8

The Credit Risk of Portfolios and Correlations

INTRODUCTION

So far, we have considered default risk and credit risk exposure on a single-borrower basis. Indeed, much of the banking theory literature views the personnel at banks and similar financial institutions (FIs) as credit specialists who, through monitoring and the development of long-term relationships with customers, gain a comparative advantage in lending to a specific borrower or group of borrowers.[1]

However, investment principles dictate that diversification reduces risk. For example, investing in a single stock will expose the investor to both market (systematic) and company-specific (unsystematic) risk, but adding other stocks into a portfolio will tend to diversify away the unsystematic component of risk, thereby reducing the investor's risk exposure. Because of this fundamental principle of modern portfolio theory (MPT), required returns do not include a premium for unsystematic risk.

The same principle arises when investing in debt instruments that are exposed to credit risk. If one borrower's risk of default is inversely related to another borrower's default probability, then combining loans to both borrowers may reduce the investor's (lender's) overall credit risk exposure. That is, if there is negative correlation across borrower default probabilities, then a portfolio of loans may have lower risk than an individual loan, all else equal. In this chapter, we discuss the issue of portfolio diversification in the general context of MPT and then examine the estimation of correlations used in assessing a portfolio's credit risk exposure.[2]

MODERN PORTFOLIO THEORY (MPT): AN OVERVIEW

Modern portfolio theory is used to derive optimal portfolios in a mean-variance framework. That is, investors attempt to maximize expected returns (mean) and minimize risk (variance). The (mean) return and risk of a portfolio of assets, under the assumption that returns on individual assets are normally distributed (or that asset managers have a quadratic utility function), are given in equations (8.1), (8.2), and (8.3). The assumption that individual asset returns are normally distributed and/or that managers of a financial intermediary exhibit a particular set of preferences (quadratic utility) toward returns implies that only two moments of the distribution of assets returns are necessary in order to analyze portfolio decisions: (1) the mean return of a portfolio and (2) its variance (or the standard deviation of the returns on that portfolio). Since MPT is forward-looking, the expected return and risk measures are by definition unobservable. As a result, portfolio returns and risks are usually estimated from historical time series of the returns and risks on individual assets.

Given these assumptions, the mean return ($\bar{R}_p$) and the variance of returns (σ_p^2) on a portfolio of n assets can be computed as:

$$\bar{R}_p = \sum_{i=1}^{n} X_i \bar{R}_i \tag{8.1}$$

$$\sigma_p^2 = \sum_{i=1}^{n} X_i^2 \sigma_i^2 + \sum_{i=1}^{n} \sum_{\substack{j=1 \\ i \neq j}}^{n} X_i X_j \sigma_{ij} \tag{8.2}$$

or

$$\sigma_p^2 = \sum_{i=1}^{n} X_i^2 \sigma_i^2 + \sum_{i=1}^{n} \sum_{\substack{j=1 \\ i \neq j}}^{n} X_i X_j \rho_{ij} \sigma_i \sigma_j \tag{8.3}$$

where $\bar{R}_p$ = the mean return on the asset portfolio
$\bar{R}_i$ = the mean return on the ith asset in the portfolio
X_i = the proportion (weight) of the asset portfolio invested in the ith asset with $i = 1, \ldots, n$
σ_i^2 = the variance of the returns on the ith asset
σ_{ij} = the covariance of the returns between the ith and jth assets, with $j = 1, 2, \ldots, n$
ρ_{ij} = the correlation between the returns on the ith and jth assets, where $-1 \leq \rho_{ij} \leq +1$

From equation (8.1), it can be seen that the mean return on a portfolio of assets ($\bar{R}_p$) is simply a weighted average (with weights X_i) of the mean returns on the individual assets in that portfolio ($\bar{R}_i$). By comparison, the variance of returns on a portfolio of assets (σ_p^2) is decomposable into two terms. The first term reflects the weighted sum of the variances of returns on the individual assets (σ_i^2), and the second term reflects the weighted sums of the covariances among the assets (σ_{ij}). Because a covariance is unbounded, it is common in MPT-type models to substitute the correlation among asset returns for the covariance term, using the statistical definition:

$$\sigma_{ij} = \rho_{ij}\sigma_i\sigma_j \tag{8.4}$$

Because a correlation is constrained to lie between plus and minus one, we can evaluate the effect of varying ρ_{ij} on asset portfolio risk. For example, in the two-asset case, if ρ_{ij} is negative, the second term in equation (8.3) will also be negative and will offset the first term, which will always be positive.[3] By appropriately exploiting correlation relationships among assets, a portfolio manager can significantly reduce risk and improve a portfolio's risk-return trade-off.[4] Computationally, the efficient frontier, or the portfolio of assets with the lowest risk for any given level of return, can be determine by solving for the asset proportions (X_i) that minimize σ_p for each given level of returns ($\bar{R}_p$). In Figure 8.1, both B and C are efficient asset portfolios in this sense.

The best of all the risky asset portfolios on the efficient frontier is the one that exhibits the highest excess return over the risk-free rate (r_f) relative to the level of portfolio risk, or the highest risk-adjusted excess return:[5]

$$(\bar{R}_p - r_f)/\sigma_p \tag{8.5}$$

This risk-return ratio is usually called the *Sharpe ratio*. Graphically, the optimal risky asset portfolio is the one in which a line drawn from the return axis, with an origin at r_f, is just tangential to the efficient frontier (this is shown as portfolio D in Figure 8.1). Because the slope of this line reflects the $(\bar{R}_p - r_f)/\sigma_p$ ratio for that portfolio, it is also the portfolio with the highest Sharpe ratio.[6]

APPLYING MPT TO NONTRADED BONDS AND LOANS

MPT has been around for over 40 years and is now a portfolio management tool commonly used by most mutual fund and pension fund managers. It has also been applied with some success to publicly traded junk bonds

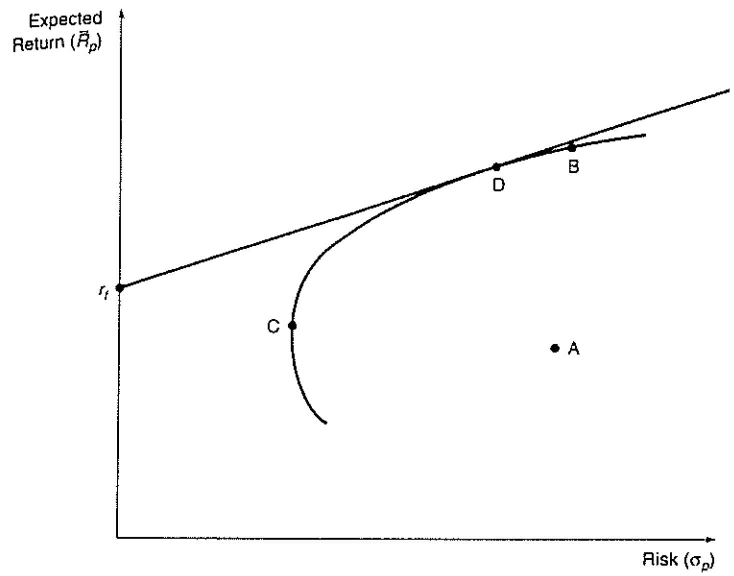

FIGURE 8.1 The Optimum Risky Loan Portfolio

when their returns have tended to be more equity-like than bond-like and when historical returns are available (see Altman and Saunders [1997]). With respect to most loans and bonds, however, there are problems with non-normal returns, unobservable returns, and unobservable correlations.[7]

Non-Normal Returns

Loans and bonds tend to have relatively fixed upside returns and long-tailed downside risks. Thus, returns on these assets tend to exhibit a strong negative skew and, in some cases, kurtosis (fat-tailedness) as well. MPT is built around a model in which only two moments—the mean and variance—are required to describe the whole distribution of returns. To the extent that the third (skewness) and fourth (kurtosis) moments of returns are important in fully describing the distribution of asset returns, the use of simple, two-moment MPT models becomes difficult to justify.[8]

Unobservable Returns and Correlations

An additional problem in the application of portfolio theory to credit risk measurement relates to the fact that most loans and corporate bonds are nontraded or are traded over-the-counter at very uneven intervals with little historical price or volume data. This makes it difficult to compute mean returns ($\bar{R}_i$), the variance of returns (σ_i^2), and the covariance (σ_{ij}) or return correlations (ρ_{ij}) using historical time series. It is to this issue that we now turn.

ESTIMATING CORRELATIONS ACROSS NONTRADED ASSETS

Implementing the MPT was impractical when it was first introduced because of the large data requirements associated with estimating all possible pair-wise correlation coefficients in real time. For example, in 2008, Kamakura provided estimates of 4.85 million pairs of correlation coefficients covering all possible pairs of 21,000 public firms across 30 countries. This database management problem was exacerbated for nontraded assets, such as loans, that are of critical importance in credit risk assessment. In addition to the database management problem associated with large numbers of pair-wise correlations, the use of historical correlations is not reliable because of the instability associated with changing debt positions and credit quality of borrowing firms.

Partly as a response to the information costs associated with the data requirements of the MPT, the capital asset pricing model (CAPM) was introduced by William Sharpe, John Lintner, and Jan Mossin.[9] Together with the work of Harry Markowitz and Merton Miller, Sharpe's work on the CAPM won the Nobel Prize in economics in 1990. Among other contributions, the CAPM transferred the correlation analysis from a data-intensive pair-wise computation to a single regression analysis in which the correlation can be computed from the regression coefficient on an overall market risk index. That is, if all assets are correlated to the systematic market risk factor, then the correlation coefficient can be expressed in terms of the asset's beta, which is obtained from the following regression analysis:

$$R_i = r_f + \beta_i R_M \tag{8.6}$$

where R_i is asset i's expected return, r_f is the risk-free rate, β_i is the estimated value of asset i's beta coefficient that measures systematic risk exposure, and R_M is the return on a market index (such as the S&P 500). Since unsystematic (company-specific) risk is diversifiable, the only risk that is priced according to the CAPM is systematic market risk. Each asset's market beta can then be used to calculate the correlation coefficient of the return on the asset ρ_{iM} as follows:

$$\rho_{iM} = \beta_i \sigma_M / \sigma_i \tag{8.7}$$

Thus, for example, if asset A has a beta of 0.5 and asset B has a beta of −0.5, and both assets have the same standard deviation of returns, then assets A and B are perfectly negatively correlated from the standpoint of their systematic risk exposures. That is, when asset A has a positive return, asset B is expected to have a negative return of equal absolute value. If asset C has a beta of 0.5 and the same standard deviation of returns, then assets A (B) and C are perfectly positively (negatively) correlated. Thus, rather than calculating correlation coefficients for every possible pair of assets, the CAPM permits the computation of correlations using a single beta for each asset.

The commercial credit risk estimation products (e.g., Moody's KMV and the Kamakura Risk Manager) calculate correlation coefficients using a process based on an asset pricing model, as expressed in simple form in equation (8.7).[10] Of course, rather than using only a single market index (such as the S&P 500), the methodologies used by the commercial products incorporate more complex asset pricing models. For example, Chen et al. (2008) outline a reduced form model that estimates a two-factor model for interest rate risk and a one-factor model for default risk. They then use the estimated functional forms to solve for the correlation across the factors.[11] As an illustration of how commercial models utilize this methodology, we now describe the portfolio models of Moody's KMV and Kamakura.[12]

MOODY'S KMV'S PORTFOLIO MANAGER

KMV's Portfolio Manager can be viewed as a full-fledged MPT optimization approach because all three key variables—returns, risks, and correlations—are calculated. However, it can also be used to analyze risk effects alone, as will be discussed below. This section explains how the three key variables that enter into any MPT model can be calculated.

Returns

In the absence of historical returns on traded loans, the (expected) excess return over the risk-free rate on the ith loan ($R_i - r_f$) over any given horizon can be set equal to:

$$R_i - r_f = [\text{Spread}_i + \text{Fees}_i] - [\text{Expected loss}_i] - r_f \tag{8.8}$$

or

$$R_i - r_f = [\text{Spread}_i + \text{Fees}_i] - [\text{EDF}_i \times \text{LGD}_i] - r_f \tag{8.8'}$$

The first component of returns is the spread of the loan rate over a benchmark rate such as the London Inter-Bank Offered Rate (LIBOR), plus any fees directly earned from the loan and expected over a given period (say, a year). Expected losses on the loan are then deducted because they can be viewed as part of the normal cost of doing banking business. In the context of a KMV-type model, where the expected default frequency (EDF) is calculated from stock returns (as in the Credit Monitor model), then, for any given borrower, expected losses will equal EDF_i times LGD_i, where LGD_i is the loss given default for the ith borrower (usually estimated from the bank's internal database). KMV deducts the risk-free rate, r_f, to present loan returns in an "excess return" format. Of course, if the bank desires, it can calculate the portfolio model using gross returns instead (i.e., not deducting r_f).[13]

Loan Risks

Again assume that the loan matures on or before the chosen credit risk horizon date. In the absence of return data on loans, a loan's risk (σ_i) can be approximated by the unexpected loss rate on the loan (UL_i)—essentially, the variability of the loss rate around its expected loss value ($EDF_i \times LGD_i$). There are a number of ways in which UL_i might be calculated, depending on the assumptions made about the maturity of the loan relative to the credit horizon, the variability of LGD, and the correlation of loan LGDs with EDFs. For example, in the simplest form, when a loan matures before the horizon, a default-only model (DM) can be employed where the borrower either defaults or doesn't default (i.e., there are no credit migrations as in a mark-to-market (MTM) model—see the discussion in Chapter 9), so that defaults are binomially distributed with a fixed LGD across all borrowers. Under these conditions, UL_i can be estimated as:

$$\sigma_i = UL_i = LGD \times \sqrt{(EDF_i)(1 - EDF_i)} \tag{8.9}$$

where $\sqrt{(EDF_i)(1 - EDF_i)}$ reflects the variability of a default rate frequency that is binomially distributed.[14]

A slightly more sophisticated DM version would allow LGD to be variable, but factors affecting EDFs are assumed to be different from those affecting LGDs, and LGDs are assumed to be independent across borrowers.[15] In this case (see Kealhofer [1995]):

$$\sigma_i = \sqrt{EDF_i(1 - EDF_i)\overline{LGD}_i^2 + EDF_i VOL_i^2} \tag{8.10}$$

where $\overline{LGD}_i$ is the expected value of borrower i's LGD, and VOL_i is the standard deviation (volatility) of borrower i's LGD.

Equation (8.10) can be generalized to solve for σ_i under a full mark-to-market (MTM) model with credit upgrades and downgrades as well as default. That is, for the case where the maturity of the loan exceeds the loan's credit horizon, the loan's risk is measured as:

$$\sigma_i = \sqrt{EDF_i(1 - EDF_i)\overline{LGD}_i^2 + EDF_i(VVOL_i)^2 + (1 - EDF_i)VVOL_i^2} \tag{8.10'}$$

where $VVOL_i$ (or valuation volatility) is the standard deviation of borrower i's MTM loan value in the nondefault state.

$VVOL_i$ can be viewed as the standard deviation of asset values and can be calculated using the methodology outlined in Chapter 4. However, in Chapter 4 we focused on the area under the valuation distribution that fell below the default point (i.e., the region in which the value of assets fell below the debt repayment). Here we examine only the distribution of asset values above the default point in order to estimate the VVOL.[16]

Another difference between Moody's KMV's Portfolio Manager (PM) and the discussion of Moody's KMV in Chapter 4 is that PM does not assume normally distributed asset portfolios. Both an analytical approximation and the Monte Carlo method are used in the MTM version of PM so as to allow for the possibility of fat tails in the distribution of portfolio returns. The analytical approximation adjusts tail probabilities based on returns, the weighted average of individual loan ULs, and minimum and maximum possible portfolio values. The analytical approximation is most accurate for the 10 basis point level of tail risk (i.e., the worst one-thousandth of all possible outcomes). Monte Carlo simulation draws states of the world to estimate whether each borrower in the portfolio defaults and, if so, what the LGD would be, conditional on the random draw of overall business factors.[17] This process is repeated 50,000 to 200,000 times to determine a frequency distribution that approximates the distribution of the portfolio's value.[18]

Correlations

One important intuition from the structural form approach is that default correlations are generally likely to be low. To see why, consider the context of the two-state DM version of a KMV-type model. A default correlation would reflect the joint probability of two firms, G and F—say, for example, General Electric and Ford—having their asset values fall below their debt values over the same horizon (say one year). In the context of Figure 8.2, the General Electric asset value would have to fall below its debt value (B_G in the figure), and the Ford asset value would have to fall below its debt value (B_F). The joint area of default is shaded, and the joint probability distribution of asset values is represented by the concentric circles. The circles are similar to those used in geography maps to describe the topographical characteristics (e.g., height) of hills. The inner circle is the top of the hill (high probability), and the outer circles are the bottom of the hill (low probability). The joint probability that asset values will fall in the shaded region is low (as shown) and will depend, in part, on the asset correlations between the two borrowers.[19] The two graphs below and to the left of the graph of the concentric circles represent the payoff on each firm's debt as a function of the market value of the firm's assets. Applying equation (8.4) to the

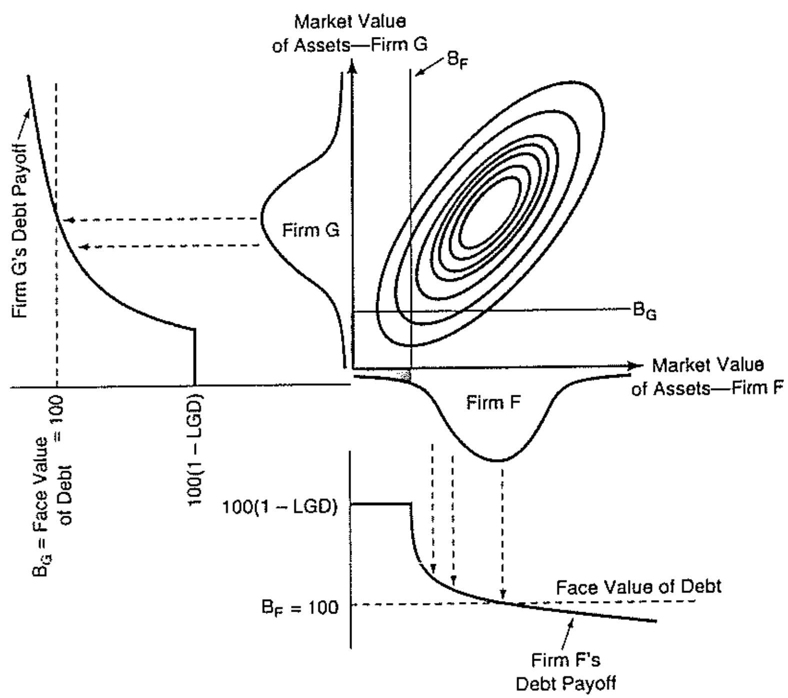

FIGURE 8.2 Value Correlation

simple binomial DM model for Ford (F) and General Electric (G) yields:

$$\rho_{GF} = \frac{\sigma_{GF}}{\sigma_F \times \sigma_G} \tag{8.4'}$$

or

$$\rho_{GF} = \frac{JDF_{GF} - (EDF_G \times EDF_F)}{\sqrt{(EDF_G)(1 - EDF_G)}\sqrt{(EDF_F)(1 - EDF_F)}} \tag{8.11}$$

The numerator of equation (8.11) is the covariance (σ_{GF}) between the asset values of the two firms, G and F. It reflects the difference between the cases where the two asset values are jointly distributed (JDF_{GF}) and where they are independent ($EDF_G \times EDF_F$).[20] The denominator reflects the standard deviation (σ) of default rates under the binomial distribution for each firm.

Although correlations may generally be low, Figure 8.2 can be used to understand the dynamics of how correlations may increase over time. For example, KMV correlations among U.S. firms have recently been rising. To see why, note that the leverage ratios of U.S. corporations have more than doubled over the past decade (i.e., in the context of Figure 8.2, B_F and B_G have both shifted up along their respective axes) and thus the cross-hatched area of joint default has expanded.[21]

Rather than seeking to directly estimate correlations using equation (8.11), Moody's KMV uses a multifactor stock-return model from which correlations are derived. The model reflects the correlation among the systematic risk factors affecting each firm and their appropriate weights. Moody's KMV's multifactor approach to calculating correlations is somewhat similar to the CreditMetrics stock-return factor approach to correlation calculation discussed more fully later in this chapter, except that KMV uses asset correlations rather than equity correlations.[22] Moody's KMV typically finds that correlations lie in the range .002 to .15. Gupton (1997) employs Moody's data over 1970–1995 to obtain implied default correlations between .0013 to .033 using CreditMetrics.[23] The low correlations obtained using all of these models are consistent with evidence showing a significant reduction in credit risk for diversified debt portfolios. Moody's KMV shows that 54 percent of the risk can be diversified away by simply choosing a portfolio composed of the debt issued by five different BBB-rated firms.[24] Barnhill and Maxwell (2001) show that diversification can reduce a bond portfolio's standard deviation from $23,433 to $8,102 ($9,518) if the portfolio consists of 100 bonds from 24 industry sectors (a single sector). Carey (1998) also finds significant diversification benefits across size, obligor concentration, and rating classification for a portfolio consisting of private placements.[25] However, some correlations can be quite high. In November 2008, Kamakura reported a 70 percent default correlation between Citigroup and Ford, as well as an 88 percent correlation coefficient between Ford and GM.

Calculating Correlations Using Moody's KMV's Portfolio Manager To estimate correlations, Moody's KMV's Portfolio Manager decomposes asset returns into systematic and unsystematic risk using a three-level structural model. Asset returns are extracted from equity returns using the Moody's KMV Credit Manager approach outlined in Chapter 4 for imputing firm asset values. Using a time series of such asset values, asset returns can be calculated. Once asset returns are estimated, the first-level decomposition into risk factors is a single index model that regresses asset returns on a composite market factor that is constructed individually for each firm. The composite market factor used in the first-level analysis comprises a weighted sum of country and industry factors. These factors are estimated at the second level of analysis and may be correlated with each other.[26]

The second level separates out the systematic component of industry and country risk, each of which is further decomposed into three sets of independent factors at the third level. These third-level factors are: (1) two global economic factors—a market-weighted index of returns for all firms and the return index weighted by the log of market values; (2) five regional factors—Europe, North America, Japan, Southeast Asia, and Australia/New Zealand; (3) seven sector factors—interest sensitive (banks, real estate, and utilities), extraction (oil and gas, mining), consumer nondurables, consumer durables, technology, medical services, and other (materials processing, chemicals, paper, steel production).[27]

For any firm i, the multifactor model can be written as:

$$R_k = \sum_{G=1,2} \beta_{kG} R_G + \sum_{R=1,\ldots,5} \beta_{kR} R_R + \sum_{S=1,\ldots,7} \beta_{kS} R_S + \sum_{I} \beta_{kI} \varepsilon_I + \sum_{C} \beta_{kC} \varepsilon_C + \varepsilon_k \tag{8.12}$$

where $\beta_{kG}, \beta_{kR}, \beta_{kS}$ = firm k's beta coefficients on global, regional, and sector factors (from the third regression level)

R_G = the return on the two independent global economic factors

R_R = the return on the five independent regional economic factors

R_S = the return on the seven independent industrial sector effects

β_{kI}, β_{kC} = firm k's beta coefficients on the country- and industry-specific systematic risk components (from the second level)

ε_I = the industry-specific effect for industry I

ε_C = the country-specific effect for country C

ε_k = firm k's company-specific risk (from the first level)

We can express the asset variance for firm k as follows:

$$\sigma_k^2 = \sum_{G=1,2} \beta_{kG}^2 \sigma_G^2 + \sum_{R=1,\ldots,5} \beta_{kR}^2 \sigma_R^2 + \sum_{S=1,\ldots,7} \beta_{kS}^2 \sigma_S^2 + \sum_{I} \beta_{kI}^2 \sigma_I^2 + \sum_{C} \beta_{kC}^2 \sigma_C + \varepsilon_k^2 \tag{8.13}$$

Equation (8.13) can be used to calculate correlations between firms j and k as follows:

$$\sigma_{jk} = \sum_{G=1,2} \beta_{jG}\beta_{kG}\sigma_G^2 + \sum_{R=1,\ldots,5} \beta_{jR}\beta_{kR}\sigma_R^2 + \sum_{S=1,\ldots,7} \beta_{jS}\beta_{kS}\sigma_S^2 + \sum_{I} \beta_{jI}\beta_{kI}\sigma_I^2 + \sum_{C} \beta_{jC}\beta_{kC}\sigma_C^2 \tag{8.14}$$

Thus, the correlation coefficient between firms j and k is:

$$\rho_{jk} = \sigma_{jk}/\sigma_j\sigma_k$$

After they are calculated, the three inputs (returns, risks, and correlations) can be employed in a number of directions. One potential use would be to calculate a risk/return efficient frontier for the loan portfolio, as shown in Figure 8.1. Reportedly, one large Canadian bank manages its U.S. loan portfolio using a Moody's KMV-type model.[28]

A second use would be to measure the risk contribution of expanding lending to any given borrower. As discussed earlier in this chapter, the risk (in a portfolio sense) of any one loan will depend not only on the risk of the individual loan on a stand-alone basis, but also on its correlation with the risks of other loans. For example, a loan, when viewed individually, might be thought to be risky, but because its returns are negatively correlated with other loans, it may be quite valuable in a portfolio context in lowering portfolio risk. The measurement of the marginal contribution to the risk of a portfolio of any particular loan is called *loan transfer pricing*.

The effects of making additional loans to a particular borrower also depend crucially on assumptions made about the balance sheet constraint. For example, if investable or loanable funds are viewed as fixed, then expanding the proportion of assets lent to any borrower i (i.e., increasing the asset i's portfolio weight, X_i) means reducing the proportion invested in all other loans (assets). However, if the funds constraint is viewed as being nonbinding, then the amount lent to borrower i can be expanded without affecting the amount lent to other borrowers. In the KMV-type marginal risk contribution calculation, the funding constraint is assumed to be binding if:

$$X_i + X_j + \cdots + X_n = 1$$

By comparison, under CreditMetrics (see Chapter 9), marginal risk contributions are calculated assuming no such funding constraint; for example, a bank can make a loan to a twentieth borrower without reducing the loans outstanding to the 19 other borrowers.

Assuming a binding funding constraint, the marginal risk contribution for the ith loan (MRC_i) can be calculated as:[29]

$$MRC_i = X_i \frac{dUL_p}{dX_i} \qquad (8.15)$$

where UL_p is the risk (standard deviation) of the total loan portfolio and X_i is the proportion of the loan portfolio lent to the ith borrower:[30]

$$UL_p = \sqrt{\sum_{i=1}^{N} X_i^2 UL_i^2 + \sum_{i=1}^{N} \sum_{\substack{j=1 \\ i \neq j}}^{N} X_i X_j UL_i UL_j \rho_{ij}} \qquad (8.16)$$

and

$$\sum_{i=1}^{N} X_i = 1$$

The marginal risk contribution can be viewed as a measure of the economic capital needed by the bank in order to make a new loan to the ith borrower because it reflects the sensitivity of portfolio risk (specifically, portfolio standard deviation) to a small percentage change in the weight of the asset (dX_i). Note that the sum of MRCs is equal to UL_p; consequently, the required capital for each loan is just its MRC scaled by the capital multiple (the ratio of capital to UL_p).[31]

KAMAKURA AND OTHER REDUCED FORM MODELS

It is also important to discuss default correlations derived from intensity-based models, such as the Kamakura Risk Manager (see Chapter 5). In these models, default correlations reflect the effect of events inducing *simultaneous* jumps in the default intensities of obligors. The causes of defaults themselves are not modeled explicitly; instead, the focus is on modeling various approaches to default-arrival intensity based on correlated *times to default*. This allows the model to answer questions such as what was the worst week, month, year, and so on, out of the past N years, in terms of loan portfolio risk? That worst period will be when correlated default intensities were the highest (defaults arrived at the same time). With joint credit events, some of the default intensity of each obligor is tied to such a marketwide event with some probability. For example, the intensity-based model of Duffie and Singleton (1998) allows for default intensities to be correlated through changes in default intensities themselves as well as joint credit events. In the Duffie and Singleton model, obligors have default intensities that mean-revert with correlated Poisson arrivals of randomly sized jumps. They then formulate individual obligor default intensity times as multivariate exponentials, which allows them to develop a model for simulating correlated defaults.

Duffie and Singleton (1998) consider a hazard function in which each asset's conditional default probability is a function of four parameters: λ, θ, k, and J.[32] That is, the intensity h of a loan's default process has independently distributed jumps in default probability that arrive at some constant intensity λ; otherwise, if no default event occurs, h returns at mean-reversion rate k to a constant default intensity θ. The jumps in intensity follow an exponential distribution with mean size of jump equal to J. Therefore, the form of the individual firm's probability of survival (conditional upon survival to date t) from time t to time s is:

$$p(t,s) = e^{\alpha(s-t)+\beta(s-t)h(t)}$$

where $\beta(t) = -(1 - e^{-kt})/k$

$\alpha(t) = -\theta[t + \beta(t)] - [\lambda/(J + k)][Jt - \ln(1 - \beta(t)J)]$

As a numerical illustration, suppose that $\lambda = .001$, $k = .5$, $\theta = .001$, $J = 5$, and $h(0) = .001$.[33] Then the arrival of a jump in default risk reduces the expected remaining life of the loan to less than three months. Thus, as a stand-alone asset, this loan is very risky. However, we must consider the credit risk of the loan in a portfolio, allowing for imperfectly correlated default arrival times. That is, the timing of sudden jumps of default arrival intensities may be imperfectly correlated across loans. For simplicity, assume that other parameters (i.e., the sizes of the jumps in default intensities) are equal and independent across loans and across time, thereby fixing the parameter values θ, k, and J.

Correlations across loan default probabilities occur because common factors affect the timing of jumps in default probabilities across assets (loans). Specifically, the intensity jump time, λ, can be separated into a common factor with intensity V_c and an idiosyncratic factor, V. Thus,

$$\lambda = vV_c + V \tag{8.17}$$

where v is the sensitivity of the timing of jumps in default intensities to common factors.[34]

These common factors, V_c, can be viewed as macroeconomic factors, similar to those used in the multifactor models discussed earlier in this chapter. The correlation coefficient between the times to the next credit event for any pair of loans can be expressed as a simple function of v, V_c, and V. Pan and Singleton (2008) use sovereign credit default swap (CDS) spreads for Mexico, South Korea, and Turkey in order to estimate a similar jump process, assuming a single factor model. Although their model performs well at the longer maturities, they find pricing errors at the short-maturity, one-year contract, consistent with the existence of a liquidity premium in bond yields (see Chapters 5 and 7). Consistent with this view expressed by market practitioners, Pan and Singleton (2008) find larger bid-ask spreads for the one-year contract.

To illustrate this using a numerical example, Figure 8.3 shows a portion of a typical sample path for the total arrival intensity h of defaults for the following parameter values: $\lambda = .002$, $\theta = .001$, $k = .5$, $J = 5$, $v = .02$, and $V_c = .05$. Using equation (8.17), we can compute

$$V = .002 - (.02)(.05) = .001$$

We can also compute the probability that loan i's default intensity jumps at time t, given that loan j's intensity has experienced a jump, as:

$$vV_c/(V_c + V) = (.02)(.05)/(.05 + .001) = 2 \text{ percent}$$

Figure 8.3 shows a marketwide credit event occurring just prior to year 2.8 on the calendar time axis. This event instigates jumps in default intensity for several firms. These defaults are represented by the small x's along the bottom of the figure. Correlations across default intensities cause a rapid

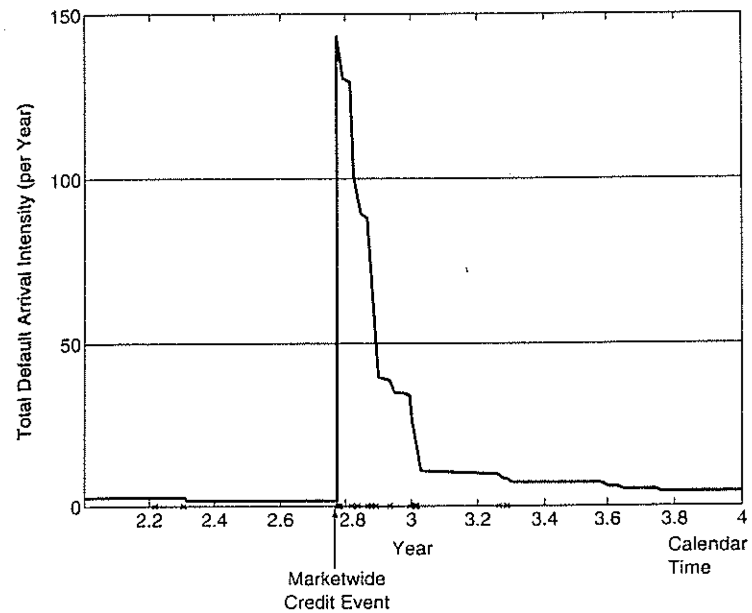

FIGURE 8.3 Correlated Default Intensity
Note: The figure shows a portion of a simulated sample path of total default arrival intensity (initially 1,000 firms). An x along the calendar time axis denotes a default event.
Source: Duffie and Singleton (1998).

increase in default risk in the period immediately surrounding the marketwide credit event. However, the mean reversion built into the intensity process (k is assumed to equal .5) causes the total arrival intensity for defaults to drop back almost to pre-event levels within one year.

Taking the scenario illustrated in Figure 8.3 as the base case, Duffie and Singleton (1998) also examine alternative correlation values: zero correlation ($v = V_c = 0$) and high correlations ($v = .02$ and $V_c = .1$). Figure 8.4 plots the probabilities of experiencing four or more defaults in any time window (of m days) for the three different assumptions about correlations: zero (low) correlation, medium correlation (the base case), and high correlation. Figure 8.4 shows the substantial impact that correlation has on the portfolio's credit risk exposure. This implies that the correlations in default risk shocks (i.e., the correlated jumps in default intensities) may make it

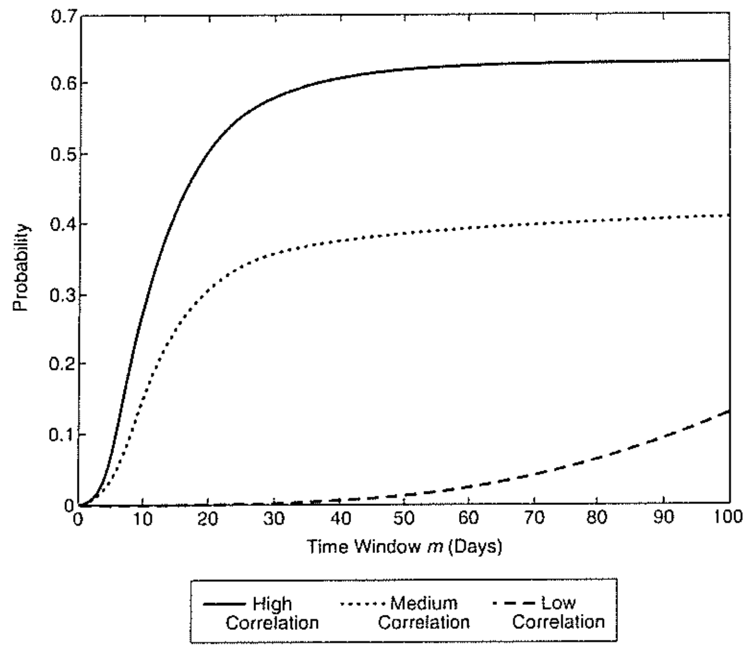

FIGURE 8.4 Portfolio Default Intensities
Note: The figure shows the probability of an m-day interval within 10 years experiencing four or more defaults (base case).
Source: Duffie and Singleton (1998).

difficult for banks to recapitalize within one year of experiencing defaults on the loans in their portfolios (see Carey [2001b]).

SUMMARY

Modern portfolio theory (MPT) provides an extremely useful framework for a loan portfolio manager considering risk/return trade-offs. The lower the correlation among loans in a portfolio, the greater the potential for a manager to reduce a bank's risk exposure through diversification. Furthermore, to the extent that a VAR-based capital requirement reflects the concentration risk and default correlations of the loan portfolio, such a portfolio may have lower credit risk than when loan exposures are considered independently additive.

In this chapter, we describe the methodologies used to measure default correlations in commercial models such as Moody's KMV and Kamakura's reduced form model. These models are derived from the academic literature that implemented the MPT using the capital asset pricing model (CAPM). Application of the CAPM to credit risk measurement has been expanded to incorporate multifactor asset pricing models to estimate the sensitivity of default risk to underlying macroeconomic and other risk factors. The default correlations for any pair of assets then can be computed using the factor loadings (the beta estimates) for each asset.