PORTFOLIO CREDIT RISK (II)

Thomas Wilson describes the tabulation of loss distributions for a macroeconomic model of credit risk

LAST MONTH, WE DEVELOPED a multifactor model for systemic default and credit migration risk based on historical macroeconomic and average default rate time series for different country/industry segments. In developing this model for systematic or non-diversifiable credit risk, we made use of several intuitive observations that credit managers very often take for granted.

We formulated each of these intuitive observations into a rigorous statistical model, which we then estimated. While the resulting distributions of correlated average default probabilities are interesting, we still need a way to tabulate explicitly the loss distribution for any arbitrary portfolio of credit risk exposures. To this end, we will now turn our attention to developing an efficient method of doing this, which is capable of handling:

☐ portfolios with large, undiversified positions and/or diversified positions;

☐ portfolios with non-constant 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) and/or retail portfolios such as mortgages, overdrafts and credit cards.

In our model[1], time is divided into discrete periods, indexed by t, and there are three steps in each period. To make these more tangible, we will consider a single-period, two-segment numerical example (see table A). The steps are:

☐ **Determine the state.** The first step during any given period is to determine the state of the world, ie, the macroeconomic health of the economy. In our simple example, there are three possible states of the economy that can occur: an economic "expansion" (with GDP growth of +1%), an "average" year (with GDP growth of 0%) and an economic "recession" (with GDP growth of –1%). Each can occur with equal probability (33.33%).

☐ **Determine the probability of default for each segment.** The second step is to translate the state of the world into conditional migration and default probabilities for each customer segment. In this example, there are two counterparty segments: a "low-beta" segment, where "beta" relates to the segment's default rates relative to the credit cycle and not to an equity cycle whose default probability reacts less strongly to macroeconomic fluctuations than a high-beta segment. These defau probabilities would be driven by the statistical systematic risk mo els described last month.

☐ **Determine loss distributions.** The third step is to determine the a tual loss distribution for the portfolio. We will demonstrate this by ta ulating the (non-discounted) loss distribution for portfolios that are co stant over their life, cannot be liquidated and have known recovery rate later, we will show how to incorporate liquid traded assets, retail po folios, random recovery rates and country defaults into the analysis.

Undiversified portfolios

The conditional loss distribution in the simple two-counterparty, thre state numerical example is tabulated by recognising that there are thr independent "draws" (or states) of the economy and that, condition on each of these states, there are only four possible default scenari (ie, A defaults, B defaults, both A and B default or neither default as shown in table B.

The conditional probability of each of these loss events for each sta of the economy is calculated by convoluting or aggregating each pos tion's marginal loss distribution under the assumption of independen for each state. Thus, the conditional probability of a $200 loss in t expansion state is 0.01%, whereas the unconditional probability of achie ing the same loss, given the entire distribution of future economic stat (ie, expansion, average, recession) is 0.1%. For the example shown i table B, the expected portfolio loss is $6.50 and the credit risk capital $100 (since this is the maximum potential loss within a 99% confiden interval across all possible future states of the economy).

Our approach for tabulating loss distributions is, therefore, fir to tabulate the conditional portfolio loss distribution for each sta of the world, given that counterparty defaults and credit migratio are independent, *conditional* on that state of the world; then, we a gregate these conditional loss distributions to an unconditional lo distribution by recognising that each was generated by an indepe dent, random draw from the possible states of the world. The im plicit assumption underlying our calculation of the conditional lo distributions is that all default and migration correlations are full determined by the systematic risk model. That is, no further info mation beyond country, industry, rating and the state of the econ my is useful in terms of predicting the default correlation betwee any two counterparties.

To underscore this point, suppose that management were confronte with two single-A counterparties in the German construction indust with the prospect of either a recession or an economic expansion in th near future. Using the traditional approach, which ignores the impa of the economy in determining the default probabilities for this an other segments, we would conclude that the counterparty default rat were correlated. Using our approach, we observe that the probabilit of default for both counterparties is significantly higher in a recessio than in an expansion since credit cycles are "caused" by macroec nomic cycles, leading to correlated defaults. However, because we a sume that default and migration correlations are fully determined b the segment's systematic risk, no other information beyond the cou terparties' country, industry and rating (the counterparties' segme tation criteria, say) are useful in determining their joint default corr lation once the state of the economy is determined. Using our approac we know that the two German, single-A construction companies bot have a higher probability of default during a recession but, given th

[1] McKinsey & Co has implemented this model within an application which will be made available free of charge. Contact author for details

A. A numerical example

1. Determine state	**State**	**GDP**	**Prob**
	Expansion	+1	33.33%
	Average	0	33.33%
	Recession	–1	33.33%
2. Determine segment probability default	**State**	**"Low-beta" Prob A**	**"High-beta" Prob B**
	Expansion	2.50%	0.75%
	Average	2.97%	3.45%
	Recession	4.71%	5.25%
3. Determine loss distributions			

B. Numerical example: two counterparty exposures

Loss distribution	Expansion A	Expansion B	Expansion A+B	Expansion Prob	Average A	Average B	Average A+B	Average Prob	Recession A	Recession B	Recession A+B	Recession Prob
	–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	32.26%	0	0	0	31.23%	0	0	0	30.10%
	Corr(A,B) = 0%				Corr(A,B) = 0%				Corr(A,B) = 0%			

Conditional corr(A,B) = 1%

Probability of loss

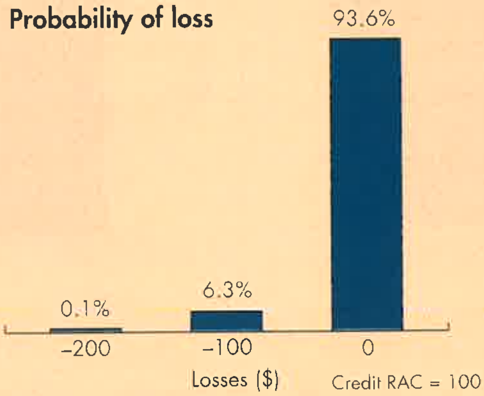

C. Numerical example: diversified counterparty exposures

NA + NB = infinity

	Loss A	B	A+B	Prob
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 Corr(A,B)			91%
	Credit RAC = 9.96			

NA = 1 and NB = infinity

	Loss A	B	A+B	Prob
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			

we are in a recession, actual defaults for the two are independent. This is the identical assumption implicitly made by other multifactor models such as CreditMetrics (JP Morgan, 1997) and KMV's portfolio model (Kealhofer, 1995i & ii), but ours extends the framework by allowing the systematic risk profile of different corporate and retail segments to be estimated using actual loss histories as well as information from equity markets, allowing the institution to capture better the correlations and momentum found in actual default data.

As an aside, if the recovery rate were known with certainty, we could adjust the conditional loss distributions appropriately, to reflect the fact that a default event did not imply that the full exposure was lost. For example, if the recovery rate was 71.57% (Standard & Poor's average recovery rate for triple-B senior secured debt), then the expected losses would be $1.85 and the credit risk capital would be $28.43. Note that the assumption of a constant recovery rate will be relaxed later.

Portfolio diversification

Intuitively, we should be able to diversify away all idiosyncratic risk as we increase the number of counterparties in each segment, leaving only systematic (non-diversifiable) risk. Put 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-factor or multifactor models in finance, eg, the behaviour of an equity portfolio under the capital asset pricing model.

Our multifactor model for systemic 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 in table C, for a well-diversified portfolio consisting of many counterparties in each segment (NA+NB = infinity), all idiosyncratic risk in each segment is diversified away, leaving only the systematic risk in each segment.

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 occurs 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 credit risk capital value of only $9.96 for the $200 million exposure in this simple example, demonstrating not only the benefit to be derived from portfolio diversification but also the fact that not all systematic risk can be diversified away.

In the second case in table C (NA = 1 and NB = infinity), all of the idiosyncratic risk is diversified away within segment B, leaving only the systematic risk component. The segment A position, on the other hand, still contains idiosyncratic risk, since it comprises only a single risk position. Thus, for each state of the economy, two possible out-

1. Example portfolio loss distribution: diversified position

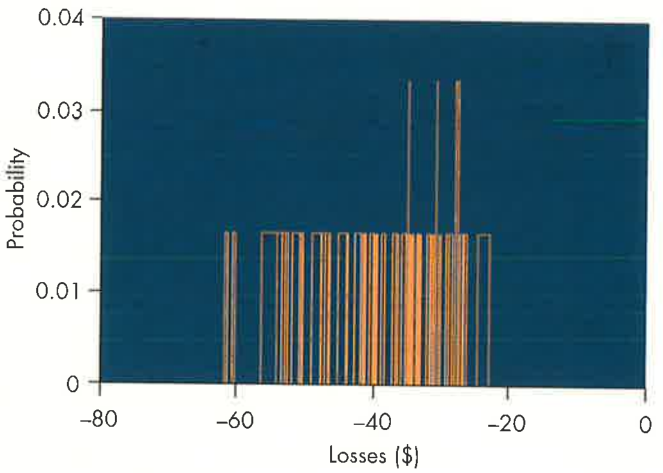

Expected loss = 37.545 Credit risk capital = 24.027 Total = 61.572

2. Example portfolio loss distribution: undiversified position

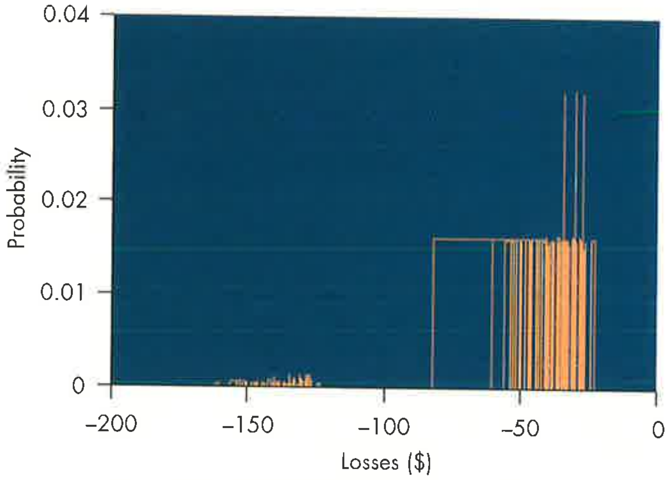

Expected loss = 41.284 Credit risk capital = 98.91 Total = 140.193

3. Non-constant or discounted exposures

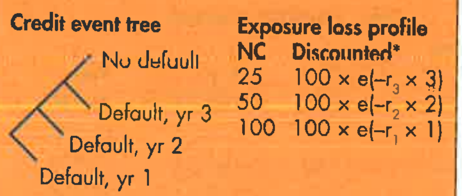

NC = non-constant
* r_i is the continuously compounded, yearly zero-coupon discount rate

4. Non-constant or discounted exposures

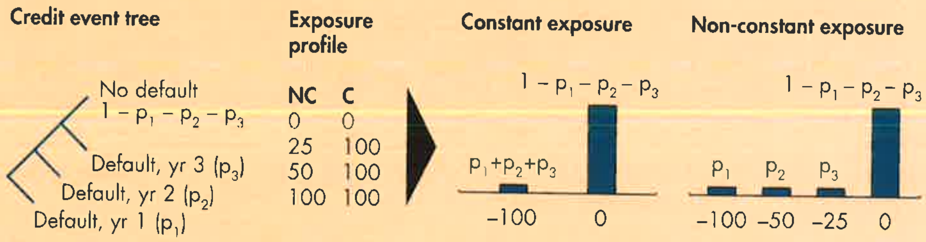

NC = non-constant; C = constant

comes 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% (a 33.33% probability that the expansion state occurs, multiplied by a 2.5% probability of default for a segment A counterparty, given that state). Regardless of whether counterparty A goes into default or not, 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 we had equal $100, 10-year exposures to single-A counterparties in each of five country segments – Germany, France, Spain, the US and the UK – at the beginning of 1996. The aggregate simulated loss distribution for this portfolio of diversified country positions, conditional on the then-current macroeconomic states for the different countries at the end of 1995, is given in figure 1. This portfolio has an expected loss of $37.5 million and would require an additional $24 million in risk capital to support it over the life of the commitments.

We now introduce a single large, undiversified exposure into the same portfolio. The impact is clearly visible in figure 2. The return distribution has become 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 centred around –140, or it will not, generating a similar, but higher probability, "cloud" centred around –40. This calculation therefore demonstrates the risk capital premium needed to support the addition of a large, undiversified exposure.

Variable exposures and discounted losses

The calculations above illustrate how to tabulate the (non-discounted) loss distributions for non-liquid 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 not typically constant (as pointed out in Wilson, 1997, Lawrence, 1995 and Rowe, 1995). Second, the calculations ignore the time value of money, so a potential loss in the future is deemed as "painful", in terms of today's value, as a loss today.

In reality, the value of potential economic loss in the event of default varies over time. This poses a further challenge for our model, even supposing that the potential exposure envelope has been calculated for all counterparties and their netting/collateral sets – a difficult enough task in its own right.

This can be seen in figure 3. If the counterparty had gone into default some time during the first year, the present value (PV) of the portfolio's loss would be $100 in the case of non-constant exposures and $100e(–$r_1$×1) in the case of discounted exposures, where r_1 is the continuously compounded, one-year, zero-coupon rate. If one counterparty were to default in year two, however, the PV of the portfolio's loss would be $50 in the case of non-constant exposures and

$\$100e(-r_2 \times 2)$ in the case of discounted exposures. Unlike the case of constant, non-discounted exposures, where the timing of the default is inconsequential, non-constant exposures or discounting of the losses imply that the timing of the default is critical in terms of tabulating the potential economic loss.

Addressing both of these issues requires us to work with marginal, rather than cumulative, default probabilities. Whereas the cumulative default probability is the probability of observing a default in any of the prior years, the marginal default probability is the probability of observing a loss in each specific year (assuming that the loss has not already occurred in a previous period).

Figure 4 illustrates the impact of non-constant loss exposures in terms of tabulating loss distributions. With constant 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 figure, implying a loss of 100 – or it does not. If the exposure varies, however, you stand to lose a different amount depending upon the exact timing of the default event. In the above example, you would stand to 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.

Until now, we have simulated only the cumulative default probabilities; the tabulation of the marginal default probabilities from the cumulative is a straightforward exercise. Once done, the portfolio loss distribution can be tabulated by convoluting the individual loss distributions, in the same manner as described earlier. The primary difference between our model and other models is that we explicitly recognise that loss distributions for non-constant exposure profiles are not binomial but multinomial, due to the fact that the timing of default is just as important in terms of tabulating the position's loss distribution.

Liquid assets

Until now, we have also assumed that the counterparty exposure must be held until maturity and that it cannot be liquidated at some market price before its maturity or expiry. This assumption is inadequate for three reasons:

☐ Many financial institutions are faced with the increasing probability that a bond name will also show up in their loan portfolio. As such, they want to measure the credit risk contribution arising from their secondary bond trading operations and integrate it into an overall credit portfolio perspective.

☐ While this assumption has been appropriate historically for many traditional asset classes (especially corporate loans and credit lines to support off-balance-sheet products), there has been a sharp increase in tools that allow credit exposures to be managed after their origination. Examples of such emerging tools include secondary and securitised loan trading; credit portfolio exchanges or "swaps"; credit derivatives; third party guarantees or insurance, etc. Having said this, it is nonetheless the case that most of the credit risk held by many retail and commercial banks – in the relatively opaque retail and mid-market segments – remains illiquid, especially in Europe and Asia.

☐ Finally, many financial institutions choose to tabulate the actual loss distribution for holding period horizons that might differ from the "liquidation" period of the commitment, reflecting, say, the periodicity of their capital allocation planning and budgeting process.

In all three cases, management is presented with two specific measurement challenges. First, as with 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 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 expiry. This assumption is commonly re-

5. Marked-to-market credit event profit/loss distribution

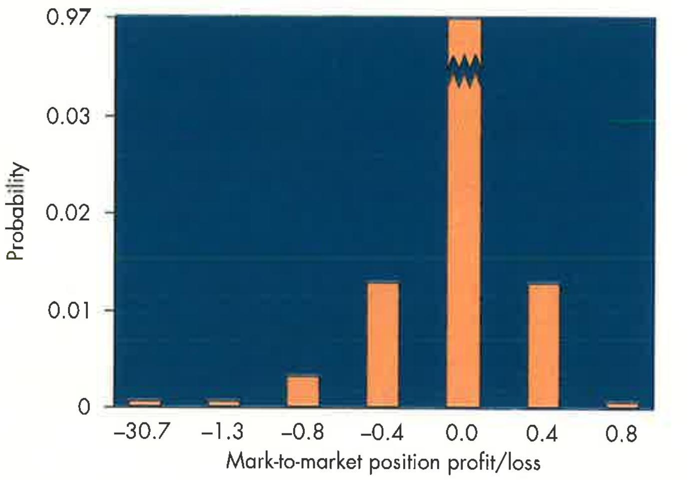

6. Distribution of defaulted bond prices

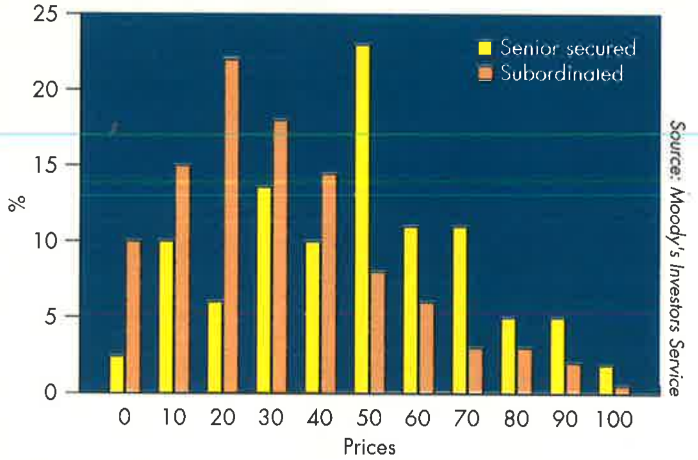

laxed either by specifying a common holding period consistent with the institution's planning and budgeting process, thereby ignoring the risk capital needed to support the positions beyond the first year, or by specifying a position-specific holding period horizon, coinciding with the time it would take management to recognise a problem and liquidate the position in an orderly manner. This latter approach requires that management has a clear perspective regarding the market liquidity for each individual position.[2]

The second challenge arises in regard to tabulating the marked-to-market value losses should a credit event occur. Until now we have defined the loss distribution only in terms of default events (although default probabilities have been tabulated using rating migrations as well). It is clear, however, that if the position can be liquidated prior to its maturity, then other credit events (eg, credit downgrades and upgrades) may affect its marked-to-market value at any time before its expiry. For example, if you lock in a single-A spread and the credit rating of the counterparty decreases to triple-B, you suffer an economic loss.

To calculate the marked-to-market distribution for positions that can be liquidated prior to maturity, we need to modify our approach in two important ways. First, we must simulate not only the cumulative default probabilities for each rating class but also store in memory the probabilities of counterparty rating migrations. This is straightforward, although it is 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 stored – 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 assumption that the rating migration probabilities for any liquidation horizon that falls between years can be approximated by some interpolation method.[3]

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, in basis points per year, 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.[4]

The results of applying this approach are shown in figure 5, where we have tabulated 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 71.57%, the average recovery rate for senior secured credits rated triple-B. As is clear from figure 5, it is inappropriate to talk about "loss" distributions in the context of marked-to-market loan or debt securities, since it is also possible to *profit* from an improvement in the counterparty's credit standing.

Although this approach allows one 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 levels and average credit spread levels. If your market risk measurement system does not capture these risks, a more complicated approach could be used, eg, simulating jointly interest rate levels, average credit spread levels and credit rating migrations. While potentially desirable, such an approach may come close to the fabled "all singing, all dancing" models that our grandmothers warned us about.

Variable recovery rates and country defaults

Throughout our discussions so far, we have also assumed that the amount that could be recovered in the event of default was a known constant, dependent only upon the rating of the counterparty. In reality, however, the actual amount that one can recover is neither constant nor is it dependent solely on the rating of the counterparty, as demonstrated by the Moody's Investors Service (1994) data plotted in figure 6. Clearly, the market's expectation of potential recovery from defaulted securities depends upon the individual security.

One way to model the impact of random recovery rates is to simulate jointly defaults and recovery rates. As is apparent from figure 6, however, we must first specify a recovery distribution for each of the relevant recovery classes. We do this for various classes (senior secured, subordinated, unsecured, etc), by using Moody's, Standard & Poor's or your own portfolio's historical experiences. For every simulation of the

2 Complicating this process is the fact that, while any credit exposure can be liquidated for a price, its actual liquidation depends on the price offered. Very often, management is far more willing to acknowledge the possibility of liquidating a credit risky position than to actually do it

3 Clearly, this assumption can be relaxed if the Markov migration probability matrix can be calculated for time horizons of higher frequency; the problem one might face is that this calculation may introduce more noise than explanatory power, due to the higher frequency of the data

4 For credit-risky exposures, such as the potential exposure from trading lines that can only be liquidated via other methods (such as credit guarantees or credit derivatives), one would have to provide a "cost of credit insurance" table or liquidation, again as a function of the exposure's maturity and credit rating

systematic risk factors, we then take a random draw on this recovery distribution to tabulate the exposure's loss in the event of default.

This technique implicitly assumes that the random recovery rates can be drawn independently from one another across different macroeconomic scenarios and counterparties. This assumption clearly breaks down, however, in two cases of interest:

☐ The first is when we are tabulating potential losses in the event of default arising from trading exposures. It is quite probable that the recovery rates for two client exposures are highly correlated, depending upon the compositions of their portfolios. For example, both might be corporate counterparties hedging liabilities with long-dated dollar interest rate swaps. The solution would again be to develop a model that jointly simulates the portfolio's potential exposure (driven by market rates) and credit events – ie, another "all singing, all dancing" model.

In addition to its inherent modelling complexity, however, this approach might also suffer from some concrete system constraints. To implement such an approach, all the counterparty's transactions, including netting and collateral information and valuation methods, would have to reside in a common system. While for many institutions this may be the ultimate objective, very few could currently implement such a system. More importantly, given the current state of financial disclosure, which provides little transparency regarding the hedge transactions concluded by a corporation, we may not know whether defaults and market rates are positively or negatively correlated for a particular counterparty.

☐ The second case is for certain asset classes, most notably mortgages and some collateralised loans, where the recovery rate is highly negatively correlated with the probability of default. Residential homeowners rarely default on their property when there is still much positive equity in the home. In these cases, it is sometimes useful to either add regional property price indexes to the set of explanatory variables used in the Logit modelling of our systematic risk indexes or to make the recovery distributions dependent upon the state of the economy.

In addition, our loss tabulation framework allows an institution easily to incorporate country or political risk. This is accomplished by associating each asset with a specific risk country and then simulating (correlated) country events. If a country event occurs, then all assets that bear that country's risk revert to their recovery distributions. ■

Thomas Wilson is a partner in the Zurich office of McKinsey & Company. He specialises in serving financial institutions in trading and credit risk management strategies, organisation and risk measurement

References

JP Morgan, 1997, CreditMetrics technical documentation

Kealhofer S, 1995i, *Managing default risk in portfolios of derivatives*, Derivative Credit Risk: Advances in Measurement and Management, Risk Publications, London

Kealhofer S, 1995ii, *Portfolio management of default risk*, proprietary documentation, KMV Corporation, San Francisco

Lawrence D, 1995, *Aggregating credit exposures: the simulation approach*, Derivative Credit Risk: Advances in Measurement and Management, Risk Publications, London

Moody's Investors Service, 1994, *Corporate bond defaults and default rates 1970–1993*, Moody's Investors Service

Rowe D, 1995, *Aggregating credit exposures: the primary risk source approach*, Derivative Credit Risk: Advances in Measurement and Management, Risk Publications, London

Wilson T, 1997, *Measuring and managing credit portfolio risk*, unpublished paper