

# Portfolio Credit Risk

---

*Thomas C. Wilson*

## INTRODUCTION AND SUMMARY

Financial institutions are increasingly measuring and managing the risk from credit exposures at the portfolio level, in addition to the transaction level. This change in perspective has occurred for a number of reasons. First is the recognition that the traditional binary classification of credits into “good” credits and “bad” credits is not sufficient—a precondition for managing credit risk at the portfolio level is the recognition that all credits can potentially become “bad” over time given a particular economic scenario. The second reason is the declining profitability of traditional credit products, implying little room for error in terms of the selection and pricing of individual transactions, or for portfolio decisions, where diversification and timing effects increasingly mean the difference between profit and loss. Finally, management has more opportunities to manage exposure proactively after it has been originated, with the increased liquidity in the secondary loan market, the increased importance of syndicated lending, the availability of credit derivatives and third-party guarantees, and so on.

In order to take advantage of credit portfolio management opportunities, however, management must first answer several technical questions: What is the risk of a given portfolio? How do different macroeconomic scenarios, at both the regional and the industry sector level, affect the portfolio’s risk profile? What is the effect of changing the portfolio mix? How might risk-based pricing at the individual contract and the portfolio level be influenced by the level of expected losses and credit risk capital?

This paper describes a new and intuitive method for answering these technical questions by tabulating the exact loss distribution arising from correlated credit events for any arbitrary portfolio of counterparty exposures, down to the individual contract level, with the losses measured on a marked-to-market basis that explicitly recognises the potential impact of defaults and credit migrations.<sup>1</sup> The importance of tabulating the exact loss distribution is highlighted by the fact that counterparty defaults and rating migrations cannot be predicted with perfect foresight and are not perfectly correlated, implying that management faces a distribution of potential losses rather than a single potential loss. In order to define credit risk more precisely in the context of loss distributions, the financial industry is converging on risk measures that summarise management-relevant aspects of the entire loss distribu-

---

*Thomas C. Wilson is a principal of McKinsey and Company.*

tion. Two distributional statistics are becoming increasingly relevant for measuring credit risk: expected losses and a critical value of the loss distribution, often defined as the portfolio's credit risk capital (CRC). Each of these serves a distinct and useful role in supporting management decision making and control (Exhibit 1).

*Expected losses*, illustrated as the mean of the distribution, often serve as the basis for management's reserve policies: the higher the expected losses, the higher the reserves required. As such, expected losses are also an important component in determining whether the pricing of the credit-risky position is adequate: normally, each transaction should be priced with sufficient margin to cover its contribution to the portfolio's expected credit losses, as well as other operating expenses.

*Credit risk capital*, defined as the maximum loss within a known confidence interval (for example, 99 percent) over an orderly liquidation period, is often interpreted as the additional economic capital that must be held against a given portfolio, above and beyond the level of credit reserves, in order to cover its unexpected credit losses. Since it would be uneconomic to hold capital against *all* potential losses (this would imply that equity is held against 100 percent of all credit exposures), some level of

capital must be chosen to support the portfolio of transactions in most, but not all, cases. As with expected losses, CRC also plays an important role in determining whether the credit risk of a particular transaction is appropriately priced: typically, each transaction should be priced with sufficient margin to cover not only its expected losses, but also the cost of its marginal risk capital contribution.

In order to tabulate these loss distributions, most industry professionals split the challenge of credit risk measurement into two questions: First, what is the joint probability of a credit event occurring? And second, what would be the loss should such an event occur?


In terms of the latter question, measuring potential losses given a credit event is a straightforward exercise for many standard commercial banking products. The exposure of a \$100 million unsecured loan, for example, is roughly \$100 million, subject to any recoveries. For derivatives portfolios or committed but unutilised lines of credit, however, answering this question is more difficult. In this paper, we focus on the former question, that is, how to model the joint probability of defaults across a portfolio. Those interested in the complexities of exposure measurement for derivative and commercial banking products are referred to J.P. Morgan (1997), Lawrence (1995), and Rowe (1995).

The approach developed here for measuring expected and unexpected losses differs from other approaches in several important respects. First, it models the actual, discrete loss distribution, depending on the number and size of credits, as opposed to using a normal distribution or mean-variance approximations. This is important because with one large exposure the portfolio's loss distribution is discrete and bimodal, as opposed to continuous and unimodal; it is highly skewed, as opposed to symmetric; and finally, its shape changes dramatically as other positions are added. Because of this, the typical measure of unexpected losses used, standard deviations, is like a "rubber ruler": it can be used to give a sense of the uncertainty of loss, but its actual interpretation in terms of dollars at risk depends on the degree to which the ruler has been "stretched" by diversification or large exposure effects. In contrast, the model developed here explicitly tabulates the actual,

Exhibit 1

LOSS DISTRIBUTION

\$100 Portfolio, 250 Equal and Independent Credits with Default Probability Equal to 1 Percent



discrete loss distribution for any given portfolio, thus also allowing explicit and accurate tabulation of a “large exposure premium” in terms of the risk-adjusted capital needed to support less-diversified portfolios.

Second, the losses (or gains) are measured on a default/no-default basis for credit exposures that *cannot* be liquidated (for example, most loans or over-the-counter trading exposure lines) as well as on a theoretical marked-to-market basis for those that *can* be liquidated prior to the maximum maturity of the exposure. In addition, the distribution of average write-offs for retail portfolios is also modeled. This implies that the approach can integrate the credit risk arising from liquid secondary market positions and illiquid commercial positions, as well as retail portfolios such as mortgages and overdrafts. Since most banks are active in all three of these asset classes, this integration is an important first step in determining the institution’s overall capital adequacy.

Third, and most importantly, the tabulated loss distributions are driven by the state of the economy, rather than based on unconditional or twenty-year averages that do not reflect the portfolio’s true current risk. This allows the model to capture the cyclical default effects that determine the lion’s share of the risk for diversified portfolios. Our research shows that the bulk of the systematic or non-diversifiable risk of any portfolio can be “explained” by the economic cycle. Leveraging this fact is not only intuitive, but it also leads to powerful management insights on the true risk of a portfolio.

Finally, specific country and industry influences are explicitly recognised using empirical relationships, which enable the model to mimic the actual default correlations between industries and regions at the transaction and the portfolio level. Other models, including many developed in-house, rely on a single systematic risk factor to capture default correlations; our approach is based on a true multi-factor systematic risk model, which reflects reality better.

The model itself, described in greater detail in McKinsey (1998) and Wilson (1997a, 1997b), consists of two important components, each of which is discussed in greater detail below. The first is a *multi-factor model of sys-*

*tematic default risk.* This model is used to simulate jointly the conditional, correlated, average default, and credit migration probabilities for each individual country/industry/rating segment. These average segment default probabilities are made conditional on the current state of the economy and incorporate industry sensitivities (for example, “high-beta” industries such as construction react more to cyclical changes) based on aggregate historical relationships. The second is a method for tabulating the discrete loss distribution for any portfolio of credit exposures—liquid and nonliquid, constant and nonconstant, diversified and non-diversified. This is achieved by convoluting the conditional, marginal loss distributions of the individual positions to develop the aggregate loss distribution, with default correlations between different counterparties determined by the systematic risk driving the correlated average default rates.


### SYSTEMATIC RISK MODEL

In developing this model for systematic or nondiversifiable credit risk, we leveraged five intuitive observations that credit professionals very often take for granted.

First, that diversification helps to reduce loss uncertainty, all else being equal. Second, that substantial systematic or nondiversifiable risk nonetheless remains for even the most diversified portfolios. This second observation is illustrated by the “Actual” line plotted in Exhibit 2, which represents the average default rate for all German corporations over the

Exhibit 2

ACTUAL VERSUS PREDICTED DEFAULT RATES  
Germany




1960-94 period; the variation or volatility of this series can be interpreted as the systematic or nondiversifiable risk of the “German” economy, arguably a very diversified portfolio. Third, that this systematic portfolio risk is driven largely by the “health” of the macroeconomy—in recessions, one expects defaults to increase.

The relationship between changes in average default rates and the state of the macroeconomy is also illustrated in Exhibit 2, which plots the actual default rate for the German economy against the predicted default rate, with the prediction equation based solely upon macroeconomic aggregates such as GDP growth and unemployment rates. As the exhibit shows, the macroeconomic factors explain much of the overall variation in the average default rate series, reflected in the regression equation’s  $R^2$  of more than 90 percent for most of the countries investigated (for example, Germany, the United States, the United Kingdom, Japan, Switzerland, Spain, Sweden, Belgium, and France). The fourth observation is that different sectors of the economy react differently to macroeconomic shocks, albeit with different economic drivers: U.S. corporate insolvency rates are heavily influenced by interest rates, the Swedish paper and pulp industry by the real terms of trade, and retail mortgages by house prices and regional economic indicators. While all of these examples are intuitive, it is sometimes surprising how strong our intuition is when put to statistical tests. For example, the intuitive expectation that the construction sector would be more adversely affected during a recession than most other sectors is supported by the data for all of the different countries analysed.

Exhibit 3 illustrates the need for a multi-factor model, as opposed to a single-factor model, for systematic risk. Performing a principal-components analysis of the country average default rates, a good surrogate for systematic risk by country, it emerges that the first “factor” captures only 77.5 percent of the total variation in systematic default rates for Moody’s and the U.S., U.K., Japanese, and German markets. This corresponds to the amount of systematic risk “captured” by most single-factor models; the rest of the variation is implicitly

Exhibit 3

TOTAL SYSTEMATIC RISK EXPLAINED




Note: The factor 2 band for Japan is 79.7; the factor 3 band for the United Kingdom is 82.1.

assumed to be independent and uncorrelated. Unfortunately, the first factor explains only 23.9 percent of the U.S. systematic risk index, 56.2 percent for the United Kingdom, and 66.8 percent for Germany. The exhibit demonstrates that the substantial correlation remaining is explained by the second and third factors, explaining an additional 10.2 percent and 6.8 percent, respectively, of the total variation and the bulk of the risk for the United States, the United Kingdom, and Germany. This demonstrates that a single-factor systematic risk model like one based on asset betas or aggregate Moody’s/Standard and Poor’s data alone is not sufficient to capture all correlations accurately. The final observation is also both intuitive and empirically verifiable: that rating migrations are also linked to the macroeconomy—not only is default more likely during a recession, but credit downgrades are also more likely.

When we formulate each of these intuitive observations into a rigorous statistical model that we can estimate, the net result is a multi-factor statistical model for systematic credit risk that we can then simulate for every country/industry/rating segment in our sample. This is demonstrated in Exhibit 4, where we plot the simulated cumulative default rates for a German, single-A-rated, five-year exposure based on current economic conditions in Germany.

Exhibit 4

**SIMULATED DEFAULT PROBABILITIES**  
Germany, Single-A-Rated Five-Year Cumulative Default Probability



**LOSS TABULATION METHODS**

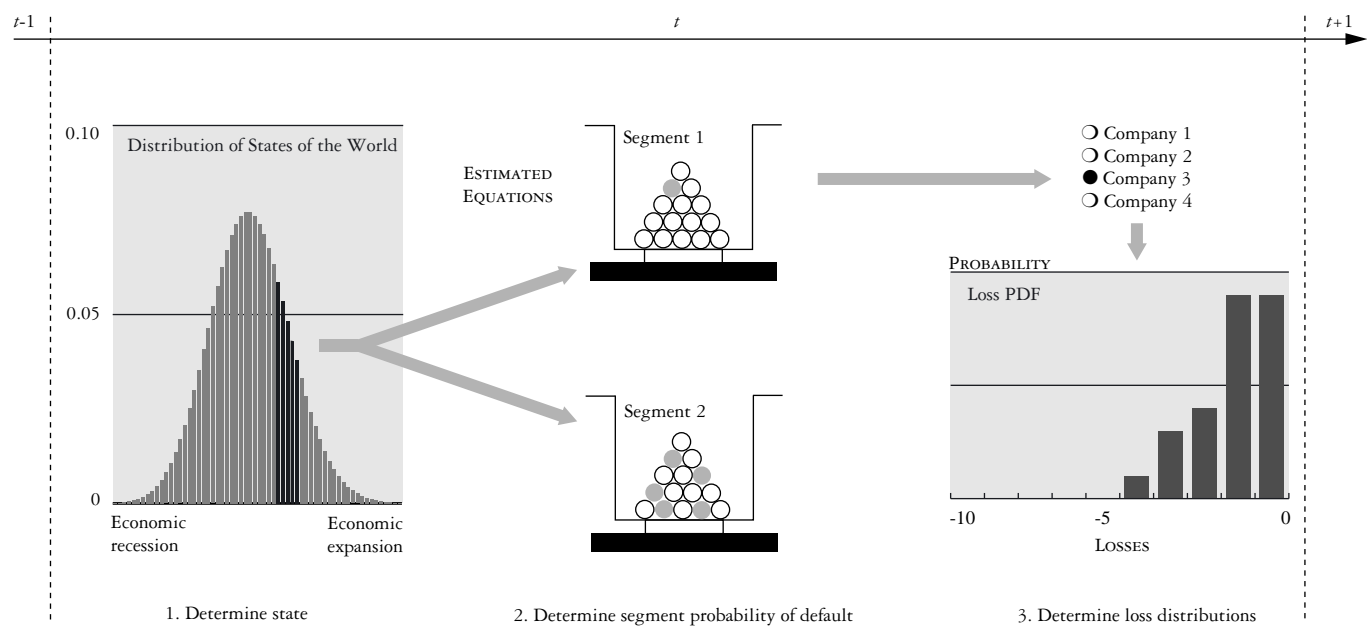
While these distributions of correlated, average default probabilities by country, sector, rating, and maturity are interesting, we still need a method of explicitly tabulating the loss distribution for any arbitrary portfolio of credit risk exposures. So we now turn to developing an efficient method for tabulating the loss distribution for

any arbitrary portfolio, capable of handling portfolios with large, undiversified positions and/or diversified portfolios; portfolios with nonconstant exposures, such as those found in derivatives trading books, and/or constant exposures, such as those found in commercial lending books; and portfolios comprising liquid, credit-risky positions, such as secondary market debt, or loans and/or illiquid exposures that must be held to maturity, such as some commercial loans or trading lines. Below, we demonstrate how to tabulate the loss distributions for the simplest case (for example, constant exposures, nondiscounted losses) and then build upon the simplest case to handle more complex cases (for example, nonconstant exposures, discounted losses, liquid positions, and retail portfolios). Exhibit 5 provides an abstract timeline for tabulating the overall portfolio loss distribution. The first two steps relate to the systematic risk model and the third represents loss tabulations.

Time is divided into discrete periods, indexed by  $t$ . During each period, a sequence of three steps occurs: first, the state of the economy is determined by simulation; second, the conditional migration and cumulative default probabilities for each country/industry segment

Exhibit 5

**MODEL STRUCTURE**



are determined based on the equations estimated earlier; and, finally, the actual defaults for the portfolio are determined by sampling from the relevant distribution of segment-specific simulated default rates. Exhibit 6 gives figures for the highly stylised single-period, two-segment numerical example described below.

1. *Determine the state:* For any given period, the first step is to determine the state of the world, that is, the health of the macroeconomy. In this simple example, three possible states of the economy can occur: an economic “expansion” (with GDP growth of +1 percent), an “average” year (with GDP growth of 0 percent), and an economic “recession” (with GDP growth of -1 percent). Each of these states can occur with equal probability (33.33 percent) in this numerical sample.

2. *Determine segment probability of default:* The second step is to then translate the state of the world into conditional probabilities of default for each customer segment based on the estimated relationships described earlier. In this example, there are two counterparty segments, a “low-beta” segment, whose probability of default reacts less strongly to macroeconomic fluctuations (with a range of 2.50 percent to 4.71 percent), and a “high-beta” segment, which reacts quite strongly to macroeconomic fluctuations (with a range of 0.75 percent to 5.25 percent).

3. *Determine loss distributions:* We now tabulate the (nondiscounted) loss distribution for portfolios that are constant over their life, cannot be liquidated, and have known recovery rates, including both diversified and non-

diversified positions. Later, we relax each of these assumptions within the framework of this model in order to estimate more accurately the expected losses and risk capital from credit events.

The conditional loss distribution in the simple two-counterparty, three-state numerical example is tabulated by recognising that there are three independent “draws,” or states of the economy and that, conditional on each of these states, there are only four possible default scenarios: A defaults, B defaults, A+B defaults, or no one defaults (Exhibit 7).

The conditional probability of each of these loss events for each state of the economy is calculated by convoluting each position’s individual loss distribution for each state. Thus, the conditional probability of a \$200 loss in the expansion state is 0.01 percent, whereas the unconditional probability of achieving the same loss given the entire distribution of future economic states (expansion, average, recession) is 0.1 percent after rounding errors. For this example, the expected portfolio loss is \$6.50 and the credit risk capital is \$100, since this is the maximum potential loss within a 99 percent confidence interval across all possible future states of the economy.

Our calculation method is based on the assumption that all default correlations are caused by the correlated segment-specific default indices. That is, no further information beyond country, industry, rating, and the state of the economy is useful in terms of predicting the default correlation between any two counterparties. To underscore this point, suppose that management is confronted with two single-A-rated counterparties in the German construction industry with the prospect of either a recession or an economic expansion in the near future. Using the traditional approach, which ignores the impact of the economy in determining default probabilities, we would conclude that the counterparty default rates were correlated. Using our approach, we observe that, in a recession, the probability of default for both counterparties is significantly higher than during an expansion and that their joint conditional probability of default is therefore also higher, leading to correlated defaults. However, because we assume that all idiosyncratic or nonsystematic risks can be diversified

Exhibit 6

NUMERICAL EXAMPLE

1. Determine state	State	GDP	Probability of Default (Percent)
	Expansion	+1	33.33
	Average	0	33.33
	Recession	-1	33.33

2. Determine segment probability of default	State	Low-Beta Probability of Default A (Percent)	High-Beta Probability of Default B (Percent)
	Expansion	2.50	0.75
	Average	2.97	3.45
	Recession	4.71	5.25

3. Determine loss distributions
---------------------------------

away, no other information beyond the counterparties' country, industry, and rating (for example, the counterparties' segmentation criteria) is useful in determining their joint default correlation. This assumption is made implicitly by other models, but ours extends the standard single-factor approach to a multi-factor approach that better captures country- and industry-specific shocks.

Intuitively, we should be able to diversify away all idiosyncratic risk, leaving only systematic, nondiversifiable risk. More succinctly, as we diversify our holdings within a particular segment, that segment's loss distribution will converge to the loss distribution implied by the segment index. This logic is consistent with other single- or multi-factor models in finance, such as the capital asset pricing model.

Our multi-factor model for systematic default risks is qualitatively similar, except that there is no single risk factor. Rather, there are multiple factors that fully describe the complex correlation structure between countries, industries, and ratings. In our simple numerical example, for a well-diversified portfolio consisting of a large number of counterparties in each segment (the NA & NB = Infinity case), all idiosyncratic risk per segment is

diversified away, leaving only the systematic risk per segment (Exhibit 8).

In other words, because of the law of large numbers, the actual loss distribution for the portfolio will converge to the expected loss for each state of the world, implying that the unconditional loss distribution has only three possible outcomes, representing each of the three states of the world, each occurring with equal probability and with a loss per segment consistent with the conditional probability of loss for that segment given that state of the economy. While the expected losses from the portfolio would remain constant, this remaining systematic risk would generate a CRC value of only \$9.96 for the \$200 million exposure in this simple example, demonstrating both the benefit to be derived from portfolio diversification and the fact that not all systematic risk can be diversified away.

In the second case (labeled NA = 1 & NB = Infinity), all of the idiosyncratic risk is diversified away within segment B, leaving only the systematic risk component for segment B. The segment A position, however, still contains idiosyncratic risk, since it comprises only a single risk position. Thus, for each state of the economy, two outcomes


*Exhibit 7*

NUMERICAL EXAMPLE: TWO EXPOSURES

1. Determine state
2. Determine segment probability of default
3. Determine loss distributions

Loss Distribution	Expansion				Average				Recession			
	A	B	A+B	Probability of Default (Percent)	A	B	A+B	Probability of Default (Percent)	A	B	A+B	Probability of Default (Percent)
-100	-100	-200	0.01	-100	-100	-200	0.03	-100	-100	-200	0.08	
-100	0	-100	0.83	-100	0	-100	0.96	-100	0	-100	1.49	
0	-100	-100	0.24	0	-100	-100	1.12	0	-100	-100	1.67	
0	0	0	32.36	0	0	0	31.23	0	0	0	30.10	
Correlation (A,B) = 0 percent				Correlation (A,B) = 0 percent				Correlation (A,B) = 0 percent				
Conditional correlation (A,B) = 1 percent												

PROBABILITY OF LOSS EVENT



NUMERICAL EXAMPLE: DIVERSIFIED EXPOSURES

1. Determine state
2. Determine segment probability of default
3. Determine loss distributions

	NA & NB = Infinity			Probability of Default (Percent)
	Loss			
	A	B	A+B	
Expansion	-2.50	-0.75	-3.25	33.33
Average	-2.97	-3.45	-6.42	33.33
Recession	-4.71	-5.25	-9.96	33.33
	Unconditional correlation (A, B)			91.00
	Credit RAC = 9.96			

	NA = 1 & NB = Infinity			Probability of Default (Percent)
	Loss			
	A	B	A+B	
Expansion	-100	-0.75	-100.75	0.83
	0	-0.75	-0.75	32.50
Average	-100	-3.45	-103.45	0.99
	0	-3.45	-3.45	32.30
Recession	-100	-5.25	-105.25	1.57
	0	-5.25	-5.25	31.80
	Credit RAC = 105.25			

are possible: either the counterparty in segment A goes bankrupt or it does not; the unconditional probability that counterparty A will default in the economic expansion state is 0.83 percent (33.33 percent probability that the expansion state occurs multiplied by a 2.5 percent probability of default for a segment A counterparty given that state). Regardless of whether or not counterparty A goes into default, the segment B position losses will be known with certainty, given the state of the economy, since all idiosyncratic risk within that segment has been diversified away.

To illustrate the results using our simulation model, suppose that we had equal \$100, ten-year exposures to single-A-rated counterparties in each of five country segments—Germany, France, Spain, the United States, and the United Kingdom—at the beginning of 1996. The aggregate simulated loss distribution for this portfolio of diversified country positions, conditional on the then-current macroeconomic scenarios for the different countries at the end of 1995, is given in the left panel of Exhibit 9.

The impact of introducing one large, undiversified exposure into the same portfolio is illustrated in the right panel of Exhibit 9. Here, we take the same five-country portfolio of diversified index positions used in the left panel, but add a single, large, undiversified position to the “other” country’s position.

The impact of this new, large concentration risk is clear. The loss distribution becomes “bimodal,” reflecting the fact that, for each state of the world, two events might occur: either the large counterparty will go bankrupt, generating a “cloud” of portfolio loss events centered around -140, or the

undiversified position will not go bankrupt, generating a similar cloud of loss events centered around -40, but with higher probability. This risk concentration disproportionately increases the amount of risk capital needed to support the portfolio from \$61.6 to \$140.2, thereby demonstrating the large-exposure risk capital premium needed to support the addition of large, undiversified exposures.

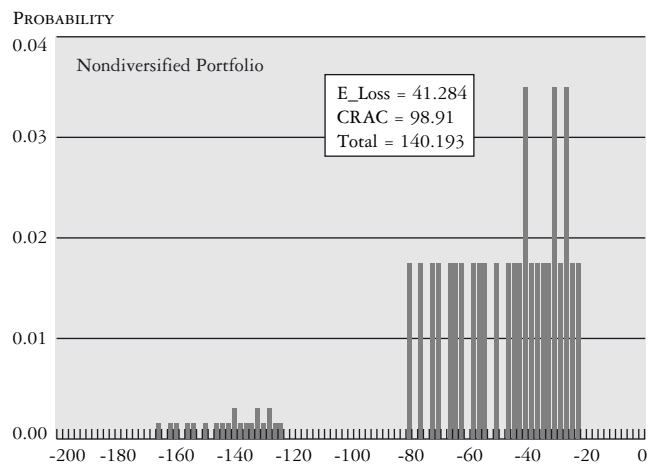

The calculations above illustrate how to tabulate the (nondiscounted) loss distributions for nonliquid portfolios with constant exposures. While useful in many instances, these portfolio characteristics differ from reality in two important ways. First, the potential exposure profiles generated by trading products are typically not constant (as pointed out by Lawrence [1995] and Rowe [1995]). Second, the calculations ignore the time value of money, so that a potential loss in the future is somehow “less painful” in terms of today’s value than a loss today.

In reality, the amount of potential economic loss in the event of default varies over time, due to discounting, or nonconstant exposures, or both. This can be seen in Exhibit 10. If the counterparty were to go into default sometime during the second year, the present value of the portfolio’s loss would be \$50 in the case of nonconstant exposures and  $\$100 * e^{-r_2 * 2}$  in the case of discounted exposures, as opposed to \$100 and  $\$100 * e^{-r_1 * 1}$  if the counterparty had gone into default sometime during the first year. Unlike the case of constant, nondiscounted exposures, where the timing of the default is inconsequential, nonconstant exposures or discounting of the losses implies that the timing of the default is critical for tabulating the economic loss.



EXAMPLES OF PORTFOLIO LOSS DISTRIBUTIONS

Portfolio Loss Distribution



Note: Business unit, book, country, rating, maturity, exposure.

Addressing both of these issues requires us to work with *marginal*, as opposed to *cumulative*, default probabilities. Whereas the cumulative default probability is the aggregate probability of observing a default in *any* of the previous years, the marginal default probability is the probability of observing a loss in each specific year, given that the default has not already occurred in a previous period.

Exhibit 11 illustrates the impact of nonconstant loss exposures in terms of tabulating loss distributions. With constant, nondiscounted exposures, the loss distribution for a single exposure is bimodal. Either it goes into default at some time during its maturity, with a cumulative default probability covering the entire three-year period equal to  $p_1 + p_2 + p_3$  in the exhibit, implying a loss of 100, or it does not. If the exposure is nonconstant, how-

ever, you stand to lose a different amount depending upon the exact timing of the default event. In the above example, you would lose 100 with probability  $p_1$ , the marginal probability that the counterparty goes into default during the first year; 50 with probability  $p_2$ , the marginal probability that the counterparty goes into default during the second year; and so on.

So far, we have been simulating only the cumulative default probabilities. Tabulating the marginal default probabilities from the cumulative is a straightforward exercise. Once this has been done, the portfolio loss distribution can be tabulated by convoluting the individual loss distributions, as described earlier. The primary difference between our model and other models is that we explicitly recognise that loss distributions for nonconstant exposure profiles are not binomial but multinomial, recognising the fact that the timing of default is also important in terms of tabulating the position's marginal loss distribution.

Exhibit 10

NONCONSTANT OR DISCOUNTED EXPOSURES


Credit Event Tree	Exposure Loss Profile	
	Nonconstant	Discounted <sup>a</sup>
No default		
Default, year three	25	$100 * e^{(-r_3 * 3)}$
Default, year two	50	$100 * e^{(-r_2 * 2)}$
Default, year one	100	$100 * e^{(-r_1 * 1)}$

<sup>a</sup> $r_1$  is the continuously compounded, per annum zero coupon discount rate.

LIQUID OR TRADABLE POSITIONS AND/OR ONE-YEAR MEASUREMENT HORIZONS

So far, we have also assumed that the counterparty exposure must be held until maturity and that it cannot be liquidated at a "fair" price prior to maturity; under such

NONCONSTANT OR DISCOUNTED EXPOSURES



circumstances, allocating capital and reserves to cover potential losses over the life of the asset may make sense. Such circumstances often arise in intransparent segments where the market may perceive the originator of the credit to have superior information, thereby reducing the market price below the underwriter’s perceived “fair” value. For some other asset classes, however, this assumption is inadequate for two reasons:

- Many financial institutions are faced with the increasing probability that a bond name will also show up in their loan portfolio. So they want to measure the credit risk contribution arising from their secondary bond trading operations and integrate it into an overall credit portfolio perspective.
- Liquid secondary markets are emerging, especially in the rated corporate segments.

In both cases, management is presented with two specific measurement challenges. First, as when measuring market risk capital or value at risk, management must decide on the appropriate time horizon over which to measure the potential loss distribution. In the previous illiquid asset class examples, the relevant time horizon coincided with the maximum maturity of the exposure, based on the assumption that management could not liquidate the position prior to its expiration. As markets become more liquid, the appropriate time horizons may be asset-dependent and determined by the asset’s orderly liquidation period.


The second challenge arises in regard to tabulating the marked-to-market value losses for liquid assets should a credit event occur. So far, we have defined the loss distribution only in terms of default events (although default probabilities have been tabulated using rating migrations as well). However, it is clear that if the position can be liquidated prior to its maturity, then other credit events (such

as credit downgrades and upgrades) will affect its marked-to-market value at any time prior to its ultimate maturity. For example, if you lock in a single-A-rated spread and the credit rating of the counterparty decreases to a triple-B, you suffer an economic loss, all else being equal: while the market demands a higher, triple-B-rated spread, your commitment provides only a lower, single-A-rated spread.

In order to calculate the marked-to-market loss distribution for positions that can be liquidated prior to their maturity, we therefore need to modify our approach in two important ways. First, we need not only simulate the cumulative default probabilities for each rating class, but also their migration probabilities. This is straightforward, though memory-intensive. Complicating this calculation, however, is the fact that if the time horizons are different for different asset classes, a continuum of rating migration probabilities might need to be calculated, one for each possible maturity or liquidation period. To reduce the complexity of the task, we tabulate migration probabilities for yearly intervals only and make the expedient assumption that the rating migration probabilities for any liquidation horizon that falls between years can be approximated by some interpolation rule.

Second, and more challenging, we need to be able to tabulate the change in marked-to-market value of the exposure for each possible change in credit rating. In the case of traded loans or debt, a pragmatic approach is simply to define a table of average credit spreads based on current market conditions, in basis points per annum, as a function of rating and the maturity of the underlying exposure. The potential loss (or gain) from a credit migration can then be tabulated by calculating the change in marked-to-market value of the exposure due to the changing of the discount rate implied by the credit migration.

MARKED-TO-MARKET CREDIT EVENT  
PROFIT/LOSS DISTRIBUTION



The results of applying this approach are illustrated in Exhibit 12, which tabulates the potential profit and loss profile from a single traded credit exposure, originally rated triple-B, which can be liquidated prior to one year. For this example, we have used a recovery rate of 69.3 percent, a proxy for the average recovery rate for senior secured credits rated triple-B. Inspection of Exhibit 12 shows that it is inappropriate to talk about “loss distributions” in the context of marked-to-market loan or debt securities, since a profit or gain in marked-to-market value can also be created by an improvement in the counterparty’s credit standing.

Although this approach allows us to capture the impact of credit migrations while holding the level of

interest rates and spreads constant, it must be seen as a complement to a market risk measurement system that accurately captures the potential profit-or-loss impact of changing interest rate and average credit spread levels. If your market risk measurement system does not capture these risks, then a more complicated approach could be used, such as jointly simulating interest rate levels, average credit spread levels, and credit rating migrations.

RETAIL PORTFOLIOS

Tabulating the losses from retail mortgage, credit card, and overdraft portfolios proceeds along similar lines. However, for such portfolios, which are often characterised by large numbers of relatively small, homogeneous exposures, it is frequently expedient to simulate *directly* the average loss or write-off rate for the portfolio under different macroeconomic scenarios based on similar, estimated equations as those described earlier, rather than migration probabilities for each individual obligor. Once simulated, the loss contribution under a given macroeconomic scenario for the first year is calculated as  $P_1 * LEE_1$ , for the second year as  $P_2 * (1 - P_1) * LEE_2$ , and so on, where  $P_i$  and  $LEE_i$  are the average simulated write-off rates and loan equivalent exposures for year  $i$ , respectively.

A bank’s aggregate loss distribution across its total portfolio of liquid, illiquid, and retail assets can be tabulated by applying the appropriate loss tabulation method to each asset class.

---

## ENDNOTE

1. This approach is embedded in CreditPortfolioView™, a software implementation of McKinsey and Company.

---

## REFERENCES

- Credit Suisse First Boston*. 1997. "CreditRisk+: Technical Documentation." London: Credit Suisse First Boston.
- Kealhofer, Stephen*. 1995a. "Managing Default Risk in Portfolios of Derivatives." In *DERIVATIVE CREDIT RISK: ADVANCES IN MEASUREMENT AND MANAGEMENT*. London: Risk Publications.
- . 1995b. "Portfolio Management of Default Risk." San Francisco: KMV Corporation.
- Lawrence, D.* 1995. "Aggregating Credit Exposures: The Simulation Approach." In *DERIVATIVE CREDIT RISK: ADVANCES IN MEASUREMENT AND MANAGEMENT*. London: Risk Publications.
- McKinsey and Company*. 1998. "CreditPortfolioView™ Approach Documentation and User's Documentation." Zurich: McKinsey and Company.
- Moody's Investors Service*. 1994. *CORPORATE BOND DEFAULTS AND DEFAULT RATES, 1970-1993*. New York: Moody's Investors Service.
- Morgan, J.P.* 1997. "CreditMetrics: Technical Documentation." New York: J.P. Morgan.
- Rowe, D.* 1995. "Aggregating Credit Exposures: The Primary Risk Source Approach." In *DERIVATIVE CREDIT RISK: ADVANCES IN MEASUREMENT AND MANAGEMENT*. London: Risk Publications.
- Wilson, Thomas C.* 1997a. "Credit Portfolio Risk (I)." *RISK MAGAZINE*, October.
- . 1997b. "Credit Portfolio Risk (II)." *RISK MAGAZINE*, November.

*The views expressed in this article are those of the author and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. The Federal Reserve Bank of New York provides no warranty, express or implied, as to the accuracy, timeliness, completeness, merchantability, or fitness for any particular purpose of any information contained in documents produced and provided by the Federal Reserve Bank of New York in any form or manner whatsoever.*