implied correlation.

[36]For the bivariate case, Schönbucher [2003a, pp. 358-359] shows that for the normal and t-copula functions, given default by one obligor, the relative jump size in the intensity function of the remaining obligor has a singularity at the origin.

[37]The adjusted intensity function is assumed to be of the form:

$$\lambda_t^{Adj} = \lambda_t - \frac{a}{t + b}$$

where a and b are coefficients that must be estimated. It is not clear why Yu did not estimate the adjustment factor directly from spreads, given its deterministic nature.

REFERENCES

Acharya, V.V., S.T. Bharath, and A. Srinivasan. "Understanding the Recovery Rates on Defaulted Securities." Working paper, London Business School, 2003.

Altman, E.I. *Corporate Financial Distress and Bankruptcy.* New York: John Wiley & Sons, 1983.

——. "Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy." *Journal of Finance*, 23 (1968), pp. 589-609.

Altman, E.I., B. Brady, A. Resti, and A. Sironi. "The Link Between Default and Recovery Rates: Theory, Empirical Evidence and Implications." Working paper, Stern School of Business, New York University, 2003.

Altman, E.I., A. Resti, and A. Sironi. "Default Recovery Rates in Credit Modeling: A Review of the Literature and Empirical Evidence." *Journal of Finance Literature*, 1 (2005), pp. 21-45.

Anderson, R., S. Sundaresan, and P. Tychon. "Strategic Analysis of Contingent Claims." *European Economic Review*, 40 (1996), pp. 871-881.

Bakshi, G., D. Madan, and F. Zhang. "Understanding the Role of Recovery in Default Risk Models: Empirical Comparisons and Implied Recovery Rates." Working paper, School of Business, University of Maryland, 2001.

Bharath, S.T., and T. Shumway. "Forecasting Default with the KMV Merton Model." Working paper, University of Michigan, 2004.

Campbell, J.Y., J. Hilscher, and J. Szilagyi. "In Search of Distress Risk." Working paper, Department of Economics, Harvard University, 2004.

Chava, S., and R.A. Jarrow. "Bankruptcy Prediction with Industry Effects." Working paper, Cornell University, 2003.

Collin-Dufresne, P., and R.S. Goldstein. "Do Credit Spreads Reflect Stationary Leverage Ratios?" *Journal of Finance*, 56 (2001), pp. 1926-1957.

Collin-Dufresne, P., R.S. Goldstein, and J.S. Martin. "The Determinants of Credit Spread Changes." *Journal of Finance*, 56 (2001), pp. 2177-2207.

Das, S.R., L. Freed, G. Geng, and N. Kapadia. "Correlated Default Risk." Working paper, Leavey School of Business, Santa Clara University, 2002.

Das, S.R., and G. Geng. "Modeling the Processes of Correlated Default." Working paper, Leavey School of Business, Santa Clara University, 2002.

Davis, M., and V. Lo. "Infectious Defaults." *Quantitative Finance*, 1 (2001), pp. 382-387.

——. "Modeling Default Correlation in Bond Portfolios." In E. Alexander, ed., *Mastering Risk,* Vol. 2. Englewood Cliffs, NJ: Prentice Hall, 2001.

Driessen, J. "Is Default Event Risk Priced in Corporate Bonds?" *The Review of Financial Studies*, 18 (2005), pp. 165-195.

Duffee, G. "Estimating the Price of Default Risk." *The Review of Financial Studies*, 12 (1999), pp. 197-226.

Duffie, D., and M. Huang. "Swap Rates and Credit Quality." *Journal of Finance*, 51 (1996), pp. 921-950.

Duffie, D., and D. Lando. "Term Structures of Credit Spreads with Incomplete Accounting Information." *Econometrica*, 69 (2001), pp. 633-664.

Duffie, D., and K. Singleton. *Credit Risk.* Princeton: Princeton University Press, 2003.

Duffie, D., and K. Wang. "Multi-Period Corporate Failure Prediction with Stochastic Covariates." Working paper, Stanford University, 2003.

Dullmann, K., and M. Trapp. "Systematic Risk in Recovery Rates – An Empirical Analysis of U.S. Corporate Credit Exposure." Working paper, Deutsche Bundesbank, Frankfurt, Germany, 2004.

Elton, E.J., M.J. Gruber, D. Agrawal, and C. Mann. "Explaining the Rate Spread on Corporate Bonds." *Journal of Finance*, 56 (2001), pp. 247-277.

Eom, Y.H., J. Helwege, and J. Huang. "Structural Models of Corporate Bond Pricing: An Empirical Analysis." *The Review of Financial Studies*, 17 (2004), pp. 499-544.

Fan, W. "Modeling Asymmetric Dependence with Copulas." Working paper, Department of Economics, Yale University, 2003.

Finger, C. "Conditional Approaches for CreditMetrics Portfolio Distributions." *CreditMetrics Monitor*, April 1999.

Frye, J. "Depressing Recoveries." *Risk*, 13(11) (2000), pp. 108-111.

——. "A False Sense of Security." *Risk*, 16(8) (2003), pp. 63-67.

Gagliardini, P., and C. Gourieroux. "Spread Term Structure and Default Correlation." Working paper, Department of Economics, University of Toronto, 2003.

Geske, R. "The Valuation of Corporate Liabilities as Compound Options." *Journal of Financial and Quantitative Analysis,* 12 (1977), pp. 541-552.

Gordy, M. "A Comparative Anatomy of Credit Risk Models." *Journal of Banking and Finance*, 24 (2000), pp. 119-149.

Gupton, G. M., D. Gates, and L. Carty. "Bank Loan Loss Given Default." Moody's Special Comment, 2000.

Guha, R. "Recovery of Face Value at Default: Theory and Empirical Evidence." Working paper, London Business School, 2002.

Hillegeist, S.A., E.K. Keating, and D.P. Cram. "Assessing the Probability of Default." *Review of Accounting Studies*, 9 (2004), pp. 5-34.

Hougaard, P. *Statistics for Biology and Health.* New York: Springer-Verlag, 2000.

Huang, J., and M. Huang. "How Much of the Corporate Treasury Spread Is Due to Credit Risk?" Working paper, The Pennsylvania State University, 2003.

Jaffee, D.M. "Cyclical Variations in the Risk Structure of Interest Rates." *Journal of Monetary Economics*, 1 (1975), pp. 309-325.

Janosi, T., R. Jarrow, and Y. Yildirim. "Estimating Expected Losses and Liquidity Discounts in Debt Prices." *Journal of Risk*, 5 (2002).

Jarrow, R.A., and D. van Deventer. "Estimating Default Correlations Using a Reduced Form Model." *Risk*, 18 (January 2005), pp. 83-87.

Jarrow, R.A., and S.M. Turnbull. "Drawing the Analogy." *Risk*, 5 (October 1992), pp. 63-70.

——. "The Intersection of Market and Credit Risk." *Journal of Banking and Finance*, 24 (2000), pp. 271-299.

——. "Pricing Derivatives on Financial Securities Subject to Credit Risk." *Journal of Finance*, 50 (1995), pp. 53-85.

——. "When Swaps Are Dropped." *Risk*, 10 (May 1997), pp. 70-75.

Jarrow, R.A., and F. Yu. "Counterparty Risk and the Pricing of Defaultable Securities." *Journal of Finance*, 56 (2001), pp. 1765-1800.

Kiefer, N.M. "Economic Duration Data and Hazard Functions." *Journal of Economic Literature*, 26 (1988), pp. 646-679.

Leland, H.E., and K. Toft. "Optimal Capital Structure, Endogenous Bankruptcy, and the Term Structure of Credit Spreads." *Journal of Finance*, 51 (1996), pp. 987-1019.

Li, D.X. "On Default Correlation: A Copula Function Approach." *The Journal of Fixed Income*, 9 (1) (2000), pp. 43-54.

Litterman, R., and J. Scheinkman. "Common Factors Affecting Bond Returns." *The Journal of Fixed Income*, 1 (1991), pp. 54-61.

Longstaff, F. A., S. Mithal, and E. Neis. "The Credit Default Swap Market." Working paper, Anderson School, UCLA, 2003.

Longstaff, F.A., and E.S. Schwartz. "A Simple Approach to Valuing Risky and Floating Rate Debt." *Journal of Finance*, 50 (1995), pp. 789-819.

Mashal, R., and M. Naldi. "Extreme Events and the Valuation of Default Baskets." *Quantitative Credit Research Quarterly*, October 2001.

——. "Pricing Portfolio Default Swaps with Counterparty Risk." *The Journal of Fixed Income*, 14(4) (2005), pp. 5-16.

Mashal, R., M. Naldi, and A. Zeevi. "Extreme Events and Multi-Name Credit Derivatives." *Risk*, October 2003, pp. 43-54.

Mella-Barral, P., and W. Perraudin. "Strategic Debt Service." *Journal of Finance*, 52 (1997), pp. 531-556.

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

Moody's Investors Service. "Default Recovery Rates of Corporate Bond Issuers." February 2003.

Ohlson, J. "Financial Ratios and the Probabilistic Prediction of Bankruptcy." *Journal of Accounting Research*, 19 (1980), pp. 109-131.

O'Kane, D., C. Pedersen, and S. M. Turnbull. "The Restructuring Clause in Credit Default Swap Contracts." Lehman Brothers, 2003.

O'Kane, D., and L. Schloegl. "Modeling Credit: Theory and Practice." Lehman Brothers, 2001.

O'Kane, D., and S. M. Turnbull. "Valuation of Credit Default Swaps." Lehman Brothers, 2003.

"Overview of the New Basle Capital Accord." Basel Committee on Banking Supervision, 2002. Available at: http://www.bis.org/publ/bcbsca02.pdf.

Pykhtin, M. "Unexpected Recovery Risk." *Risk*, 16(8) (2003), pp. 74-78.

Schönbucher, P.J. *Credit Derivatives Pricing Models*. London: John Wiley, 2003a.

——. "Information Driven Default Contagion." Working paper, Department of Mathematics, ETH, Zurich, 2003b.

Schönbucher, P.J., and D. Schubert. "Copula Dependent Default Risk in Intensity Models." Working paper, Department of Statistics, University of Bonn, 2001.

Schuermann, T. "What Do We Know About Loss Given Default?" In *Credit Risk Models and Management*. London: Risk Books, 2004.

Servigny, A., and O. Renault. "Default Correlation: Correlation Evidence." *Risk*, 16(7) (2002), pp. 90-94.

Shumway, T. "Forecasting Bankruptcy More Accurately: A Simple Hazard Model." *Journal of Business*, 74(1) (2001), pp. 101-124.

Turnbull, S.M. "The Pricing Implications of Counterparty Risk for Non Linear Products." Working paper, Bauer College of Business, University of Houston, October 2004.

——. "Pricing Loans Using Default Probabilities." *Economic Notes*, 32(2) (2003), pp. 197-217.

Wilson, T. "Portfolio Credit Risk." *Risk*, 10(9) (1997), pp. 111-116.

Yang, J. "An Empirical Investigation of the Effects of Macroeconomic Variables on Default Risk." Working paper, University of Toronto, 2003.

Yu, F. "Accounting Transparency and the Term Structure of Credit Spreads." *Journal of Financial Economics*, 75(1) (2005), pp. 53-84.

——. "Correlated Defaults in Reduced Form Models." Working paper, University of California, Irvine, 2003.

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

To order reprints of this article, please contact Ajani Malik at amalik@iijournals.com or 212-224-3205.

©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 Instituitonal Investor PLC.
The most recent two editions of this title are only ever available at http://www.euromoneyplc.com