Collateral damage

Most credit risk models focus on default probability, while making simple recovery assumptions for collateralised loans. As Jon Frye shows here, this is a mistake, because the same factors that increase default rates can also decrease the value of loan collateral

If a borrower defaults on a loan, a bank's recovery may depend on the value of the loan collateral. The value of collateral, like the value of other assets, fluctuates with economic conditions. If the economy experiences a downturn, a bank can experience a double misfortune: many obligors default, and the value of collateral is damaged.

Conventional credit models overlook the effect of economic conditions on collateral. They allow default to vary from year to year, but they hold fixed the average value of collateral and the average level of recovery.

The distinctive feature of the credit model presented here is that an economic downturn causes damage to the value of collateral. When systematic collateral damage enters the credit model, the capital allocated to a highly collateralised loan can double or triple.[1]

Taking collateral damage into account complicates a credit capital model. However, the results of the model can be well approximated by a function of expected loss alone. Expected loss can therefore be used as the basis of a credit capital estimate. This estimate is simpler, and can be more accurate than using the results of a conventional credit model that ignores the role of collateral damage.

Credit capital model

The credit capital model uses the conditional approach suggested by Finger (1999) and Gordy (2000). The variables in the model depend on a systematic risk factor, a random variable representing the good years and bad years of the economy. The co-variation between two variables stems from their mutual dependence on the systematic factor. Two variables that relate strongly to the systematic factor relate strongly to each other and therefore have a strong correlation.

Exposure of \$1 is assumed to each obligor j. At the end of a one-year analysis horizon, the value of collateral is a random number characterised by three positive parameters: its amount, μ_j; its volatility, σ_j; and its sensitivity to X, the systematic risk factor, also known as its "loading", q_j:

$$\text{Collateral}_j = \mu_j\left(1 + \sigma_j C_j\right) \quad \text{and} \tag{1}$$

$$C_j = q_j X + \sqrt{1 - q_j^2}\, Z_j \tag{2}$$

where X and Z_j have independent standard normal distributions.

Equation (2) implies that C_j has a standard normal distribution. When the systematic factor exceeds zero, both C_j and Collateral_j tend to be greater than average, but that also depends on an idiosyncratic risk factor, Z_j, which affects only the collateral of obligor j. Equation (1) shows each unit of collateral value has a normal distribution with mean equal to 1.00 and standard deviation equal to σ_j.

The overall financial condition of the obligor, A_j, also depends on the systematic risk factor via a positive loading, p_j:

$$A_j = p_j X + \sqrt{1 - p_j^2}\, X_j \tag{3}$$

where X_j have standard normal distributions independent of each other, X and Z_j.

When X exceeds zero, obligor j tends to prosper. A_j also depends on the idiosyncratic variable X_j, which affects the fortunes of obligor j and nothing else. A_j may take on a wide range of values, having a standard normal distribution. This specification ignores the influences that may exist between X_j and collateral, and/or between Z_j and A_j. These non-systematic influences have a relatively minor effect on credit capital.

The correlation between two obligors depends on their loadings on the systematic risk factor X:

$$\text{Corr}\left[A_j, A_k\right] = \text{Cov}\left[p_j X + \sqrt{1 - p_j^2}\, X_j, p_k X + \sqrt{1 - p_k^2}\, X_k\right] = p_j p_k \tag{4}$$

An obligor defaults if its financial condition falls below a threshold. Let D_j represent the default event:

$$D_j = 1 \text{ if } A_j < \Phi^{-1}\left(PD_j\right);\ D_j = 0 \text{ otherwise} \tag{5}$$

where PD_j represents the probability of default for obligor j, and Φ^{-1} is the inverse cumulative standard normal distribution. Equation (5) thus ensures obligor j defaults with probability PD_j: $\text{Prob}\,[D_j = 1] = E\,[D_j] = PD_j$.

If default occurs, the bank can recover, properly discounted and net of foreclosure expenses, no more than the loan amount:

$$\text{Recovery}_j = \text{Min}\left[1, \text{Collateral}_j\right], \text{ that is, } LGD_j = \text{Max}\left[0, 1 - \text{Collateral}_j\right] \tag{6}$$

If default occurs and collateral value exceeds exposure, the bank has no loss. If default occurs and collateral value is less than zero, the bank may lose more than \$1. Taking the average of loss given default (LGD) over all possible outcomes produces the expected loss given default (ELGD). Solving backward, the ELGD of a loan implies the level of μ_j.

For simplicity, the model includes losses due only to default and not due to downgrade or changes in pricing spreads. The amount lost to obligor j is then the product of the default event and the loss given default:

$$\text{Loss}_j = D_j LGD_j \tag{7}$$

Taking the sum over all j equals the total credit loss. Monte Carlo simulation or other means can then determine the distribution for random realisations of the systematic factor X and of the idiosyncratic factors X_j and Z_j.

Capital models are used by some banks to target the credit ratings they receive from rating agencies. A bank that targets an investment-grade rating might wish to hold enough credit capital to absorb the loss that arises in 99.9% of Monte Carlo simulation runs. We may speak of a target solvency of 99.9% or of a target insolvency of 0.1%. The latter equals α, the final parameter of the credit capital model.

In equation (7), D_j depends on X and X_j, and LGD_j depends on X and Z_j. Conditional on a realisation $X = x$, these factors are independent:

$$E\left[\text{Loss}_j \middle| X = x\right] = E\left[D_j \middle| X = x\right] \times E\left[LGD_j \middle| X = x\right] \tag{8}$$

[1] *We give broad meanings to two terms. "Collateral" here includes all the assets a bank obtains as a consequence of default, including, but not limited to, the assets pledged as collateral in a loan document. "Capital" refers to equity capital and to accumulated loan loss reserves, both of which help banks weather stressful periods*

1. Effect of X on a loan: PD = 5%, ELGD = 10%

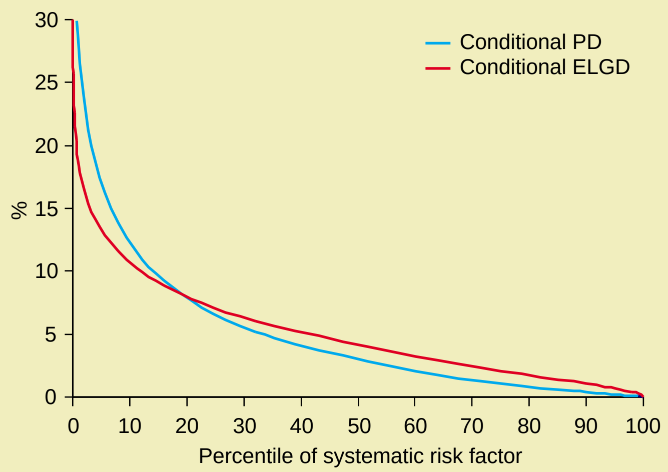

A. Expected performance of two loans

Overall expectation

	PD	×	ELGD	=	EL
First loan	5.0%		10.0%		0.5%
Second loan	1.0%		50.0%		0.5%

Expectation in an economic slump; target $\alpha = 0.1\%$

	PD	×	ELGD	=	Capital
First loan	45.4%		26.1%		11.8%
Second loan	18.4%		60.2%		11.1%

Economic slump in a conventional credit model

	PD	×	ELGD	=	Capital
First loan	45.4%		10.0%		4.5%
Second loan	18.4%		50.0%		9.2%

Thus, given a state of the economy as represented by $X = x$, the conditional expected loss for a loan equals the product of its conditional PD and its conditional ELGD.[2] An increase in x causes both conditional PD and conditional ELGD to decrease. Therefore, when X is at percentile α, conditional EL is at percentile $(1 - \alpha)$.

We want to find the increase in target capital when a particular loan is added to the credit portfolio. Marginal capital depends on the portfolio to which the loan is added. We assume that the portfolio is large enough to be fully diversified. Then the marginal capital of a loan can be treated as equal to the mathematical expectation of loss conditional on $X = \alpha$. Therefore, a bank that has an insolvency target of 0.1% can substitute the 0.1 percentile of the standard normal, $x = -3.09$, in equation (8) to obtain credit capital for a loan characterised by the five parameters σ_j, p_j, q_j, PD_j and μ_j.

The novel feature of the model is that LGD depends on the state of the economy. This feature is not present in conventional credit models. For example, CreditMetrics first determines obligor default and then independently determines LGD. Any risk in recovery is purely idiosyncratic, equivalent to forcing $q = 0$. CreditRisk+ assumes LGD is a known amount. The variance of recovery is zero, equivalent to forcing $\sigma = 0$. The capital model presented here can mimic these models. If q is set to zero (or if σ is set to zero), conditional ELGD does not respond to the state of the economy, but becomes, instead, a constant.

Quantifying collateral damage

Next we find representative values of the parameters σ, p and q. We then demonstrate the importance of collateral damage for an example loan. Finally, repeat the analysis for loans having a range of ELGD.

The parameter σ might apply to a number of assets pledged as collateral: inventory, receivables, negotiable instruments, title documents, intangibles etc. In addition, other assets with uncertain values may be awarded to the bank as a general creditor. For this mixture of assets, no single estimate of σ can be entirely satisfactory. To obtain a representative estimate, Moody's Investors Service provides summary data for 98 senior secured bank loans. On these, average recovery equals 70.26%, with standard deviation equal to 21.33%. The implied estimate of σ equals $21.33/70.26 = 30\%$. However, the recovery achieved on a loan is apt to be closer to the recovery expected on the same loan than to the overall average recovery. Therefore, the raw estimate is apt to overstate σ. We take $\sigma = 20\%$ as the representative value, with robustness checks using $\sigma = 15\%$ and $\sigma = 25\%$.

We take 0.5 as the representative value of p. Equation (4) then implies that any pair of obligors has a correlation of $(0.5)^2 = 25\%$. This accords with the average level of asset correlation suggested by the CreditMetrics Technical Document (1997). Checks of robustness are performed at $p = 0.4$ and $p = 0.6$.

No studies known to the author provide an estimate of q, the loading of collateral on the systematic factor. It appears, however, that the representative value of q is greater than or equal to p. First, all assets tend to decline with the systematic factor, whether or not they may become the source of recovery on a bank loan. If this overall systematic effect were the only channel of influence, one would suppose that $q = p$.

Two additional channels of influence increase the effect of X on collateral in an economic slump. A low value of X leads to financial distress for many obligors, and for some banks. An obligor in financial distress might devote fewer resources to resolving customer complaints, maintaining equipment and safeguarding its fixed investments. The affected assets – accounts receivable, vehicles and real estate – serve as collateral. Thus, the assets a bank obtains may already have been degraded by previous attempts to extract value. A bank must also anticipate the effects of its own distress. In the circumstances envisioned by the capital model, a bank has nearly exhausted its capital cushion. It then faces unusual pressure to liquidate assets even if it cannot obtain the best price. The value a bank actually realises from collateral may therefore be even more depressed than the values of other assets. Two channels of influence – the actions of distressed obligors before they default, and the actions of the distressed bank itself after it receives collateral – make collateral values unusually sensitive to an economic slump. It is difficult to see an opposing influence that would selectively protect collateral from systematic risk.[3] We assume that representative q is equal to p, with robustness checks of $q = p + 0.10$ and $q = p + 0.20$.

Using the representative values of the parameters σ, p and q, we take the example of a relatively low-rated obligor that has provided a high level of collateral. Specifically, figure 1 analyses a loan having $PD = 5\%$ and $ELGD = 10\%$. The horizontal axis is calibrated to the percentiles of X. The two lines represent the two factors on the right-hand side of equation (8). The average of conditional PD equals the unconditional PD of 5%. The PD-weighted average of conditional ELGD equals the unconditional ELGD of 10%.

Thus, figure 1 shows how the overall levels of PD and ELGD distribute conditionally across states of the economy. For most states, one sees relatively benign levels of both variables. But in a severe economic slump, both conditional PD and conditional ELGD rise with a vengeance. To prepare for this adverse circumstance, banks hold capital.

[2] Conditional EL, conditional PD and conditional ELGD refer to the expectation conditional on the realization of X. Given X, the expectation is taken across all j. When conditioning is clear from context, the modifier may be suppressed. Otherwise, when the variables are not designated as "conditional", they have the usual meaning of an all-inclusive expectation

[3] Some specific collateral, such as cash or Treasury securities, has a low value of q, but these cases are far from the norm

Suppose a bank targets its own insolvency at $\alpha = 0.10\%$. In figure 1, this point appears 0.1% of the distance along the axis, almost at the extreme left. The corresponding levels of conditional PD and conditional ELGD are 45.4% and 26.1%, respectively. According to equation (8), the product of these two equals credit capital: $45.4\% \times 26.1\% = 11.8\%$.

If the bank uses a conventional credit model to allocate capital for this loan, it makes a significant error. It ignores the increase in LGD that comes about in the economic slump. Specifically, it assigns capital equal to $45.4\% \times 10\% = 4.5\%$. The accurate target is 2.61 times this allocation, because the effect of collateral damage is to increase ELGD from its overall value of 10% to its value in an economic slump, 26.1%. Statistical data on the performance of bank loans would be preferable to such a prediction by a model. But data do not exist regarding LGD in an economic slump of this severity. Until we have such data, it would be risky to assume an economic slump has no effect on LGD.

Capital and expected loss

This section compares the loan of figure 1 with a second loan having equal expected loss. The equality of expected loss implies a near-equality of capital. The comparison is extended to loans with a range of combinations characteristics, and the same conclusion is found. The conclusion is then checked, both for a range of values of the model parameters σ, p, q and α, and for a fundamental change in the specification of the model.

The loan in the previous example has $PD = 5\%$ and $ELGD = 10\%$. Table A compares that loan with a second loan having the same EL, but having $ELGD = 50\%$. The middle of the table shows the first loan is more affected by collateral damage. The lower the ELGD of a loan, the greater potential it has to rise in an economic slump. An economic slump affects the ELGD of each loan, but it has a greater proportional effect on the first one, having a lower ELGD.

The two loans also differ in PD. The second loan has lower PD. In an economic slump, the PD of the second loan rises about 18-fold, while the PD of the first loan rises only about ninefold. The lower the PD of an obligor, the greater potential it has to rise in an economic slump. An economic slump affects the PD of each loan, but it has a greater proportional effect on the second one, having a lower PD.

The economic slump raises both ELGDs and both PDs. Of the two ELGDs, the proportional effect is greater for the first loan. Of the two PDs, the proportional effect is greater for the second loan. The product of the two effects is nearly the same, and so the two loans require nearly the same capital. In fact, by a narrow margin the first loan requires greater capital (11.8%) than the second loan (11.0%). This is contrary to the verdict of a conventional model. As shown at the bottom of table A, the conventional model reverses the ranking and presents an alarmingly rosy view of the low-ELGD loan.

The two loans in table A have the same EL and require nearly the same capital. (That the two loans divide EL differently between PD and ELGD has relatively little importance.) The relationship between EL and capital generalises readily, as shown in the context of a stylised bank internal rating system.

Many banks use one-dimensional internal risk rating systems. These systems initially assign a risk rating based on the characteristics of the obligor. The rating might be upgraded based on the amount of collateral securing a particular loan. A collateralised loan to a poorer-rated obligor then has the same rating as an uncollateralised loan to a better-rated obligor. This resembles the relationship of the two loans in table A. The first loan has relatively lower ELGD, and the second loan has relatively lower PD. The product of ELGD and PD can therefore be nearly equal for the two loans. When this is so throughout every rating grade, the rating system can be characterised as an EL system, even if it is not intentionally based on expected loss.

EL rating systems appear to be the norm. After a thorough study, Treacy & Carey (1998) of the Federal Reserve characterise the rating systems of nearly all large US banks as measuring EL. (They characterise some as measuring PD as well, separately.) We therefore assume the interplay between PD and ELGD results in uniform EL within a rating grade.

We establish grades for $EL = 0.025\%$, 0.05%, 0.1%, 0.2%, 0.4%,

2. Capital in a conventional credit model

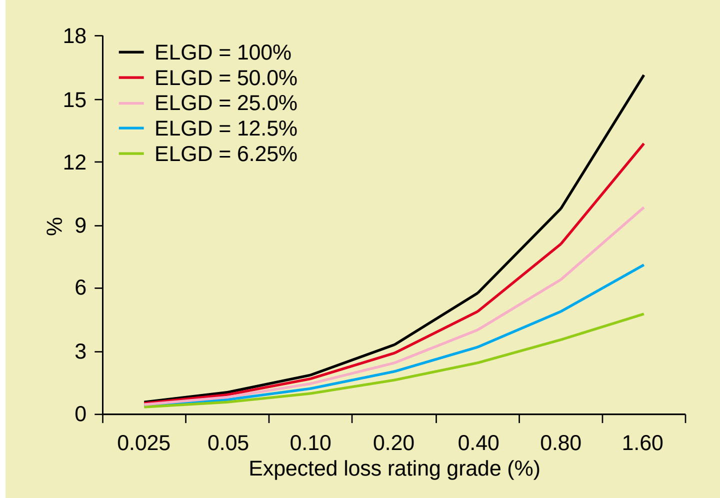

3. Capital including collateral damage

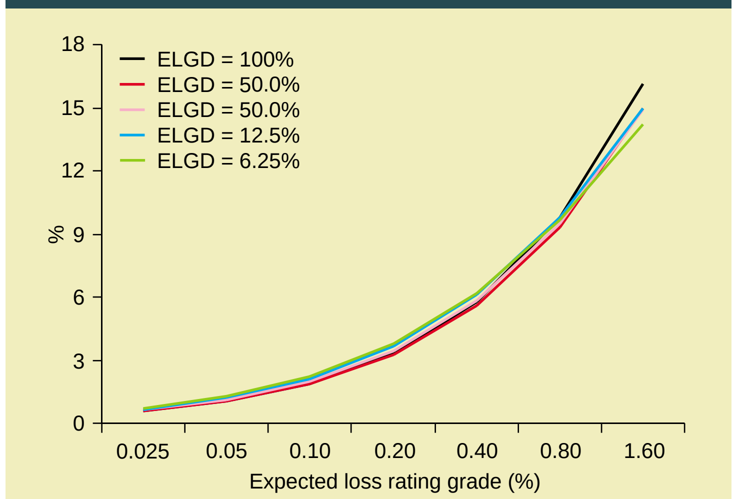

0.8% and 1.6%. This range includes most bankable assets. Within any rating grade, a conventional model allocates less capital to loans having lower ELGD, as shown in figure 2. (The target is low investment grade, $\alpha = 0.5\%$.) The five lines depict capital for five levels of ELGD. The top line depicts $ELGD = 100\%$, and the bottom line depicts $ELGD = 6.25\%$. In the rating grade where $EL = 0.1\%$, these correspond to $PD = 0.1\%$ and $PD = 1.6\%$, respectively.

Within any EL grade, the conventional model allocates more capital to loans with higher ELGD and lower PD. That is because the conventional model splits EL into PD and ELGD – and then looks only at the systematic risk in the PD fraction. In an economic slump, the greatest proportional increase in default occurs in obligors with the lowest PD. Therefore, within an EL rating grade, the conventional model concludes that more capital is required for loans where the probability of default is low.

The difference between figure 2 and figure 3 is the effect of collateral damage. The lower the ELGD, the more collateral and the more systematic risk is held by the bank. Therefore, the lines with the lowest ELGD rise the most in the transition from figure 2 to figure 3. Not only do they rise, they rise to approximately the same level. The result is that credit capital is approximately a function of expected loss alone.

Examination of figure 2 and figure 3 leads to three conclusions. First, the effect of collateral damage increases capital for all loans (except for

4. Robustness with conditional beta recovery

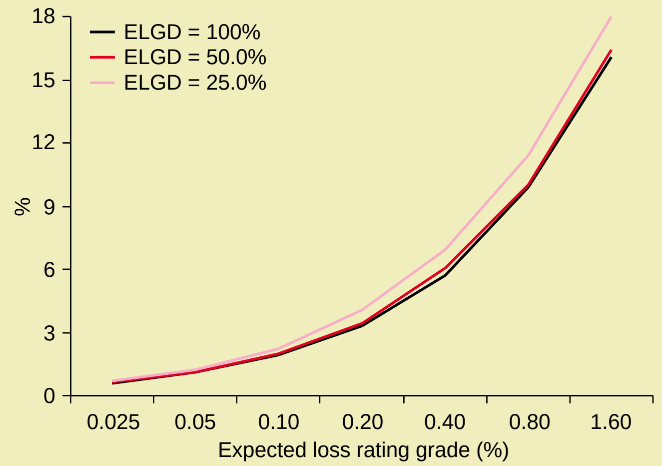

ELGD = 100%) and markedly increases capital for low-ELGD loans. Second, the low-ELGD lines in figure 3 are much closer to the line representing ELGD = 100% than they are to their own representations in figure 2. It appears more accurate to adjust the inputs for a low-ELGD loan – to adjust ELGD to 100% and to adjust PD downward, maintaining the same level of EL – than it is to simply accept the results of the conventional model using the unadjusted input. Third, a function of expected loss provides an estimate of credit capital that is more accurate than using both PD and ELGD in a conventional credit model.

These conclusions are tested for robustness. Thirty-six combinations of parameter values are used to recreate figure 2 and figure 3. Three levels of σ (15%, 20% and 25%) are combined with six combinations of values for p and q ({0.4, 0.4}, {0.4, 0.5}, {0.4, 0.6}, {0.5, 0.5}, {0.5, 0.6} and {0.6, 0.6}) and examined at two settings of α (0.1% and 0.5%). Each of the 36 pairs of charts supports all three of the conclusions stated above.[4]

These conclusions are not only robust for a range of parameter values, but also for a change in the mathematical specification of the model. In this specification there is no explicit role for collateral. Instead, recovery is modelled directly as a beta distribution, as is done in CreditMetrics. The model allows collateral damage to enter by conditioning recovery on the value of C_j. Specifically, replacing equation (1) we have:

$$\text{Recovery}_j = \text{BetaInv}\left[\Phi\left(C_j\right), \text{mean} = \mu,\ \text{s.d.} = \sigma\right] \qquad (9)$$

Figure 4 shows the results of the conditional beta recovery model. Figure 4 resembles the normal model of figure 3, except that the beta recovery model allocates slightly more capital to low-ELGD loans. Robustness checks of the beta recovery model also resemble the robustness checks of their normal model counterparts.[5] The conclusion – that capital depends principally on expected loss – is therefore robust with respect both to changes in parameter values and to a change in the model specification.

To account for the effects of collateral damage in a credit portfolio, the best solution would be to correctly model the effects of PD, ELGD, σ, p and q along the lines suggested in this article or in some other way. Most banks would find that they have neither the data nor the systems to adopt this approach in the near term. A second-best solution is to estimate capital as a function of expected loss, stratifying the portfolio by uniform σ, p and q as in the above.

For current users of a conventional credit model, a second-best solution appears to be an appropriate adjustment of the inputs. The suggestion is to adjust PD downward and to adjust ELGD to unity, keeping the product equal to the EL of the original loan.[6] The adjusted loan should contribute approximately the risk of the original loan including the risk of collateral damage. This adjustment is available to users of both CreditMetrics and CreditRisk+.

Conclusion

The credit capital model presented here takes note of an effect well known to bankers: the credit cycle can produce a double misfortune involving greater-than-average default frequency and poorer-than-average recoveries. Of the two misfortunes, conventional credit models analyse the first and ignore the second. They can therefore assign alarmingly little capital to well-collateralised loans.

The effect of economic conditions on loan recoveries complicates the capital model. However, the results of the full model are well approximated by a function of expected loss. This conclusion holds for an alternative model specification and for a robust range of parameter values.

These results contain several messages. To bank lending and credit policy officers, the results repeat a message most often heard following large credit losses: collateral should not lead to complacency, because collateral value can decline at exactly the moment that a bank gains ownership. To bank portfolio credit analysts who use models to estimate portfolio risk, the results warn that all sources of systematic risk must be included. Lacking that, the inputs to existing models should be adjusted for more accurate results. To bank supervisors attempting to assess credit risk, the results suggest that a simple estimate of credit capital can be expedient and accurate.

Naturally, bank credit models should be expanded to cover as many sources of risk as possible. An estimate based on expected loss would not be completely accurate or ideal. However, the expected loss approach may provide a better estimate than some current credit models. Until models evolve to incorporate the systematic risk of both default and recovery, a credit capital estimate based on expected loss may be the best solution. ■

Jon Frye is head of the models team in the capital markets group at the Federal Reserve Bank of Chicago. The views expressed are his and do not necessarily reflect those of the Federal Reserve Bank of Chicago. He would like to thank many readers for helpful comments on earlier versions, especially Lisa Ashley, Robert Bliss, Richard Cahill, Paul Calem, Matthew Foss, Michael Gordy, David Jones, Catherine Lemieux, Michael Lesiak, Carol Lobbes, Laura McGrew, Perry Mehta, James Nelson, Edmund Waggoner and participants at the 1999 FRB-Chicago Capital Markets Conference
Comments on this article can be posted on the technical discussion forum on the Risk Web site at http://www.riskpublications.com/risk

[4] An Excel workbook with these results is available from the author at Jon.Frye@chi.frb.org
[5] Robustness checks with lower q and/or lower σ can include lower levels of ELGD. They reach the same conclusions about the relationship of EL to credit capital
[6] Many CreditRisk+ users adjust model inputs now. They adjust exposure (rather than PD) downward as they adjust ELGD to unity, keeping the product equal to that of the original loan. This affects ELGD but not PD, so the EL of the proxy loan differs from the EL of the original. The difference in EL leads to the understatement of capital seen in figure 2. Mechanically, the understatement comes about because the downward adjustment of exposure dominates the upward adjustment of ELGD

REFERENCES

Credit Suisse Financial Products, 1997
CreditRisk+: a credit risk management framework

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

Gordy M, 2000
A comparative anatomy of credit risk models
Journal of Banking & Finance, January, pages 119–149

Gupton G, C Finger and M Bhatia, 1997
CreditMetrics technical document
JP Morgan

Moody's Investors Service, 1999
Global credit research
Historical Default Rates of Corporate Bond Issuers, 1920–1998

Treacy W and M Carey, 1998
Credit risk rating at large US banks
Federal Reserve Bulletin, November, pages 897–921

Weighting for risk

Basel has recognised that collateral and seniority give banks an advantage when an obligor defaults. Here, Jon Frye argues that the proposal may encourage banks to lend on the collateral – a practice that could threaten their own survival – and proposes a possible remedy

The advanced approach of the new Basel capital Accord seeks to improve on the current Accord by providing banks with better incentives. For example, collateral and security give a bank an advantage in a default situation. The new Accord rewards a bank that obtains collateral and security, because both affect loss given default (LGD), and LGD affects required capital. Since the current Accord is relatively insensitive to risk, making capital sensitive to LGD can represent an improvement.

But the new Accord goes too far. Its preference for low-LGD lending is so strong that it encourages banks to make low-LGD loans with reduced regard for default risk. This practice, known as "lending on the collateral", gives primary consideration to collateral and LGD, rather than to the borrower's ability to repay. This style of lending has been the traditional speciality of commercial finance companies, which generally have more capital than banks. But under Basel II, banks will find that lending on collateral requires less capital than a more customary bank loan.

This exaggerated incentive can be corrected with a small change in the risk weight function, to make it have a non-linear response to LGD. This article presents one possibility, which is preferred because of its simplicity. The preferred function agrees with Basel II over a wide range of conditions, but becomes distinctly more conservative when LGD is low.

Risk weights miss the risk

A principal goal of Basel II is to make regulatory capital requirements more compatible with risk. If this goal is to be achieved, a loan having greater risk should require more capital. Yet this does not seem to be the case. The apparent distortion of incentives is illustrated by two hypothetical loans.

The first hypothetical loan is to an obligor having a one-year probability of default equal to 20%. If such an obligor has a public rating, it is probably lower than B–. To bolster its creditworthiness, this obligor offers substantial over-collateralisation, so that in the event of default the bank expects to lose only 5% of the outstanding amount. Stated as inputs to a risk weight function, this loan has a probability of default (PD) equal to 20% and LGD equal to 5%.

The second hypothetical loan represents a more common lending situation. The obligor has a probability of default equal to 1%. The lending facility is senior but unsecured, and the bank estimates LGD equal to 50%.

Assuming the PDs and LGDs are estimated accurately, which of the two loans has greater risk? Most bankers, and most bank supervisors, would identify the first loan as riskier. In part this is because it represents the practice of "lending on the collateral". The trouble with lending on collateral is that when the obligor defaults, the value of collateral can fall below expectations. From a capital and risk perspective, prudent bankers would make the second loan in preference to the first.

This instinct is confirmed by a survey of bank practices. The survey was conducted by Risk Management Associates (RMA, formerly Robert Morris Associates), as part of its response to the Basel Committee. The survey asked each bank in the RMA capital working group to quantify capital for a one-year commercial loan, for various ranges of PD and LGD. Reading from the RMA matrix, the average bank finds the first loan to have distinctly more risk. The mean levels of capital are 5.8% and 3.7%, respectively.[1]

The Basel II risk weights take the opposite view, and encourage making the first loan in preference to the second. That would impose capital requirements of 5% and 10% respectively.[2]

The reversal of incentive is partly explained in a footnote in the internal ratings-based (IRB) document (Basel, 2001). There, it is assumed that a bank can lose no more than LGD. Since the LGD of the first loan is 5%, capital can be no more than 5%. There is an "LGD ceiling" that is binding.

What this overlooks is that LGD varies from year to year. Especially in a high-default year, LGD tends to exceed its average.[3] The conditions that push the default rate higher also tend to push LGD higher. Therefore, in a high-default year a bank can lose more than average LGD.

It might then appear that the entire problem with the Basel II risk weights is the LGD ceiling, but this is not the case. In fact, the ceiling is hardly binding. Without the ceiling, capital for the first loan equals 5.3%. Thus, even if the LGD ceiling was eliminated, the Basel II risk weights strongly favour what appears to be the riskier of the two loans. (The LGD ceiling is ignored in the remainder of this article.) Even after adding capital for operational risk, Basel II allocates less than 8% capital to the first loan. Thus, the Basel II risk weights provide an incentive for lending on the collateral, whether it is compared with the second hypothetical loan or with the average loan today.

The Basel II favouritism for the low-LGD style of loans extends to other cases. For example, an obligor whose probability of default equals 10% might provide collateral to reduce LGD to 10%. Again the lender feels protected by the collateral, but the inherent quality of the borrower is low. Basel II requires capital of only 7.7% for such a loan, in preference to the second of our hypothetical loans. Furthermore, the Basel II preference for low-LGD loans extends to other PDs. For our hypothetical loans, even if the probabilities of default were one-tenth, or one-hundredth, as great as mentioned, Basel II would still encourage banks to make the first loan in preference to the second.

Now consider the risk weight function itself. For corporate exposures in IRB, the risk weight function is found in paragraph 156:

$$\text{RW} = \left(\text{LGD}/50\right) \times \text{BRW}\left(\text{PD}\right) \qquad (1)$$

where BRW is the Basel II risk weight function. The BRW function is non-linear and concave. If PD falls by half, BRW also falls – but it falls by less than half. Because of this concave relation, BRW produces conservative risk weights at even low levels of PD.

By contrast, the relation between risk weight and LGD is strictly pro-

[1] *These levels of capital are low, compared with the Basel II levels, for reasons stated by Robert Morris Associates. The present concern is simply to find which loan has greater risk*

[2] *The calculations and functions referenced in this article are provided in a spreadsheet on request. Please send an e-mail to* Jon.Frye@chi.frb.org*, using as the subject field "Weighting for Risk"*

[3] *This has been detected in Frye (2000b) and in an unpublished study by Edward Altman, and has been observed at banks*

1. Preferred risk weight/Basel II

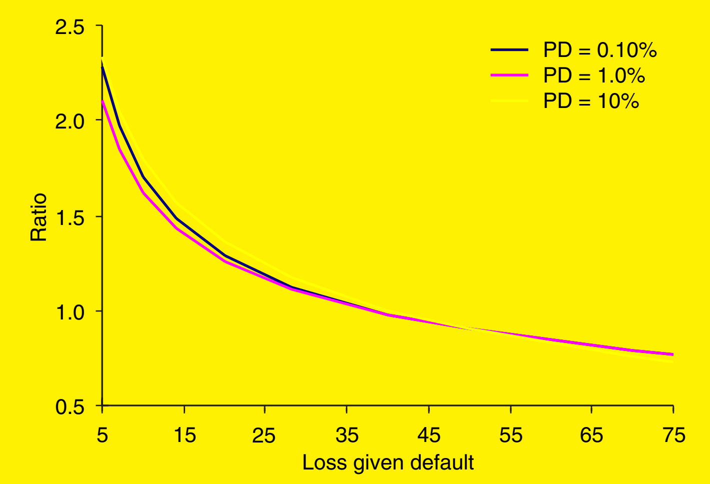

2. Preferred capital and Basel II capital (PD = 1%)

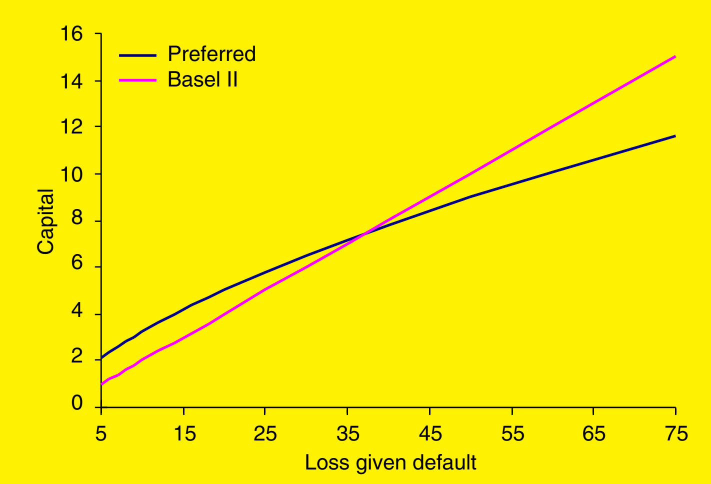

portional. If LGD falls by half, the risk weight also falls by half. This proportional relationship means that equation (1) cannot produce a conservative risk weight when LGD is low. It is the proportional response to LGD that leads to the low capital requirement on the first hypothetical loan, and to the exaggerated preference for low-LGD lending.

The preferred alternative

Ideally, the problem diagnosed in the previous section would be corrected after a thorough analysis of LGD. This analysis would produce the concave risk weight function of LGD that best fits historical experience and projections of an adverse year. At present, however, no such analysis can be fully satisfactory, because of the scarcity and generally poor quality of data on LGD.

Even though the LGD data has flaws, the risk weight function need not assume a simple proportional response to LGD. An improved function would track Basel II through moderate levels of LGD, but it would become more conservative when LGD is low. This would require a concave function. The concave function used in the preferred approach is the one already at hand, BRW. This approach simply moves LGD from outside the BRW function to inside:

$$\text{Preferred RW} = \text{K} \times \text{BRW}\left(\text{PD} \times \text{LGD}/50\right) \qquad (2)$$

The factor K would be chosen such that the total capital requirement for the banking industry would be the same in equation (2) as in (1). In this article, K is assumed to equal 90%, but a careful study might arrive at a different factor.

The preferred risk weight function (2) is explicitly *ad hoc* – it is addressed specifically to this problem that arises with low LGDs, rather than being the result of a theoretical exercise. Though it does not claim to arise from theory, it claims to produce better results. Specifically:

(1) The preferred risk weights restore the ranking of the hypothetical pair of loans.

(2) The preferred risk weights are distinctly more conservative than Basel II when LGD is low. They approximate Basel II when LGD is moderate.

(3) The preferred risk weights agree, better than the Basel II risk weights, with the actual behaviour of LGD.

The first of these is easily established. For the hypothetical pair of loans, the preferred risk weights imply capital of 13.9% and 9.0%, respectively. The preferred weights recognise the first hypothetical loan as substantially more risky.

Note the way that the ranking is restored. For the second loan, the preferred risk weights reach almost the same judgement as Basel II. But for the first loan, the preferred approach allocates 13.9% capital rather than 5.3%, or about 2.6 times as much.

To establish (2) more generally, figure 1 shows the ratio of the preferred risk weight to the Basel II risk weights, for three levels of PD. When the ratio is near 1.0, the preferred risk weights are near the Basel II risk weights. Figure 1 shows that the ratio is near 1.0 for moderate levels of LGD. Specifically, for LGD between 25 and 70, the ratio is between 0.75 and 1.25, which means that the preferred weight is within 25% of Basel II. This holds for all three levels of PD, 0.1%, 1.0% and 10%. Figure 1 also shows that the preferred risk weights are distinctly more conservative than Basel II when LGD is low. For example, if LGD equals 5%, the preferred risk weight is more than double the Basel II risk weight.

In figure 1, it is perhaps surprising that the three lines, reflecting three very different levels of PD, are close to each other across the entire spectrum of LGD. This simply means that when it comes to the ratio depicted in figure 1, the level of PD has little effect.

The preferred risk weights can therefore be seen as multiples of the Basel II risk weights, where the size of the multiple depends principally on LGD. For moderate LGD, the multiple is near 1.0. If LGD falls below 25%, the multiple rises rapidly, which results in a more conservative capital requirement for low-LGD style lending.

This behaviour establishes claim (2). The preferred risk weights are in broad agreement with Basel II for moderate LGD, but become distinctly more conservative for low LGD.

In passing, figure 1 shows a practical effect of the preferred risk weights as they would contrast with Basel II. Foundation IRB institutions, which cannot recognise LGD below 40%, would find capital requirements moderately reduced. Advanced IRB institutions would find capital decreased for some assets and increased for others.

Figure 2 compares the capital requirements for a borrower with PD = 1.00%.[4] The straight line shows the the linear response of capital under Basel II. The concave, less steeply sloped line shows capital under the preferred approach. Either approach offers an incentive to reduce LGD. The preferred approach is more conservative than Basel II when LGD is low.

Claim (3) compares the two risk weight functions with data from defaulted debt losses. This is the subject of the next section, which begins by highlighting the central feature observed in the data.

LGD responds to adversity

The Basel II risk weight function embeds an assumption regarding LGD in the adverse year. The adverse year is simply the year in which bank credit loss is at its 99.5 percentile. Data from the adverse year does not exist. For the default rate, Basel II uses a model to project its value in the adverse year.

The issue being raised is LGD in the adverse year. If a particular kind

[4] *As figure 1 suggests, other levels of PD produce a similar appearance, except for a change in the range of the vertical axis*

of loan has an LGD of 10%, what can we expect in the adverse year? This section shows that Basel II expects that all LGDs to rise 56% above their long-term averages. The preferred risk weights embed a different assumption, which makes the response of LGD sensitive to the level of LGD. A low LGD can therefore rise by more than 56% in the adverse year. This section then shows that the latter assumption – that the response of LGD is not uniform – is a better fit for loan and bond data.

To see the assumption embedded in the Basel II function, multiply (1) by 8% to obtain capital, substitute the BRW function from IRB paragraph 171, and gather together the constant factors:

$$\begin{aligned}\text{Capital} &= 8\% \times (\text{LGD}/50) \times 976.5 \times N\left[1.118 \times G(\text{PD}) + 1.288\right] \\ &\quad \times \left(1 + 0.047 \times (1 - \text{PD})/\text{PD}^{0.44}\right) \\ &= 1.56 \times \text{LGD} \times N\left[1.118 \times G(\text{PD}) + 1.288\right] \\ &\quad \times \left(1 + 0.047 \times (1 - \text{PD})/\text{PD}^{0.44}\right) \qquad (3) \\ &= (A) \times (B) \times (C) \times (D)\end{aligned}$$

Each of the four factors in this expression has a well-defined identity. Factor A is a number that makes regulatory capital equal to 8% for a specified reference loan ($\text{PD} = 0.7\%$ and $\text{LGD} = 50\%$). Factor B is expected LGD. Factor C is the function that projects the default rate in the adverse year, employing considerable theory and assumptions.[5] Factor D is an *ad hoc* adjustment from the one-year loans assumed by factor C to the three-year loans assumed by Basel II.

To show how equation (3) works in a concrete example, we recalculate the capital required by Basel II for the first hypothetical loan discussed above:

$$\text{Capital} = 1.56 \times 5.0\% \times 63.6\% \times 1.08 = 5.3\% \qquad (4)$$

Assuming that factor D does its job of correcting for the difference between three-year loans and the one-year analysis horizon, the Basel II function is saying that to protect itself against default loss, a bank should hold capital equal to:

$$\text{Capital} = 1.56 \times \text{LGD} \times \text{Default rate in the adverse year} \qquad (5)$$

In contrast to this amount of capital, the default loss in the adverse year equals LGD times the default rate. Therefore, the amount of capital in (5) is accurate only if, at every point along the LGD spectrum, LGD in the adverse year is 1.56 times its expectation.

This assumption does not appear explicitly in the Basel II documents and probably did not play a role in calibrating the risk weights. Implicitly, however, the Basel II risk weights assume that in the adverse year all LGDs respond the same, by rising 56% above their long-term averages. An LGD of 10% would rise to 15.6%, an LGD of 50% would rise to 78% and an LGD of 75% would rise to 117%.

Of course, average LGD cannot rise above 100%, so we seem to have arrived at a logical contradiction in the Basel II approach. The contradiction might be overcome with a bit of reinterpretation, but the fundamental problem – a proportional response to LGD – remains.

A different assumption is made by the preferred risk weights. It can be simply stated, since the Basel II risk weights implicitly assume that LGD rises by 56%. For the preferred approach, the response of LGD to the adverse year is as follows:

$$\text{Adverse year LGD}/\text{Expected LGD} = 1.56 \times \text{Preferred RW}/\text{Basel II RW} \qquad (6)$$

Equation (6) shows that the response of LGD to the adverse year equals 1.56 times the ratio that was charted in figure 1. That ratio depends primarily on LGD. Thus, the preferred approach assumes that LGD response depends on LGD, while Basel II assumes that LGD response is insensitive to LGD. The two assumptions will next be compared with the available data from defaulted debt recoveries.

Very little data is available at present. Some studies find that LGD responds to a high-default period, but to distinguish the alternatives we need to know the difference between the response of low-LGD assets and the response of high-LGD assets. The author is aware of only one study of bank loans that addresses this difference. This internal bank study compares the LGD of several collateral types within two periods, an overall period spanning 1989 to 1999, and a high-default sub-period spanning 1989 to 1991. The data appears in table A.[6]

A. Response of loan LGDs to a high-default period

Collateral type	Average LGD overall period	Average LGD high-default period	Response to high-default period
(1)	16%	50%*	213%*
(2)	20%	56%*	180%*
(3)	22%	37%	68%
(4)	30%	53%	77%
(5)	38%	59%	55%
(6)	40%	58%	45%

* Based on small number of defaults

B. Response of bond LGDs to a high-default period

Seniority	Average LGD overall period	Average LGD high-default period	Response to high-default period
Senior secured	37%	50%	35%
Senior unsecured	50%	58%	16%
Subordinated	63%	71%	13%

Although table A reflects the experience of only one bank, it tells an important story from the standpoint of choosing a set of risk weights. First, LGDs do not respond uniformly to a high-default period, but rather, as the last column shows, there is a range of response from 45% to 213%. Second, lower LGDs generally respond more strongly than do higher LGDs. Third, most of the collateral types show a response greater than 56%.

All three features contradict the assumption implicit in the Basel II risk weights. The greater-than-56% responses, combined with the increased responsiveness of low-LGD loans, should raise warning flags for Basel II, because they suggest the potential for a serious understatement of regulatory capital.

As well as investing in loans, banks invest in bonds, and risk weights apply to them. Bonds and loans differ in many ways, but bonds are generally subordinated to loans and therefore have greater LGDs. The bond LGD data comes from the Moody's Default Risk Service database, and reflects the same selection criteria established in Frye (2000b). For bonds of three seniority levels, table B displays both the average LGD on the overall period of 1983–1997 and the average LGD in the high-default sub-period of 1990–1991.

As with the loan data, bond LGDs do not respond uniformly to a high-default period, and lower LGDs show a stronger response. These features undermine the Basel II assumption of a uniform response. In contrast with the loan data, the bond LGD responses are much less than the 56% assumed by Basel II. Unless loan LGDs respond much more strongly than bond LGDs, this feature suggests that Basel II may be too conservative for high-LGD loans, such as subordinated lending under foundation IRB, which has LGD set to 75%.

[5] The "N" function denotes the cumulative distribution of a standard normal variable, and the "G" function denotes its inverse. Factor (C) is derived in numerous publications. Using the notation of equation (3) in Frye (2000b), the Basel II risk weights employ $\rho^2 = 0.20$ *and* $X = -2.5758$*, which represents the 0.5 percentile of the standard normal*

[6] The complete list of collateral types includes most of the common categories. Some of these collateral types had only a small number of observations in the low-default sub-period and were eliminated from the analysis. No identification is made of the collateral types appearing in table B

3. Response of LGD to high-default period

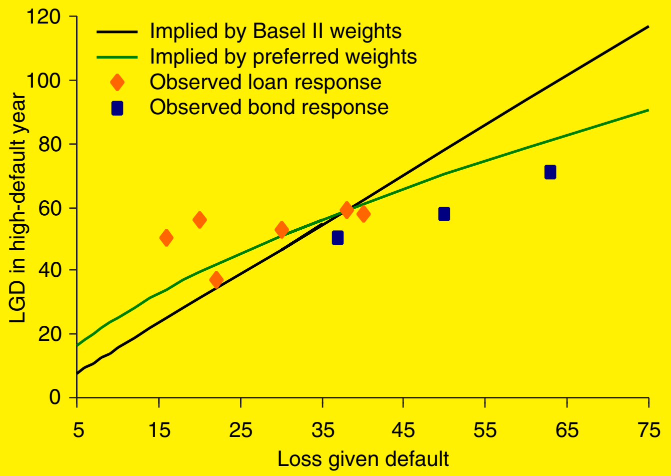

Thus, both in loans and in bonds, LGD responds to high-default periods. Assets with lower LGDs show a greater response, contrary to Basel II, which imagines that every LGD responds by 56% to the adverse year. This suggests the Basel II function is too conservative for high-LGD lending and not conservative enough for low-LGD lending.

Figure 3 brings together the elements discussed in this section: the loan data and the bond data appear as points, and the assumptions embedded in Basel II and in the preferred approach appear as lines. The horizontal axis shows expected LGD. The vertical axis shows LGD in the high-default period (for the loan and bond data) and LGD projected for the adverse year (for the two sets of risk weights). The straight line shows the linear assumption implicit in Basel II: regardless of level, each LGD responds to the adverse year by rising 56%. The concave, less steeply sloped line shows the assumption implicit in the preferred approach.

The pattern of the data points agrees better with the preferred approach. In fact, the data pattern appears less steeply sloped than either projection. This suggests little danger that the preferred approach overstates the response of LGD.

In using this data to help choose between the two sets of risk weights, we must extrapolate twice: to lower LGD assets, and to more adverse financial conditions. The Basel II risk weights perform the extrapolation by assuming that LGDs of all levels will respond to the adverse year by rising 56%. The preferred risk weights project a low percentage response for high LGDs, and a high percentage response for low LGDs. This non-uniform response provides a better match for the data we have and for the conditions we can imagine.

Objections and responses

A number of objections might be raised to the proposal presented above. This section responds to several of them.

☐ **Objection**: The "flaw" in Basel II has little effect, because low-LGD loans do not exist. **Response**: Many banks are devising LGD ratings and quantification systems, and some of these systems have buckets that map to LGD = 5% and LGD = 10%. Especially if the risk weights of Basel II are adopted, banks will have a strong incentive to analyse their historical data to discover the characteristics that have led to low average LGDs, and then to pursue business that has those characteristics.

☐ **Objection**: Low LGD loans exist, but bank regulators will not allow classification into very low LGD grades. **Response**: Regulators will allow banks to use any level of LGD that they can adequately support.

☐ **Objection**: Loans having low LGD also have low collateral risk, so the potential to respond to the adverse year is not great. **Response**: Some obligors pledge cash or Treasury bill collateral on some loans, and regulators may choose to allow lower capital for these low-response loans. But regulators should not conclude that every low-LGD loan is a low-response loan. An increasing share of low-LGD exposure is to asset-backed lending, which achieves a low LGD through collateralisation. The collateral is rarely if ever cash or Treasury bills, and it probably has risk equal to other collateral.

☐ **Objection**: It is too soon to adopt the preferred risk weights, because we have insufficient data to be certain that they are optimal. **Response**: Regulation will not wait for the data. In particular, regulation cannot wait to experience the adverse year. A judgement must be made, aided by the evidence that is available. Part of the evidence, which will not change with the passage of time, is that Basel II encourages banks to lend on collateral.

☐ **Objection**: The preferred risk weights cannot be adopted, because a good model does not support them. **Response**: A good model, and good statistical analysis, lead to a risk weight function of the same form. As discussed in Frye (2000b), capital involves the product of functions that project LGD and PD to their adverse levels. To a high degree of accuracy, the product of these functions can be approximated by a function of the product of LGD and PD. Conversely, figure 1 shows that the risk weights presented in this article are approximately the product of separate functions of PD and LGD. No doubt, the regulatory risk weight function will evolve as theory develops and data accumulates. But for the time being, the simplicity of using a single function makes it appear preferable.

Conclusion

If regulatory risk weights provide a meaningful incentive to banks, the weights that appear in the consultative document will have distorting effects on bank lending. They will encourage banks to engage in a form of lending more appropriate to specialised, better-capitalised finance companies. They will encourage banks to accept, through reliance on collateral, greater systematic risk than is prudent. Most importantly, if the level of LGD rises sharply at the same time as the rate of default, they might leave banks with insufficient capital to survive the adverse year when it arrives.

A simple change to the risk weight function could avoid these ills. One possibility has been stated and shown to be distinctly more conservative for low-LGD loans. It gives banks an incentive to reduce LGD, but a more moderate incentive. It agrees with the data we have from loans and bonds, and with reasonable ideas of what is apt to happen in an adverse financial environment. It is by no means the only remedy, nor can it be proven, with the paucity of data currently available, to be the best remedy. But it represents an improvement. ■

Jon Frye is senior economist in the policy group at the Federal Reserve Bank of Chicago. **He thanks Mark Carey and Marc Saidenberg for valuable conversations. For helpful comments on earlier versions, he thanks Mike Atz, Matthew Foss, Dale Klein, Cathy Lemieux, Dennis McLaughlin and the members of the FRB-Chicago regulatory capital comment group. The views expressed are the author's and do not necessarily represent the views of the management of the Federal Reserve Bank of Chicago or the views of the Federal Reserve System**
Comments on this article can be posted on the technical discussion forum on the Risk website at http://www.risk.net

REFERENCES

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

Frye J, 2000a
Collateral damage
Risk April, pages 91–94

Frye J, 2000b
Depressing recoveries
Risk November, pages 106–111

Moody's Investors Service, 2000
Default risk service database

Robert Morris Associates, 2000
Response to the Basel Committee's consultative paper on a new capital adequacy framework