PORTFOLIO CREDIT RISK (I)

In the first of two articles, **Thomas Wilson** describes a macroeconomic model of default rates

FINANCIAL INSTITUTIONS ARE increasingly measuring and managing the risk from credit exposures at a portfolio level, as well as the usual transaction level. This is happening for various reasons: the rise of portfolio credit risk management; the fall in profitability of traditional credit products; and increased opportunities (through credit derivatives) to manage the institution's exposure after it has been originated.

However, to exploit these opportunities, management must answer several questions. What is the risk of a given portfolio? How do different macroeconomic scenarios – at both regional and industry sector level – affect the portfolio's risk profile? What is the effect of changing the portfolio mix? How could risk-based pricing at the individual contract and portfolio levels be influenced by the level of expected losses and credit risk capital?

In this article, and its sequel next month, we describe a new, intuitive and practical method that can provide answers to these questions by tabulating the exact loss distribution from correlated credit events. The importance of tabulating the exact loss distribution is highlighted by the fact that counterparty defaults and rating migrations (upgrades and downgrades) cannot be predicted and are not perfectly correlated, implying that management faces a distribution of potential losses rather than a single potential loss (see figure 1). To make the definition of credit risk more precise in the context of loss distributions, the financial industry is converging on risk measures that summarise relevant aspects of the entire loss distribution. Two distributional statistics have become increasingly relevant for measuring credit risk: expected loss and some critical value of the loss distribution, often defined as the portfolio's credit risk capital. Each measure serves a distinct and useful role in supporting management decision-making and control.

Expected losses, illustrated as the mean of the distribution in figure 1, are often the basis for management's reserve policies: the higher the expected losses, the higher the reserves that must be set aside. As such, expected losses are also an important component for determining whether the pricing of the credit risk position is adequate. Usually, each transaction or netted credit portfolio should be priced with enough margin to cover its contribution to the portfolio's expected credit losses, as well as to other operating expenses.

Credit risk capital, defined as the maximum loss within a known confidence interval (eg, 99%) 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, to cover its unexpected credit losses. As it would be uneconomic to hold capital against *all* potential losses (this would imply that equity is held against 100% of all credit exposures), some level of capital must be chosen to support the portfolio of transactions in most, but not all, cases. Just as with expected losses, credit risk capital is also important for determining if the credit risk of a particular transaction is appropriately priced: typically, each transaction must be priced with enough margin to cover not only its expected losses but also the cost of its marginal risk capital contribution.

Most industry professionals split the challenge of credit risk measurement into two questions: what is the (joint) probability of a credit event occurring and what is the loss should it occur?

Measuring potential losses, given a credit event, is simple for many standard commercial banking products. For example, the exposure of a $100 million unsecured loan is roughly equal to $100 million, subject to any recoveries and discounting effects. For derivatives portfolios, committed but unutilised lines, collateralised transactions etc, it is more difficult. In this article, we focus on tabulating these loss distributions for an arbitrary portfolio, *given* the level of exposure. Readers interested in the complexities of exposure measurement for derivatives should consult JP Morgan (1997), Lawrence (1995) and Rowe (1995).

We will develop an approach here for measuring expected and unexpected losses that differs from other approaches in several important ways. First, it models the actual discrete loss distribution, and therefore the amount of risk capital required to support the portfolio, according to the number and size of credits, rather than using a normal distribution or mean-variance approximation. This is important because, in the case of an unbalanced, concentrated portfolio, the loss distribution is discrete and multinomial, rather than continuous and unimodal, and it is highly skewed rather than symmetric. In addition, its shape changes dramatically as other positions are added.

Because of this, the usual measure of unexpected losses, 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 tabulates the actual, discrete loss distribution for any given portfolio, thus also allowing explicit and accurate tabulation of a "large exposure premium" in terms of risk-adjusted capital for less diversified portfolios.

Second, the losses (or gains) are measured on a *discounted default basis* for credit exposures that *cannot* be liquidated (eg, most mid-market or retail loans and some corporate over-the-counter trading lines), as well as on a marked-to-market basis (recognising the economic impact of both credit migrations and defaults) for those that can be liquidated before the maximum maturity of the exposure (eg, traded corporate loans or debt securities). This implies that the model can be used to cover the credit exposures from liquid secondary market positions, as well as from illiquid commercial positions. Also, because we model both the average default and migration rates and correlations for single, rated counterparty names, as well as average loss rates on portfolios of non-rated retail counterparties, we can integrate retail exposures such as mortgage or credit card pools, which often make up the lion's share of a commercial bank's risk profile, into the same risk measurement framework.

The third difference is that the tabulated loss distributions are

1. Loss distribution (illustrative)

$100 portfolio, 250 equal and independent credits with prob(default) = 1%
Expected losses = –1.0; standard deviation = 0.63; credit risk capital = –1.8

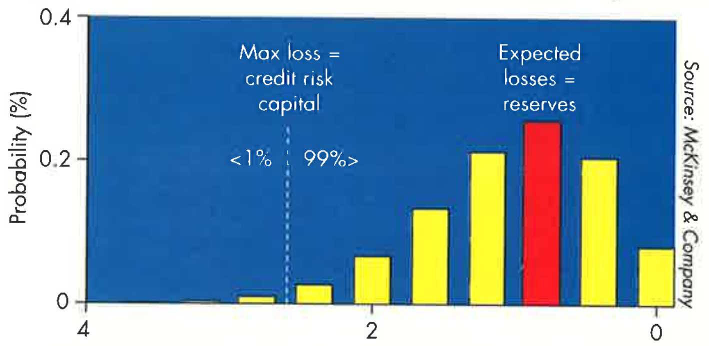

2. Actual v. predicted default rates in Germany

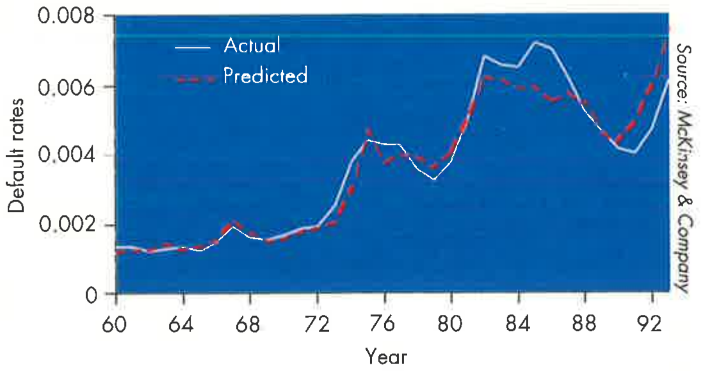

3. Germany's industrial average default rates: 1960–1992

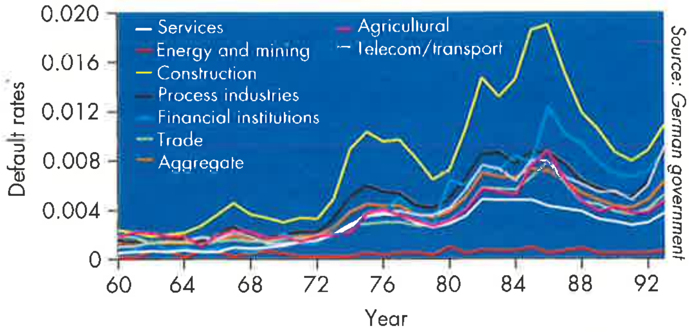

4. Default and migration probabilities

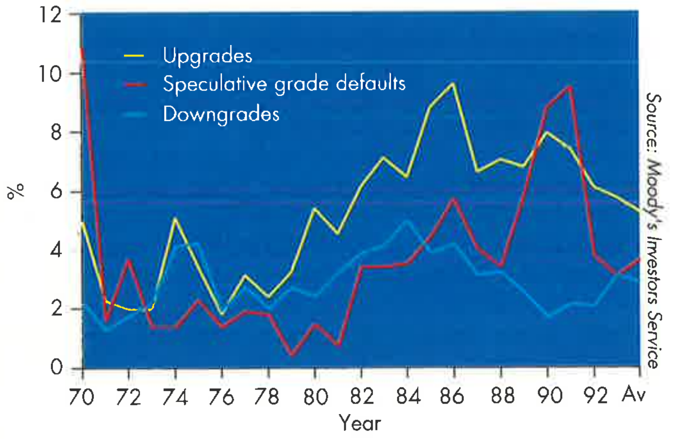

5. Total systematic risk explained

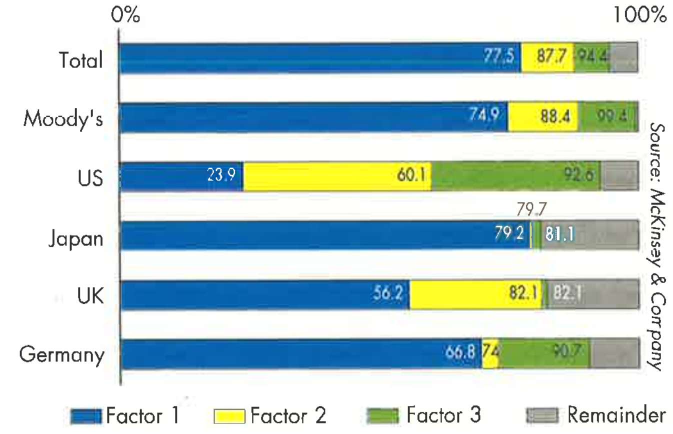

conditional on the current state of the economy, rather than being based on unconditional or 20-year averages that do not reflect the portfolio's true current risk. This allows us to capture the cyclical default effects that constitute the largest source of risk for diversified portfolios. Specific country and industry influences are recognised on the basis of empirical relationships, rather than assuming a constant correlation between counterparty segments or correlations determined by equity price movements, thereby allowing the model to mimic the actual, historical default correlations between corporates in different industries and regions as well as retail portfolios.

Some models, including many developed in-house, rely on one systematic risk factor to capture default correlations or, equivalently, on a constant correlation between all counterparties. Our approach is based on a true multifactor systematic risk model, which better captures reality. Finally, the loss distributions can be tabulated to incorporate not only random recovery rates, but also country events that can trigger technical default on an obligation, regardless of the counterparty's credit quality.

Intuition behind the model

In terms of modelling the joint default behaviour or systematic risk of a portfolio of counterparties over time, the model developed here leverages several intuitive observations on the behaviour and impact of individual, as well as aggregate, credit events. These observations are more complex than they might first appear; for a more complete justification and analysis, see Wilson (1997). Next month, we will show how to tabulate losses for an arbitrary portfolio of exposures based on these systematic risk indexes.

The first observation is that diversification helps to reduce loss uncertainty, all else being equal. Second, there is still substantial loss uncertainty – the systematic risk – for even the most diversified portfolios. Third, a portfolio's systematic risk is mostly driven by the "health" of the macroeconomy – in recessions, one expects defaults (and downgrades) to increase.

This relationship between average default rates and the state of the macroeconomy is shown in figure 2, which plots the actual default rate versus one predicted using only macroeconomic aggregates such as GDP growth and unemployment rates. The macroeconomic factors explain much of the overall variation in the average default rate series, as reflected by Logit regression equation R^2 (see equation (1)) of more than 90% for most of the countries investigated in Wilson (1997), including the US, Germany, Japan and France as well as domestic retail mortgage portfolios in many of the same countries. Similar analysis also shows that different sectors of the economy react differently to macroeconomic shocks and cycles. For example, the intuitive expectation that the construction sector would be more adversely affected during a recession than other sectors is supported by data for a wide variety of different countries analysed in Wilson (1997).

More specifically, we model speculative default rates by estimating the parameters of a Logit function where the dependent variable is the probability of default for speculative grade counterparties and the independent variable is a segment-specific index depending on current macroeconomic variables, such as:

$$p_{j,t} = 1 / \left[1 + \exp\left(y_{j,t}\right)\right] \tag{1}$$

where $p_{j,t}$ is the probability of default for a speculative grade counterparty in segment j at time t, and $y_{j,t}$ is the segment-specific macroeconomic index, whose parameters must be estimated. The Logit functional form was chosen over linear and exponential representations for two reasons: first, because it offered, on average, a better fit (as measured by R^2) and second, because, for any value of the index y, it yields a probability within the interval [0,1]. The latter is important not only because it is a desirable property for a probability, but also because we will be simulating these index values over a multi-year horizon and extreme economic cycles using Monte Carlo methods: without this, we cannot guarantee the resulting probability of default will be between [0,1].

The macroeconomic index, which measures the health of the country's macroeconomy, is determined by such variables as overall unemployment, GDP growth rate, the rate of government spending, regional housing price indexes and unemployment rates for mortgage portfolios and so on. More specifically, the index takes the following form:

$$y_{j,t} = \beta_{j,0} + \beta_{j,1}X_{j,1,t} + \beta_{j,2}X_{j,2,t} + \beta_{j,3}X_{j,3,t} + \upsilon_{j,t} \qquad (2)$$

where $y_{j,t}$ is the index value specific for the j^{th} segment at time t, $\beta_j = (\beta_{j,0}, \beta_{j,1}, \beta_{j,2}, \beta_{j,3})$ is a set of regression coefficients to be estimated for the j^{th} segment, $X_{j,t} = (X_{j,1,t}, X_{j,2,t}, X_{j,3,t})$ is the set of explanatory variables at time t for the j^{th} segment (eg, GDP growth rates and unemployment) and $\upsilon_{j,t}$ is a random variable assumed to be independent and identically normally distributed, for example:

$$\upsilon_{j,t} \sim N(0, \sigma_j) \text{ or } \upsilon_t \sim N(0, \Sigma_\upsilon) \qquad (3)$$

where υ is the vector of stacked index innovations and Σ_υ is the $j \times j$ covariance matrix of the index innovations. This Logit formulation can be viewed as a multifactor model for determining segment-specific average speculative default rates: the systematic risk component is captured by the (weighted) influence of the macroeconomic variables with a segment-specific "surprise" captured by the error term υ. To make the model segment-specific, we estimate individual Logit functions for each country/industry/retail pool segment whenever data is available[1], allowing the explanatory variables to differ between segments.

The fourth observation is that these relationships hold not only for each country but also for different segments within each country. Figure 3 plots the average default rates by segment for Germany, all of which move very strongly with the overall German credit cycle but with different amplitudes, construction being the most volatile (or highest "beta") industry and energy and mining being the least volatile (or lowest "beta") industry. Similar, strong regression results also hold for these industries.

The final observation is also both intuitive and empirically verifiable: rating migrations are linked to the macroeconomy – both default and credit downgrades are more likely in a recession. Our starting point is an unconditional Markov transition matrix (which we denote (ϕM)), calculated using rating agency data, own experiences or JP Morgan's estimated transition matrix (JP Morgan, 1997). This matrix gives the average probability of migrating from one credit class to another within a year and is estimated using several years' data across many countries and industries (see table A). The matrix is therefore unconditional: the probabilities, like the ratings, will vary with the economic cycle.

A. Unconditional credit migration probabilities

Initial rating	Final rating AAA	AA	A	BBB	BB	B	CCC	Default
AAA	0.894	0.098	0.006	0.002	0.001	0.000	0.000	0.000
AA	0.009	0.909	0.071	0.008	0.001	0.002	0.000	0.000
A	0.001	0.026	0.901	0.060	0.008	0.004	0.000	0.001
BBB	0.001	0.003	0.063	0.851	0.063	0.015	0.002	0.003
BB	0.000	0.002	0.006	0.074	0.789	0.102	0.012	0.015
B	0.000	0.001	0.004	0.006	0.061	0.830	0.038	0.061
CCC	0.002	0.000	0.002	0.010	0.015	0.120	0.660	0.191
Default	0.000	0.000	0.000	0.000	0.000	0.000	0.000	1.000

Source: Moody's Investors Service

Figure 4 plots rating migrations against average speculative default rates – as we have shown, a good proxy for the economy's health. One important observation from this figure is that rating migrations are not only volatile, but are also correlated with changes in the average speculative default rate: if the speculative default rate is higher than average, then downward migrations increase and upward migrations decrease; these observations are confirmed by the correlations between the two.

These empirical relationships between the speculative default rates and the credit migration, as well as default rates for other rating classes discussed earlier, can be used to make the *un*conditional Markov rating transition matrix *conditional* on the current speculative default rate (see Wilson (1997) for more details).

Intuitively, if the ratio of the actual speculative default rate to its unconditional mean ($SDP_t/\phi SDP$) equals one, then the conditional transition matrix is equal to the unconditional matrix. If this ratio is greater than one (ie, more defaults than average), then more of the probability mass is shifted into downgraded and default states; and vice versa if the ratio is less than one.[2]

Based on these conditional one-year Markov transition matrices, we can calculate the cumulative conditional probabilities of migrating from one rating class to another over any (yearly) time frame using the following formula:

$$M_t = \Pi_{i=1..t} M(SDP_t / \phi SDP) \qquad (4)$$

where M_t represents the t-year conditional cumulative rating distribution for a given future path of the speculative default rates, $M(SDP_t/\phi SDP)$ is the conditional one-year rating transition matrix dependent upon the speculative default rate, SDP_t is the realised speculative default rate at time t, ϕSDP is the average speculative default rate and $\Pi_{i=1..t}$ is the multiplication operator.

In terms of our modelling strategy, the multiple factors driving our multifactor model of systematic credit risk are the correlated macroeconomic variables (eg, GDP growth, unemployment rates) for the various countries in our sample, as well as industry-specific "shocks". Figure 5 shows the need for a multifactor, rather than a single-factor model, for systematic risk. A factor analysis of country-average default rates – a good surrogate for systematic risk by country – shows that the first "factor" captures only 77.5% of the total variation in systematic default rates for Moody's, the US, UK, 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 assumed to be independent and uncorrelated. Unfortunately, the first factor only

[1] Typically, data comes from own experiences or public statistics. In addition, one of the rating agencies will make such data commercially available soon

[2] An alternative starting point may have been to model the one-year probability of default for each rating class rather than the one-year Markov rating transition matrix. But we chose the latter approach because we are interested in measuring the credit risk of multi-year engagements; for these, credit migrations leading to later defaults can be just as important as immediate defaults, implying that modelling one-year default rates would have been inappropriate

6. Simulations of speculative rating default probability for Germany

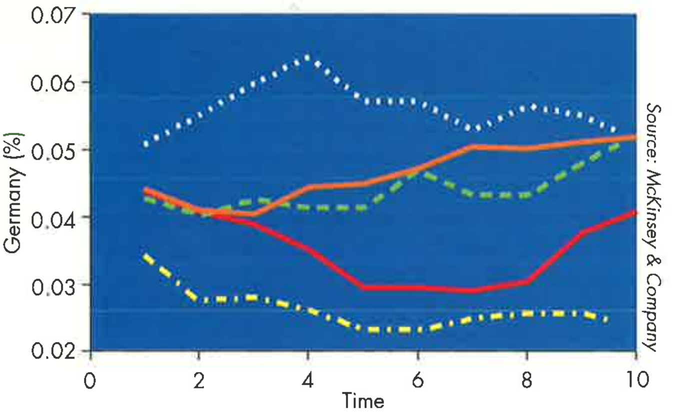

explains 23.9% of the US systematic risk index, 56.2% of the UK's and 66.8% of Germany's. The substantial correlation remaining is explained by the second and third factors, covering an extra 10.2% and 6.8% respectively of the total variation and the bulk of the risk for the US, the UK and Germany. This shows that a single-factor systematic risk model, eg, one based on asset betas or agency data alone, is not enough to capture all correlations.

The multiple factors driving our multifactor model of systematic credit risk are, in fact, the correlated macroeconomic variables for the various countries in our sample, as well as industry-specific "shocks".

Multifactor systemic risk model

Using these, as well as other historical relationships, we can construct a multifactor model for the evolution of default and credit migration probabilities. We do this in three steps: first, describing the evolution of the global macroeconomy; second, mapping this into country/industry-specific one-year speculative default rates or default indexes; finally, mapping these into country and industry-specific cumulative rating migration probabilities.

Defining the state of the economy

At the beginning of any period t, the economy's health is drawn from a (conditional) distribution. Using our Logit models for determining the country/industry-specific speculative default rates, the economy's condition is defined by the various macroeconomic explanatory variables discussed earlier (eg, GDP innovations, changes in housing price indexes, etc). For simplicity, we use a set of univariate, auto-regressive equations of order two (AR(2)) to model the development of the individual macroeconomic time series describing the economy's health.[3] More precisely, we assume that the evolution of each of the macroeconomic series is governed by:

$$X_{j,i,t} = k_{i,0} + k_{i,1}X_{j,i,t-1} + k_{i,2}X_{j,i,t-2} + \varepsilon_{j,i,t} \qquad (5)$$

[3] In practice, we actually use univariate ARIMA(p,d,q) processes. While it is clear that a better modelling strategy might have been to estimate a vector autoregressive moving average (VARMA(p,q)) model or even a multi-country structural equation model, we chose the independently estimated ARIMA representation because of its simplicity. Further work on improving the macroeconomic models should be undertaken

where $X_{j,i,t}$ is the value of the i[th] macroeconomic variable in the j[th] segment at time t, $k_{i,j}$ (where j = 1..3) are the three constants to be estimated for each of the i macroeconomic variables and $\varepsilon_{j,i,t}$ is the error term, assumed to be distributed $N(0,\sigma_i)$.

The constant parameters (k) are estimated using historical data, ie, data for which $t < 0$. Once the constant parameters have been estimated, the $\varepsilon_{j,i,t}$ terms for $t \geq 0$ can be considered as the forecast errors or "surprises" for the t-step ahead forecast and are the variables which need to be simulated. As the forecast "surprises" are correlated (eg, a surprisingly poor showing for GDP growth will most likely be associated with an unpleasant surprise in terms of unemployment rates), we also need to estimate the joint covariance matrix for the forecast error terms to simulate the joint development of the different macroeconomic variables. More specifically, based on the assumptions made earlier, it follows that:

$$\varepsilon \sim N(0, \Sigma_\varepsilon) \qquad (6)$$

where ε is the stacked vector of errors from each of the i AR(2) equations and Σ_ε is the covariance matrix of ε.

Defining segment default indexes

Based on this state of the economy, we determine the average speculative default rate for each country/industry segment. The macroeconomic forecast and simulated forecast error terms are used to drive the Logit functions described in equations (1) and (2); these, in turn, are used to construct the t-step ahead cumulative rating distributions using the empirical regularities specified in equation (4).

Combining the earlier equations, we derive the speculative default rates for all of the j country/industry segments, given a particular realisation of the various macroeconomic innovations and innovations to the Logit functions. More specifically, we define a system of equations governing the joint evolution of the country/industry-specific speculative default rates and associated macroeconomic variables:

$$p_j(y_{j,t}) = 1/\left[1+\exp(y_{j,t})\right]$$

$$y_{j,t} = \beta_{j,0} + \beta_{j,1}X_{j,1,t} + \beta_{j,2}X_{j,2,t} + \beta_{j,3}X_{j,3,t} + \upsilon_{j,t}$$

$$X_{j,i,t} = k_{i,0} + k_{i,1}X_{j,i,t-1} + k_{i,2}X_{j,i,t-2} + \varepsilon_{j,i,t}$$

$$E = \begin{pmatrix}\upsilon \\ \varepsilon\end{pmatrix} \sim N(0,\Sigma), \qquad \Sigma \equiv \begin{bmatrix}\Sigma_\upsilon & \Sigma_{\upsilon,\varepsilon} \\ \Sigma_{\varepsilon,\upsilon} & \Sigma_\varepsilon\end{bmatrix}$$

where E is the $(j+i)\times 1$ vector of innovations affecting the system of equations given above and Σ is the $(j+i)\times(j+i)$ covariance matrix of macroeconomic variable forecast errors (υ), and segment-specific speculative default rate "shocks" (ε).

Defining A as the $n\times n$ Cholesky decomposition of Σ so that $\Sigma = AA'$, we can now simulate the joint speculative default rates across all segments over some time horizon, T. This is done by:

- drawing a sequence of realisations for z_t, t = 1 to T, of (j+i) N(0,I) random variables, where I is the $(j+i)\times(j+i)$ identity matrix;
- calculating the realisations E_t, incorporating the correlations between

7. Simulated default probabilities

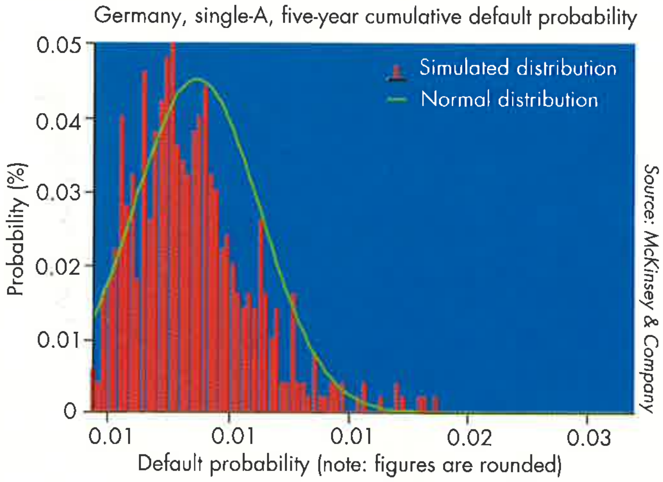

the various macroeconomic and segment specific default innovations, using the relationship $E_t = A^*z_t$; and

☐ calculating P(t) using the above system of equations.

Figure 6 gives an example of five simulations of the speculative default rates for Germany over a 10-year horizon. For example, we can interpret the uppermost simulation as representing an immediate recession followed by a gradual recovery and the lowermost simulation run as being an immediate and sustained improvement in the global economy.

Tabulating the cumulative migration distributions

Once the speculative default rates for each country/industry segment have been simulated, we then calculate the unique Markov transition matrix for each segment for any time horizon by using equation (4). Using the relationships outlined in equation (4) and the simulated speculative default rates, we can calculate the rating distribution for any initial rating at different points in time in the future, conditional on the simulated macroeconomic cycle over that time horizon. Figure 7 shows the histogram of simulated cumulative default probabilities for a five-year, single-A counterparty in Germany, based on 500 simulations.

Similar histograms can be tabulated for both cumulative default probabilities and cumulative credit rating migration probabilities for any maturity, rating class, country and industry combination. Interestingly, the histogram in figure 7 does not look "normal": it is truncated at zero (a natural phenomenon, as probabilities should not in general be negative) and, because of this, it is skewed. In general, this phenomenon occurs most with higher rated, shorter maturity exposures, which tend on average to have lower cumulative default probabilities. The distributions for lower rated, longer maturity portfolios tend to be more symmetric in shape (since the mean is quite a bit higher).

Systematic risks

Given the simulated default probability distribution for each of the segments, by rating, we can tabulate several interesting distribution parameters such as the expected conditional default probabili-

B. Cumulative conditional default probability

Maturity 1	2	3	4	5	6	7	8	9	10	Rating
0.000	0.008	0.035	0.093	0.193	0.345	0.562	0.863	1.252	1.741	AAA
0.000	0.055	0.178	0.382	0.684	1.094	1.628	2.315	3.143	4.122	AA
0.121	0.364	0.780	1.385	2.206	3.241	4.493	5.999	7.697	9.579	A
0.575	1.610	3.153	5.121	7.483	10.135	13.013	16.148	19.367	22.633	BBB
2.872	7.087	12.160	17.510	22.941	28.209	33.227	38.094	42.599	46.750	BB
11.798	22.170	31.473	39.438	46.340	52.253	57.333	61.858	65.755	69.120	B
26.999	44.020	55.545	63.445	69.169	73.446	76.760	79.494	81.726	83.582	CCC
100	100	100	100	100	100	100	100	100	100	D

Source: McKinsey & Company

C. Cumulative excess loss probability

Maturity 1	2	3	4	5	6	7	8	9	10	Rating
0.00	0.01	0.02	0.06	0.12	0.24	0.41	0.63	0.83	1.08	AAA
0.00	0.04	0.10	0.22	0.36	0.64	1.00	1.43	1.77	2.17	AA
0.05	0.17	0.36	0.68	1.03	1.67	2.40	3.19	3.72	4.30	A
0.25	0.74	1.35	2.24	3.05	4.39	5.68	6.85	7.38	7.91	BBB
1.26	3.06	4.57	6.29	7.35	9.05	10.27	11.05	10.92	10.79	BB
5.16	7.79	8.84	9.87	9.97	10.67	10.88	10.72	9.96	9.31	B
7.07	8.37	7.94	7.75	7.11	7.03	6.78	6.42	5.83	5.36	CCC
0	0	0	0	0	0	0	0	0	0	D

Source: McKinsey & Company

ty and maximum possible default probability for each segment, by rating. Table B shows the conditional expected cumulative default rates for the German index portfolio for exposures of any maturity up to 10 years and any initial rating from triple-A to default, conditional on the macroeconomic environment in Germany in 1995. It is interesting to note that if the *long run average* levels of the macroeconomic explanatory variables are used for the initial conditions for simulating the AR(2) processes given in equation (5), then the cumulative default probabilities given in table B would be *equal* to the unconditional cumulative default probabilities implied by the migration matrix given in table A. Thus, our approach explicitly allows one the choice of whether to recognise credit cycles explicitly.

Table B *could* represent the conditional expected portfolio loss rates if, and only if, four conditions are met: the portfolio is diversified to the extent that it mimics the German index portfolio (eg, a large number of small German exposures in all industries); recovery rates are zero; the exposure for each individual counterparty is constant over its entire life; and the losses are treated on a non-discounted basis. As these assumptions are restrictive in practice for an arbitrary portfolio of credit exposures, we present the cumulative expected default probabilities here only to illustrate the properties of the systematic risk model; we will relax each of these restrictive assumptions next month, demonstrating then how to tabulate the actual loss distributions for any arbitrary portfolio.

According to this table, the average conditional cumulative default rates for a three-year, single-A diversified German index position is 0.780, conditional on the current macroeconomic environment. The last row of this table is equal to 100 because it represents counterparties that are currently in default and the model assumes that default is an absorbing state. In a similar manner, we can also tabulate other critical values of the index loss distribution.

In table C, we have tabulated the additional probability of default within a 99% confidence interval, above and beyond the expected default probability, for the German index portfolio conditional on the current macroeconomic environment. As with the previous exhibit, this table *could* be interpreted as the credit risk capital for a portfolio of German exposures if and only if the portfolio mimicked the German index, had no recovery potential in the case of default, had constant exposures over the life of the commitments and if the losses were to be treated on a non-discounted basis. Again, we will demonstrate how to calculate the credit risk capital from an arbitrary portfolio of exposures in part two of this article.

Tables B and C show two interesting and connected phenomena: the excess maximum default probability above the expected probability of default do not increase monotonically by rating category or by maturity as one might expect given the close relationship between maximum possible default rates and credit risk capital. These observations are reflected most clearly in the declining cumulative excess loss probabilities for the triple-C rating category for maturities longer than two years and for the profile by rating for eight-year exposures, which peaks at the double-B class and declines thereafter.

D. Total risk

Maturity 1	2	3	4	5	6	7	8	9	10	Rating
0.00	0.01	0.06	0.15	0.31	0.58	0.97	1.49	2.08	2.82	AAA
0.00	0.09	0.28	0.60	1.05	1.74	2.63	3.75	4.91	6.29	AA
0.17	0.54	1.14	2.06	3.23	4.91	6.89	9.19	11.41	13.87	A
0.83	2.35	4.50	7.37	10.54	14.53	18.69	23.00	26.75	30.54	BBB
4.13	10.15	16.73	23.80	30.29	37.25	43.50	49.14	53.52	57.54	BB
16.96	29.96	40.31	49.31	56.31	62.92	68.21	72.58	75.71	78.43	B
34.07	52.39	63.49	71.19	76.28	80.48	83.54	85.91	87.55	88.94	CCC
100	100	100	100	100	100	100	100	100	100	D

Source: McKinsey & Company

The explanation for this phenomenon is straightforward: for lower rated exposures, the uncertainty of loss relative to the expected loss actually decreases with maturity after some point, reflecting the fact that more of the exposure will be lost with increasing certainty as the maturity is extended. This interesting effect disappears, however, if one focuses on the total maximum risk, defined as the total maximum default probability within a 99% confidence interval, which, like expected losses, increases monotonically in maturity for all rating classes (see table D). ■

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 will appear in next month's issue