IRB approach explained

At the end of this month, the consultation period for the new Basel Accord on bank capital will end. We have prepared a technical section this month devoted to various issues surrounding Basel II. In the first paper, Tom Wilde sheds light on the assumptions and parameters underlying the internal ratings-based approach

Beyond the contentious multipliers and add-ons present in the internal ratings-based approach (IRB, Basel, 2001), lie equations that have generally been welcomed. The Basel Committee on Banking Supervision has been applauded for bringing bank capital more into line with credit risk modelling theory. However, the IRB calculations do lack some transparency. Here, I show how the base risk weights of the IRB approach are determined and discuss the granularity adjustment, which fits into the same framework. The interested reader should also refer to Gordy (2000b), which reports the actual modelling work on which the IRB approach is based.

A brief history of credit risk modelling

The context of the IRB approach is the theory of credit risk modelling. In recent years, commercial or publicly available models of credit risk have appeared, notably KMV's CreditMonitor, JP Morgan's[1] CreditMetrics, Credit Suisse First Boston's CreditRisk+ and McKinsey's CreditPortfolioView. The ideas behind these models are due in part to earlier authors, notably Merton (see, eg, Merton, 1974) and Vasicek (1987).

Hickman & Koyluoglu (1998) pointed out that the models all derive from a common framework in which part of each model deals with systematic risk, and a "second stage" assesses the additional unsystematic component. Systematic risk has the same meaning as in the capital asset pricing model, being the sensitivity of the overall condition of the portfolio to the economy, which cannot be diversified away. Hickman & Koyluoglu also showed that the different choices made in the publicly available models can broadly be compensated for by appropriate parameterisation (see also Gordy, 2000a, which contains an in-depth analysis of CreditMetrics and CreditRisk+).

Risk managers, regulators, consultants and academics have now all contributed to this theory so there is now an accepted general framework for measuring credit risk. Finger (1999), Belkin, Suchower & Forest (1998) and Lucas *et al* (1999) are some further important contributions to the theory.

The common framework

We summarise the common framework from Hickman & Koyluoglu (1998). Suppose there are systematic factors, or drivers of default, represented mathematically by random variables $X_1, \ldots, X_n$, such that the probability of default of any obligor A in the portfolio is given by some specified function[2]:

$$\text{Probability of default} = P_A(X_1, \ldots, X_n)$$

The variables $X_1, \ldots, X_n$ "represent" systematic factors that affect default rates, and which can be real (such as interest rates, share price indexes or macroeconomic data) or formal (such as factors derived from principal component analysis). For a given model, the functional dependence $P_A(X_1, \ldots, X_n)$ typically has a fixed functional form, with additional parameters specified for each obligor reflecting their intrinsic credit quality and sensitivity to the variables. The specification of this function and of the variables $X_1, \ldots, X_n$ is the source of most of the differences between the commercially available credit risk models.

Default probabilities therefore depend on factors that can also affect other obligors, and in this way systematic risk is modelled. Normally, this is the main risk inherent in the portfolio. However, the systematic factors only affect the probability that an obligor will default; there is still an element of chance about the actual number of defaults or the amount of loss. This element of pure luck is the source of unsystematic risk. Unlike systematic risk, it will only be important for a portfolio that is small or has large exposures.

If $n = 1$, the model is called a one-factor model. Only one random variable models the systematic element of uncertainty in the portfolio. The one-factor models are the simplified relatives of the commercial models, and are the models used to set the IRB risk weights. Specifically, the simplified form of CreditMetrics is used to calculate base risk weights, and CreditRisk+ assists in the calculation of the granularity adjustment.

By using one-factor models, there is no need to use a commercial model directly, and the parameterisation issues for the more detailed models are avoided as much as possible. However, there is a less obvious additional advantage: one-factor models give rise to additive capital requirements. This property is important because it is the reason why the risk weights in the IRB approach do not have to be calibrated with reference to any particular test portfolio; subject to the modelling assumptions, they are reasonably valid for all portfolios.

Before using a credit risk model, the Basel Committee has needed to set parameters whose values must be determined by other considerations than modelling alone.[3] Thus capital for credit risk will be calculated at 99.5% confidence over a one-year time horizon, and including expected loss. The concepts presented below do not depend in any way on these particular parameters but we will find it convenient to use them to avoid cumbersome "over-general" notation.

A risk or capital weight is a factor to be applied to, say, loans, so that the total across all loans adds up to the capital required for the whole portfolio, ie:

$$\text{Capital requirement} = \Sigma(\text{Exposure} \times \text{capital weight}) \qquad (1)$$

or, equivalently, in the language of risk weights (which are capital requirements scaled up by 12.5):

$$\text{Capital requirement} = \Sigma(\text{Exposure} \times \text{risk weight}) \times 8\% \qquad (2)$$

The problem is that capital was not defined as a sum of weights, but as the result of a portfolio calculation. Can total capital, in fact, be expressed as suggested by equation (1) above? The weights used would have to be valid across a wide range of portfolios.[4] In short, the Committee must reconcile the following two objectives:

- capital requirements should add up, to a reasonable approximation, to

[1] *Now RiskMetrics Group*
[2] *We use* P *for the random variable representing default rates and* p *for the average of* P *over values of the variables. In most cases, the form of* P *is chosen so that its average* p *is an explicit parameter to be input*
[3] *It is valid to debate the choices they have made but that is not the purpose of this article*
[4] *Calculations using an "example portfolio" do not establish that certain weights are appropriate or not. What is needed is a general theory that shows certain weights will indeed hold in general*

99.5% percentile loss

☐ capital weights must depend only on properties of the individual transaction, as in equations (1) and (2).

However, these objectives are not contradictory if a one-factor model is used. This is the key advantage of one-factor models.

But to take advantage of the additive behaviour of one-factor models, capital must be calculated by looking only at systematic risk. The unsystematic component, which is normally much smaller, is not tractable in the same way. The granularity adjustment deals with this component.

Systematic risk and its contributions

Normally, in finance, systematic risk is measured via covariance ("betas" in the classic portfolio theory), but the framework described above makes available a whole statistical distribution analogous to the loss distribution but measuring only systematic risk.

Assume we are modelling the loss distribution from a portfolio Π, using the common framework and the notation above. Let E_A be the average loss given default (for simplicity, we take exposure to mean loss amount given default, so that average recovery is already factored in) and $\sigma(E_A)$ be its proportional standard deviation (the standard deviation in dollars is then $E_A\sigma(E_A)$). Now suppose we know the values of the systematic variables $X_1, \ldots, X_n$. This is like knowing the state of the economy. Then the average loss from the portfolio, and its variance[5],are:

$$\mu = \sum_A E_A P_A\left(X_1, \ldots, X_n\right) \tag{3}$$

$$\sigma^2 = \sum_A E_A^2\left(P_A\left(1 - P_A\right) + P_A\sigma^2\left(E_A\right)\right) \tag{4}$$

The actual loss is drawn from a statistical distribution with this mean and variance. Systematic risk distribution means that, if the portfolio is sufficiently fine grained, the variance conditional on $X_1, \ldots, X_n$ will be small, and so (eg, by Chebyshev's inequality) the conditional distribution of losses will just be the conditional average loss "with certainty". The uncertainty remaining then arises from the fact that we do not know the values of the variables $X_1, \ldots, X_n$. The distribution of losses is the distribution of the conditional mean loss in equation (3). This random variable is called the systematic loss of the portfolio Π.

From the original portfolio Π construct a portfolio Π_m by replacing each A with m new obligors $A_1, \ldots, A_m$ each having exposure E_A/m but the same default probability $P_A(X_1, \ldots, X_n)$. As m tends to infinity, the portfolios form a sequence that we think of as having a limit Π_∞ called the systematic portfolio. Π_∞ is what the Basel Committee calls an "infinitely fine grained" portfolio. We can check that the conditional variance (4) of Π gets divided by m in this process, while the mean (3) is unaffected. So the notional portfolio Π_∞ has as its total loss distribution just the systematic distribution (3), ie, the random variable μ defined by (3) is the systematic loss distribution of Π.

For any statistic for the original portfolio Π, there is the corresponding equivalent systematic statistic, which is the statistic evaluated on Π_∞, or equivalently on the systematic loss distribution (3). Of particular importance are the systematic percentiles, such as the systematic loss 99.5% percentile, which is defined as the 99.5% confidence point on the systematic loss distribution. These are the quantities used to define capital in the IRB approach.

Putting equation (1) another way, the capital weight should be the marginal contribution made to the 99.5% point on the loss distribution by each exposure (think of equation (1) for the original portfolio, and for the portfolio with one exposure removed). Analogously, the systematic risk contribution is the contribution made to the systematic 99.5% point. We write RC_A and SRC_A for these two risk contributions. The systematic loss distribution is given by (3) as the distribution of μ. This looks promisingly like a sum of random variables, one for each obligor. Suppose now that we are using a one-factor model, ie, $n = 1$. Then the systematic loss distribution is just:

$$\sum_A E_A P_A\left(X\right)$$

Now provided this is an increasing function of X (in practice each $P_A(X)$ will be increasing by suitable choice of X, so this condition is not serious), the 99.5th percentile of this distribution is just:

$$\sum_A E_A P_A\left(X_{99.5\%}\right)$$

and therefore the difference made to this by obligor A is just:

$$SRC_A = E_A P_A\left(X_{99.5\%}\right)$$

We have achieved our goal of a model in which systematic risk contributions depend only on the properties of an individual obligor.

The IRB risk weights

We apply the theory above to the IRB risk weights. These were derived using the one-factor form of CreditMetrics. In this model, the systematic factor X is a standard normal random variable (the systematic variable has a "real" interpretation as minus the normalised systematic component of asset return for each obligor).[6] Using $N()$ to denote the cumulative normal density function, the functional dependence is given by the well-known Vasicek formula[7]:

$$P_A = N\left(\frac{N^{-1}\left(p_A\right) + X\rho_A^{1/2}}{\left(1 - \rho_A\right)^{1/2}}\right) \tag{5}$$

where ρ_A is the "asset R – squared" for obligor A and p_A is the unconditional default probability (the formula does not make it obvious that p_A is the average of P_A but this is clear from the derivation of the Vasicek formula).[8] The systematic risk contribution at 99.5% confidence is derived by setting X equal to its 99.5% value 2.576, and multiplying by the exposure. Thus:

$$SRC_A = E_A N\left(\frac{N^{-1}\left(p_A\right) + 2.576\rho_A^{1/2}}{\left(1 - \rho_A\right)^{1/2}}\right) \tag{6}$$

Using the Committee's choice of $\rho_A = \rho = 20\%$ (Basel, 2001, paragraph 172), we can confirm the IRB risk weights for corporates directly. We calculate:

$$\left(1 - \rho_A\right)^{-1/2} = 0.8^{-1/2} = 1.118 \text{ and } 2.576\,\rho_A^{1/2} / \left(1 - \rho_A\right)^{1/2} = 1.288$$

These are the coefficients in the middle term in Basel (2001, paragraph 171). The retail risk weights (Basel, 2001, paragraph 310) can be derived similarly with $\rho = 8\%$.

The calibration factor

The full IRB risk weight is given as:

$$976.5 \times N\left(1.118N^{-1}\left(p_A\right) + 1.288\right) \times \left(\text{Maturity adjustment}\right)$$

(see Basel, 2001, paragraph 171). The middle term is equation (6). The last term is an adjustment representing the additional capital required against fair-value changes based on a loan of three years' maturity. Its form is derived empirically (the Committee refer to "judgemental pooling of information") and we do not discuss it here, except to note that it becomes one for a loan of explicit one-year maturity (to see this, calculate the risk weight in Basel, 2001, paragraph 159, for an asset with maturity $M = 1$. The third term of the present base risk weight cancels with the maturity adjustment.)

The first term is a "calibration adjustment". This number is mostly due to conventions about the way risk weights are stated, so to see the real

5 The variance formula is obtained by directly applying the definition of variance for a single obligor A, with loss equal to the (uncertain) loss given default in the event of default, or zero otherwise. The portfolio variance is just the sum across obligors because they are independent conditional on the X_{1-n}

6 Not quite lognormal assets, as apparently claimed in Basel (2001, paragraph 172) but the point is minor

7 See, eg, Hickman & Koyluoglu (1998) or Finger (1999)

8 Thus p is the unconditional chance of default over the next year, including good and bad economic outcomes. This can be compared with guidance from the New Accord on default probabilities. A "long run average" (Basel, 2001, paragraph 217) roughly corresponds to observing p_A from a time series of historic outcomes for obligors starting out in the same condition as A (these are realisations of P, not p), but p is arguably not "forward looking" as required by paragraph 218 as it depends only on the spot characteristics of A

scaling factor we must disentangle these. What factor would give a "pure modelling" result, ie, just amount to equation (6)? For an asset with loss given default equal to 50%, the capital requirement factor free of any scaling should be just $50\% \times N(1.118N^{-1}(p_A) + 1.288)$. To achieve this effect, the base risk weight should be set at:

$$100/0.08 \times 50\% \times N\left(1.118N^{-1}(p_A)+1.288\right)$$

We have $100/0.08 \times 50\% = 625$, so the Basel Committee seems to have employed a scaling factor of $976.5/625 = 1.56$ approximately, but that is not quite right. Later, 4% of the capital calculated here is rebated against the granularity adjustment (see Basel, 2001, paragraph 432). There is no free lunch, however. This same 4% is first produced by grossing up the baseline charge (Basel, 2001, paragraph 171) by $1/0.96$, ie, so that only 96% of the number calculated covers all the base risk, leaving 4% free to cover granularity risk. Thus the actual multiplier is:

$$976.5/625 \times 0.96 = 149.99\%$$

The Committee does not mention this multiplier explicitly, preferring to present the calibration in terms of the capital requirement for a specific asset (Basel, 2001, paragraph 172). This may reflect their original approach, but one can assume they have also rounded up to arrive at 1.5 exactly.

The granularity adjustment

We have shown how the base risk weights quantify systematic risk exactly according to the credit risk modelling framework introduced above. But the question of unsystematic risk remains. Put another way, the risk weights add up to the 99.5% point of the loss distribution of Π_∞, not Π itself. In general, this should be expected to be close, but an underestimate of total risk, and possibly a material understatement for a small portfolio or one with large exposures. Figures 1 and 2 show actual and systematic loss distributions for a typical large and small portfolio. The distributions are similar for the large but very different for the small portfolio.

The task to be performed by the granularity adjustment (Basel, 2001, chapter 8) is to adjust the 99.5% point on the systematic loss distribution to try to achieve the corresponding point of the actual loss distribution.

The adjustment (before rebating 4% of the base risk weight charge already surcharged within the base risk weight) is given as $GA = GSF/n^*$ where n^* is the reciprocal of the "Herfindahl index". The granularity scaling factor GSF is given in Basel (2001, paragraph 457) by:

$$GSF = \left(0.6 + 1.8LGD_{AG}\right)\left(9.5 + 13.75PD_{AG}/F_{AG}\right)$$

and is a scaled version of the formula given in paragraph 456:

$$\beta = \left(0.4 + 1.2LGD_{AG}\right)\left(0.76 + 1.1PD_{AG}/F_{AG}\right) \tag{7}$$

where the scaling factor is 18.75 (= $1.5/0.08$, arising from the transition to assets from capital, and the further multiplier of 1.5 consistent with the factor already applied to the base risk weights).

Unlike the base risk weights, the granularity adjustment is based on CreditRisk+ with one factor. Here, default rates are parameterised depending on a gamma distributed random variable X, with the following relation replacing (5):

$$P_A = p_A\left(1 - \omega_A + \omega_A X\right)$$

Here X has mean 1 and variance $\sigma(X)$, and the parameter ω_A is called the "factor loading". It describes how X affects the default rate of A, where $0 \le \omega_A \le 1$. The systematic risk contribution for obligor A is therefore $SRC_A = E_A p_A(1 - \omega_A + \omega_A X_{99.5\%})$ in place of (6).

Since two models are being used, they must be calibrated somehow. The Committee recognises that this is not a perfect situation but chooses the best place to stitch the models together, namely at the 99.5% confidence level (Basel, 2001, paragraph 446). Thus the weight ω_A is chosen so that CreditRisk+ agrees with the CreditMetrics model used for the base risk weights, ie, the systematic risk contributions from the two models coincide at 99.5%. This requirement is presented as:

$$F_A = p_A\left(1 - \omega_A + \omega_A X_{99.5\%}\right) - p_A = p_A \omega_A\left(X_{99.5\%} - 1\right)$$

where F_A is "systematic risk sensitivity" defined in Basel (2001, paragraph 431), ie, the excess of the CreditMetrics systematic risk contribution over the default probability.

1. Loss distribution – large portfolio

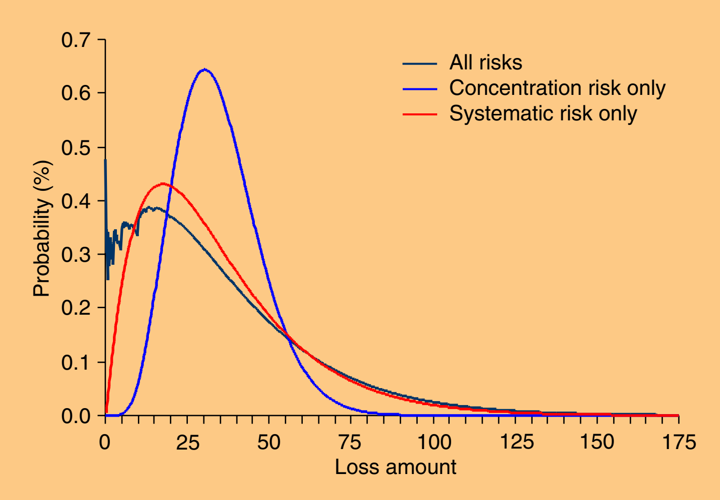

2. Loss distribution – small portfolio

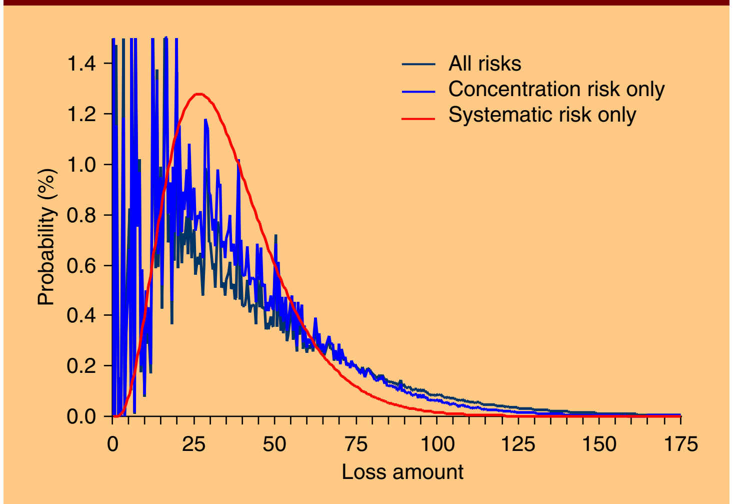

The gamma distributed X is parameterised with $\sigma(X) = 2$ (Basel, 2001, paragraph 445) and this leads to $X_{99.5\%} = 12.007$ (this is easy to check using GAMMAINV(0.995,0.25,4) in Excel). So the condition for equal percentiles is:

$$p_A \omega_A \sigma(X) = F_A \frac{\sigma}{\left(X_{99.5\%} - 1\right)} = F_A \frac{2}{\left(12.007 - 1\right)} = 0.182F_A \tag{8}$$

as given in Basel (2001, paragraph 454).

To understand the granularity adjustment, we rephrase the Herfindahl index in terms of a more universal concept, the variance of the loss distribution. From (4), the variance of the portfolio loss distribution conditional on X is:

$$\sigma^2\left(\Pi \middle| X\right) = \sum_A E_A^2\left(P_A(X)\left(1 - P_A(X)\right) + P_A(X)\sigma^2\left(E_A\right)\right)$$

So, the total variance of the loss distribution is:

$$\sigma^2(\Pi) = \sigma^2\left(\sum_A E_A P_A\right) + \sum_A E_A^2\left(p_A\left(1 - p_A\right) - \sigma^2\left(P_A\right) + p_A\sigma^2\left(E_A\right)\right)$$

using p_A as the average of P_A. The left-hand summand is the variance of the systematic loss variable μ given by equation (3). The other summand is called the unsystematic variance:

$$\sigma^2_{Unsys} = \sum_A E_A^2 \left(p_A\left(1-p_A\right) - \sigma^2\left(P_A\right) + p_A \sigma^2\left(E_A\right)\right) \quad (9)$$

In our one-factor CreditRisk+model we have:

$$\sigma^2\left(P_A\right) = p_A^2 \omega_A^2 \sigma^2\left(X\right) \quad (10)$$

and so substituting into equation (9):

$$\sigma^2_{Unsys} = \sum_A \left(E_A^2\left(p_A - p_A^2 - p_A^2\omega_A^2\sigma^2\left(X\right)\right) + E_A^2 p_A \sigma^2\left(E_A\right)\right) \quad (11)$$

The IRB approach chooses a specific form for recovery rate volatility, (Basel, 2001, paragraph 447), which, in our notation[9], is:

$$\sigma^2\left(E_A\right) = \frac{LGD_A\left(1 - LGD_A\right)}{4LGD_A^2} = \frac{\left(1 - LGD_A\right)}{4LGD_A}$$

To make our notation comparable with Basel (2001), we write our E_A as $EAD_A \times LGD_A$, where exposure before recovery is EAD_A. Then equation (11), using equation (8), becomes:

$$\sigma^2_{Unsys} = \sum_A EAD_A^2\left(LGD_A^2\left(p_A - p_A^2 - 0.033F_A^2\right) + 0.25 p_A LGD_A\left(1 - LGD_A\right)\right)$$

where 0.033 is the square of 0.182. This is given in Basel (2001, paragraph 452), except that our summation is directly over obligors.[10] Also, by ignoring terms of order higher than one in default probabilities, the denominator of the fraction for A_b is very nearly $PD_{AG}LGD_{AG}(0.25 + 0.75LGD_{AG})$. We substitute all this in the sum in Basel (2001, paragraph 455). The easiest way to do this is to assume that each risk grade contains only one obligor, so $H_b = 1$. We obtain, very nearly:

$$H^* = \frac{1}{n^*} = \frac{4\sigma^2_{Unsys}}{\left(\sum_A EAD_A\right)\left(1 + 3LGD_{AG}\right)EL} \quad (12)$$

where:

$$EL = \left(\sum_A EAD_A\right) LGD_{AG} PD_{AG} = \sum_A E_A p_A$$

is the mean loss from the portfolio. So, the Herfindahl index H^* contains much the same information as the unsystematic variance of the loss distribution. An odd term $1 + 3LGD_{AG}$ seems to have been stripped out, but this exact term is reinstated as one of the brackets in the equation for β, as can be seen by examining equation (7).

What does the IRB equation for the granularity adjustment look like in terms of unsystematic variance? Using equations (7) and (12), we can rewrite the adjustment as:

$$GA = \left(\sum_A EAD_A\right)\beta / n^* = 1.6 \times \left(0.76 + 1.1PD_{AG}/F_{AG}\right) \times \sigma^2_{Unsys}/EL$$

Thus the only information the granularity adjustment uses about the unsystematic risk is the unsystematic variance. This is not surprising – in general, no other reliable information about the unsystematic risk to use. A significant part of the unsystematic risk arises from the supposed recovery rate volatility, for which we used an assumption about variance, but no other information.

There is an alternative presentation of these facts. The actual adjustment is the result of the careful numerical work described in Gordy (2000b), but if one had nothing else to go on what would be one's first estimate of the granularity adjustment? Many practitioners would have scaled portfolio systematic unexpected loss by the ratio of standard deviations, ie, applied the formula:

$$GA_{estimate} = UL \times \left(\left(\sigma^2_{Total}/\sigma^2_{Sys}\right)^{1/2} - 1\right) \cong UL \times \sigma^2_{Unsys} / 2\,\sigma^2_{Sys}$$

In the one-factor CreditRisk+ model, we can use equation (12) to work this out. To do this, we also need to calculate the systematic component of variance, which we obtain from equation (10) bearing in mind that $\sigma(X) = 2$. After some manipulation, one obtains an equation similar to the actual adjustment:

$$GA_{estimate} = \left(9.47PD_{AG}/F_{AG}\right)\left(0.4 + 1.2LGD_{AG}\right)/n^*$$

This is a moderately successful approximation to the actual granularity adjustment and tends to overestimate, though it is close enough to provide a "reasonableness check". The crossover point where this expression agrees exactly to the actual adjustment expression is $0.76 = 8.37PD_{AG}/F_{AG}$, which happens for a default probability of about 30 basis points. The two expressions are within 25% of each other up to about a default probability of 1.30%.

Conclusion

The IRB risk weights and granularity adjustment are derived from a general framework of ideas that deserves to be widely understood as part of the basic language of risk management. The IRB approach is not necessarily right or the best possible, but at least rules derived from a general framework are better than a patchwork of inconsistent special cases. But this clear conceptual approach really only exists for the IRB risk weights and granularity adjustment – if the same approach had penetrated throughout the new Accord then arguably there could have been different proposals in several areas. For example, the definition of default and discussion of "through the cycle" versus "point in time" default probabilities, the treatment of retail assets, the treatment of credit risk mitigation (the "*w* factor") and joint default risk, and the treatment of counterparty risk are all areas that can be considered in the light of the ideas above, at least as one valid point of view, and this approach has not been adopted as universally as it could have been. ■

Tom Wilde is the London-based head of credit risk portfolio modelling at Credit Suisse First Boston. The first half of this article is based on an unpublished document presented by the International Swaps and Derivatives Association to US regulators in September 1999. The second half, on the new Accord's granularity adjustment, is new
Comments on this article can be posted on the technical discussion forum on the Risk website at http://www.risk.net

[9] *The division by* LGD^2 *is needed because we refer to proportional volatility of* LGD
[10] *There is no need for the intervening risk grades in Basel (2001), and, unfortunately, if the recovery rate is not constant, then, as a result of this formulation, the result can depend somewhat on the particular choice of grades. To avoid the issue assume recovery rates are constant across grades – after all, grades may be chosen at will and need not contain more than one obligor*

REFERENCES

Basel Committee on Banking Supervision, 2001
The new Basel capital accord
IRB Consultative Document, Bank for International Settlements, January

Belkin B, S Suchower and L Forest Jr, 1998
The effect of systematic credit risk on loan portfolio value at risk and loan pricing
CreditMetrics Monitor, March

Finger C, 1999
Conditional approaches for CreditMetrics portfolio distributions
CreditMetrics Monitor, April

Gordy M, 2000a
A comparative anatomy of credit risk models
Journal of Banking and Finance 24, pages 119–149

Gordy M, 2000b
A risk factor model foundation for ratings based capital rules
Working paper, September

Hickman A and U Koyluoglu, 1998
Reconcilable differences
Risk October, pages 56–62

Lucas A, P Klaassen, P Spreij and S Straetmans, 1999
An analytic approach to credit risk of large corporate bond and loan portfolios
Research memorandum, vrije Universiteit, Amsterdam, February

Merton R, 1974
On the pricing of corporate debt: the risk structure of interest rates
Journal of Finance 29, pages 449–470

Vasicek O, 1987
Probability of loss on loan portfolio
KMV Corporation