Industry Practices in Credit Risk Modeling and Internal Capital Allocations: Implications for a Models-Based Regulatory Capital Standard

Summary of Presentation

David Jones and John Mingo

I. Why Should Regulators Be Interested in Credit Risk Models?

Bank supervisors have long recognized two types of shortcomings in the Basle Accord's risk-based capital (RBC) framework. First, the regulatory measures of "capital" may not represent a bank's true capacity to absorb unexpected losses. Deficiencies in reported loan loss reserves, for example, could mask deteriorations in banks' economic net worth. Second, the denominator of the RBC ratios, total risk-weighted assets, may not be an accurate measure of total risk. The regulatory risk weights do not reflect certain risks, such as interest rate and operating risks. More importantly, they ignore critical differences in credit risk among financial instruments (for example, all commercial credits incur a 100 percent risk weight), as well as differences across banks in hedging, portfolio diversification, and the quality of risk management systems.

These anomalies have created opportunities for "regulatory capital arbitrage" that are rendering the formal RBC ratios increasingly less meaningful for the largest, most sophisticated banks. Through securitization and other financial innovations, many large banks have lowered their RBC requirements substantially without reducing materially their overall credit risk exposures. More recently, the September 1997 Market Risk Amendment to the Basle Accord has created additional arbitrage opportunities by affording certain credit risk positions much lower RBC requirements when held in the trading account rather than in the banking book.

Given the prevalence of regulatory capital arbitrage and the unstinting pace of financial innovation, the *current* Basle Accord may soon become overwhelmed. At least for the largest, most sophisticated banks, it seems clear that regulators need to begin developing the next generation of capital standards now—before the current framework is completely outmoded. "Internal models" approaches to prudential regulation are presently the only long-term solution on the horizon.

The basic problem is that securitization and other forms of capital arbitrage allow banks to achieve *effective* capital requirements well below the *nominal* 8 percent Basle standard. This may not be a concern—indeed, it may be desirable from a resource allocation perspective—when,

David Jones is an assistant director and John Mingo a senior adviser in the Division of Research and Statistics of the Board of Governors of the Federal Reserve System.

in specific instances, the Basle standard is way too high in relation to a bank's true risks. But it *is* a concern when capital arbitrage lowers *overall* prudential standards. Unfortunately, with the present tools available to supervisors, it is often difficult to distinguish these cases, especially given the lack of transparency in many off-balance-sheet credit positions.

Ultimately, capital arbitrage stems from the disparities between true economic risks and the "one-size-fits-all" notion of risk embodied in the Accord. By contrast, over the past decade many of the largest banks have developed sophisticated methods for quantifying credit risks and internally allocating capital against those risks. At these institutions, credit risk models and internal capital allocations are used in a variety of management applications, such as risk-based pricing, the measurement of *risk-adjusted* profitability, and the setting of portfolio concentration limits.

II. The Relationship between PDF and Allocated Economic Capital

Before discussing various credit risk models per se, it may be helpful to describe how these models are used within banks' capital allocation systems. Internal capital allocations against credit risk are based on a bank's estimate of the probability density function (PDF) for credit losses. Credit risk models are used to estimate these PDFs (see chart). A risky portfolio is one whose PDF has a relatively long, fat tail—that is, where there is a significant likelihood that actual losses will be substantially higher than expected losses, shown as the left dotted line in the chart. In this chart, the probability of credit losses exceeding the level X is equal to the shaded area under the PDF to the right of X.

The estimated capital needed to support a bank's credit risk exposure is generally referred to as its "economic capital" for credit risk. The process for determining this amount is analogous to VaR methods used in allocating economic capital against market risks. Specifically, the economic capital for credit risk is determined in such a way that the estimated probability of *unexpected* credit losses exhausting economic capital is less than the bank's "target insolvency rate." Capital allocation systems generally assume that it is the role of reserving policies to cover *expected* credit losses, while it is the role of equity capital to cover *credit risk*, or the *uncertainty* of credit losses. Thus, required economic capital is the amount of equity over and above expected losses necessary to achieve the target insolvency rate. In the chart, for a target insolvency rate equal to the shaded area, the required economic capital equals the distance between the two dotted lines.

In practice, the target insolvency rate is usually chosen to be consistent with the bank's desired credit rating. For example, if the desired credit rating is AA, the target insolvency rate might equal the historical one-year default rate for AA-rated corporate bonds (about 3 basis points).

To recap, economic capital allocations for credit risk are based on two critical inputs: the bank's target insolvency rate and its estimated PDF for credit losses. Two banks with identical portfolios, therefore, could have very different economic capital allocations for credit risk, owing to differences in their attitudes toward risk taking, as reflected in their target insolvency rates, or owing to differences in their methods for estimating PDFs, as reflected in

The Relationship between PDF and Allocated Economic Capital Losses

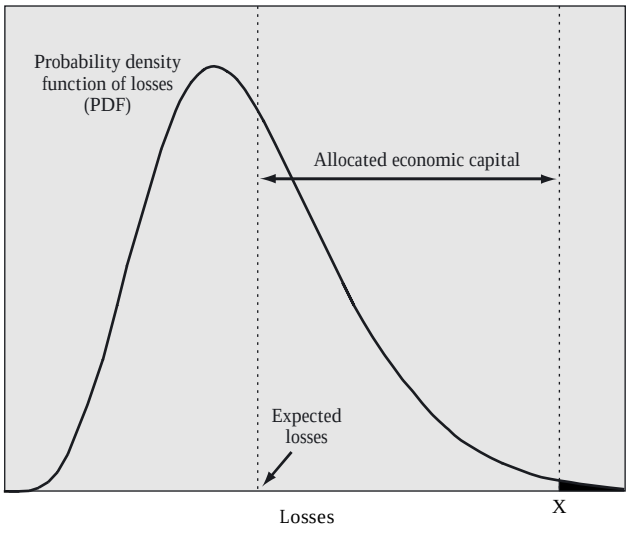

Note: The shaded area under the PDF to the right of X (the target insolvency rate) equals the cumulative probability that unexpected losses will exceed the allocated economic capital.

their credit risk models. Obviously, for competitive equity and other reasons, regulators prefer to apply the same minimum soundness standard to all banks. Thus, any internal models approach to regulatory capital would likely be based on a bank's estimated PDF, not on the bank's *own* internal economic capital allocations. That is, the regulator would likely (a) decide whether the bank's PDF estimation process was acceptable and (b) at least implicitly, set a regulatory maximum insolvency probability (rather than accept the bank's target insolvency rate if such a rate was deemed "too high" by regulatory standards).

III. Types of Credit Risk Models

When estimating the PDF for credit losses, banks generally employ what we term either "top-down" or "bottom-up" methods (see exhibit). Top-down models are often used for estimating credit risk in consumer or small business portfolios. Typically, within a broad subportfolio, such as credit cards, all loans would be treated as more or less homogeneous. The bank would then base its estimated PDF on the *historical* credit loss rates for that subportfolio taken as a whole. For example, the variance in subportfolio loss rates over time could be taken as an estimate of the variance of loss rates associated with the current subportfolio. A limitation of top-down models, however, is that they may not be sensitive to changes in the subportfolio's composition. That is, if the quality of the bank's card customers were to change over time, PDF estimates based on that portfolio's historical loss rates could be highly misleading.

Where changes in portfolio composition are a significant concern, banks appear to be evolving toward bottom-up models. This is already the predominant method for measuring the credit risks of large and middle-market customers. A bottom-up model attempts to quantify credit risk at the level of each individual loan, based on an explicit credit evaluation of the underlying customer. This evaluation is usually summarized in terms of the loan's internal credit rating, which is treated as a proxy for the loan's probability of default. The bank would also estimate the loan's loss rate in the event of default, based on collateral and other factors. To measure credit risk for the portfolio as a whole, the risks of individual loans are aggregated, taking into account correlation effects. Unlike top-down methods, therefore, bottom-up models explicitly consider variations in credit quality and other compositional effects.

Overview of Risk Measurement Systems

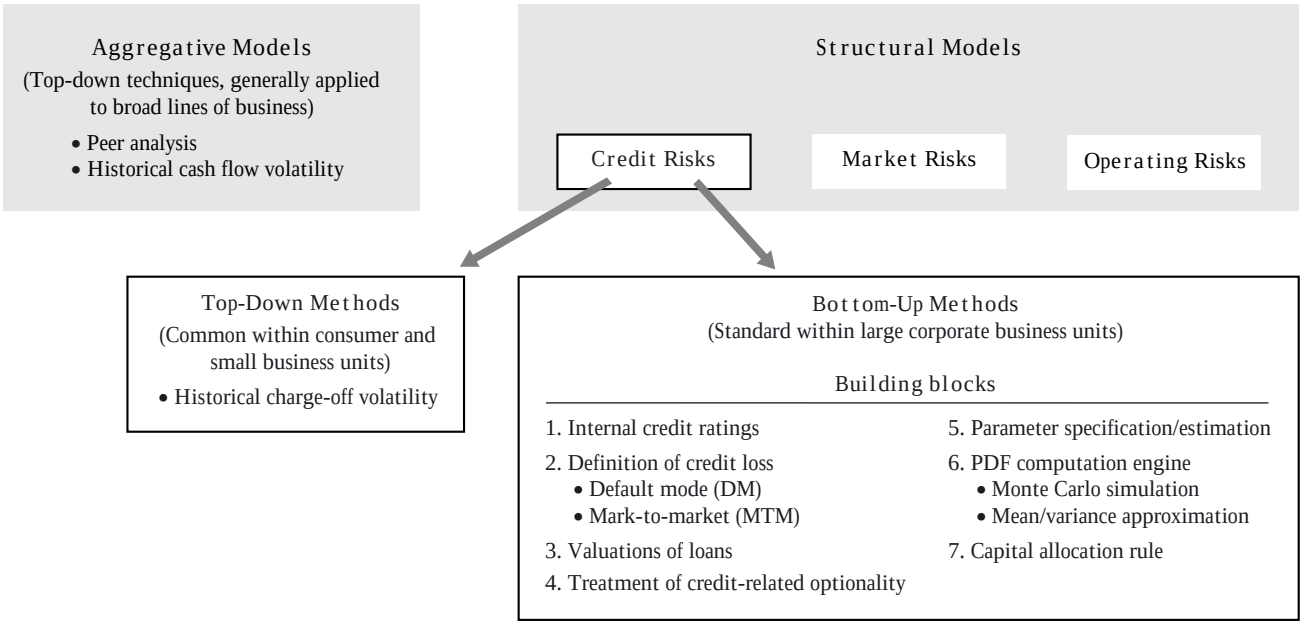

IV. Modeling Issues

The remainder of this summary focuses on four aspects of credit risk modeling: the conceptual framework, credit-related optionality, model calibrations, and model validation. The intent is to highlight some of the modeling issues that we believe are significant from a regulator's perspective; the full version of our paper provides significantly greater detail.

A. Conceptual Framework

Credit risk modeling procedures are driven importantly by a bank's underlying definition of "credit losses" and the "planning horizon" over which such losses are measured. Banks generally employ a one-year planning horizon and what we refer to as either a default-mode (DM) paradigm or a mark-to-market (MTM) paradigm for defining credit losses.

1. Default-Mode Paradigm

At present, the default-mode paradigm is by far the most common approach to defining credit losses. It can be thought of as a representation of the traditional "buy-and-hold" lending business of commercial banks. It is sometimes called a "two-state" model because only two outcomes are relevant: nondefault and default. If a loan does *not* default within the planning horizon, no credit loss is incurred; if the loan defaults, the credit loss equals the difference between the loan's book value and the present value of its net recoveries.

2. Mark-to-Market Paradigm

The mark-to-market paradigm generalizes this approach by recognizing that the economic value of a loan may decline even if the loan does not formally default. This paradigm is "multi-state" in that "default" is only one of several possible credit ratings to which a loan could migrate. In effect, the credit portfolio is assumed to be marked to market or, more accurately, "marked to model." The value of a term loan, for example, typically would employ a discounted cash flow methodology, where the credit spreads used in valuing the loan would depend on the instrument's credit rating.

To illustrate the differences between these two paradigms, consider a loan having an internal credit rating equivalent to BBB. Under both paradigms, the loan would incur a credit loss if it were to default during the planning horizon. Under the mark-to-market paradigm, however, credit losses could *also* arise if the loan were to suffer a downgrade short of default (such as migrating from BBB to BB) or if prevailing credit spreads were to widen. Conversely, the value of the loan could increase if its credit rating improved or if credit spreads narrowed.

Clearly, the planning horizon and loss paradigm are critical decision variables in the credit risk modeling process. As noted, the planning horizon is generally taken to be one year. It is often suggested that one year represents a reasonable interval over which a bank—in the normal course of business—could mitigate its credit exposures. Regulators, however, tend to frame the issue differently—in the context of a bank under *stress* attempting to unload the credit risk of a significant portfolio of deteriorating assets. Based on experience in the United States and elsewhere, *more* than one year is often needed to resolve asset-quality problems at troubled banks. Thus, for the banking book, regulators may be uncomfortable with the assumption that capital is needed to cover only one year of unexpected losses.

Since default-mode models *ignore* credit deteriorations short of default, their estimates of credit risk may be particularly sensitive to the choice of a one-year horizon. With respect to a three-year term loan, for example, the one-year horizon could mean that more than two-thirds of the credit risk is potentially ignored. Many banks attempt to reduce this bias by making a loan's estimated probability of default an increasing function of its maturity. In practice, however, these adjustments are often made in an ad hoc fashion, so it is difficult to assess their effectiveness.

B. Credit-Related Optionality

In contrast to simple loans, for many instruments a bank's credit exposure is not fixed in advance, but rather depends on future (random) events. One example of such "credit-related optionality" is a line of credit, where optionality reflects the fact that drawdown rates tend to increase as a

customer's credit quality deteriorates. As observed in connection with the recent turmoil in foreign exchange markets, credit-related optionality also arises in derivatives transactions, where counterparty exposure changes randomly over the life of the contract, reflecting changes in the amount by which the bank is "in the money."

As with the treatment of optionality in VaR models, credit-related optionality is a complex topic, and methods for dealing with it are still evolving. At present, there is great diversity in practice, which frequently leads to very large differences across banks in credit risk estimates for similar instruments. With regard to virtually identical lines of credit, estimates of stand-alone credit risk can differ as much as a tenfold. In some cases, these differences reflect modeling assumptions that, quite frankly, seem difficult to justify—for example, with respect to committed lines of credit, some banks implicitly assume that future drawdown rates are *independent* of future changes in a customer's credit quality. Going forward, in our view the treatment of credit-related optionality needs to be a priority item, both for bank risk modelers and their supervisors.

C. Model Calibration

Perhaps the most difficult aspect of credit risk modeling is the calibration of model parameters. To illustrate this process, note that in a default-mode model, the credit loss for an individual loan reflects the combined influence of two types of risk factors—those determining whether or not the loan defaults and, in the event of default, risk factors determining the loan's loss rate. Thus, implicitly or explicitly, the model builder must specify (a) the expected probability of default for each loan, (b) the probability distribution for each loan's loss-rate-given-default, and (c) among all loans in the portfolio, *all* possible pair-wise correlations among defaults and loss-rates-given-default. Under the mark-to-market paradigm, the estimation problem is even more complex, since the model builder needs to consider possible credit rating migrations short of default as well as potential changes in future credit spreads.

This is a daunting task. Reflecting the longer term nature of credit cycles, even in the best of circumstances—assuming parameter stability—many years of data, spanning multiple credit cycles, would be needed to estimate default probabilities, correlations, and other key parameters with good precision. At most banks, however, data on historical loan performance have been warehoused only since the implementation of their capital allocation systems, often within the last few years. Owing to such data limitations, the model specification process tends to involve many crucial simplifying assumptions as well as considerable judgment.

In our full paper, we discuss assumptions that are often invoked to make model calibration manageable. Examples include assumptions of parameter stability and various forms of independence within and among the various types of risk factors. Some specifications also impose normality or other parametric assumptions on the underlying probability distributions.

It is important to note that estimation of the extreme tail of the PDF is likely to be highly sensitive to these assumptions and to estimates of key parameters. Surprisingly, in practice there is generally little analysis supporting critical modeling assumptions. Nor is it standard practice to conduct sensitivity testing of a model's vulnerability to key parameters. Indeed, practitioners generally presume that all parameters are known with certainty, thus ignoring credit risk issues arising from parameter uncertainty or model instability. In the context of an internal models approach to regulatory capital for credit risk, sensitivity testing and the treatment of parameter uncertainty would likely be areas of keen supervisory interest.

D. Model Validation

Given the difficulties associated with calibrating credit risk models, one's attention quickly focuses on the need for effective model *validation* procedures. However, the same data problems that make it difficult to calibrate these models also make it difficult to validate the models. Owing to insufficient data for out-of-sample testing, banks generally do not conduct statistical back testing on their estimated PDFs.

Instead, credit risk models tend to be validated indirectly, through various market-based "reality" checks.

Peer-group analysis is used extensively to gauge the reasonableness of a bank's overall capital allocation process. Another market-based technique involves comparing actual credit spreads on corporate bonds or syndicated loans with the break-even spreads implied by the bank's internal pricing models. Clearly, an implicit assumption of these techniques is that prevailing market perceptions and prevailing credit spreads are always "about right."

In principle, *stress testing* could at least partially compensate for shortcomings in available back-testing methods. In the context of VaR models, for example, stress tests designed to simulate hypothetical shocks provide useful checks on the reasonableness of the required capital levels generated by these models. Presumably, stress-testing protocols also could be developed for credit risk models, although we are not yet aware of banks actively pursuing this approach.

V. Possible Near-Term Applications of Credit Risk Models

While the reliability concerns raised above in connection with the *current* generation of credit risk models are substantial, they do not appear to be insurmountable. Credit risk models are progressing so rapidly it is conceivable they could become the foundation for a new approach to setting formal regulatory capital requirements within a reasonably near time frame. Regardless of how formal RBC standards evolve over time, within the short run supervisors need to improve their existing methods for assessing bank capital adequacy, which are rapidly becoming outmoded in the face of technological and financial innovation. Consistent with the notion of "risk-focused" supervision, such new efforts should take full advantage of banks' own internal risk management systems—which generally reflect the most accurate information about their credit exposures—and should focus on encouraging improvements to these systems over time.

Within the relatively near term, we believe that there are at least two broad areas in which the inputs or outputs of bank's internal credit risk models might usefully be incorporated into prudential capital policies. These include (a) the *selective* use of internal credit risk models in setting formal RBC requirements against certain credit positions that are not treated effectively within the current Basle Accord and (b) the use of internal credit ratings and other components of credit risk models for purposes of developing specific and practicable *examination* guidance for assessing the capital adequacy of large, complex banking organizations.

A. Selective Use in Formal RBC Requirements

Under the current RBC standards, certain credit risk positions are treated ineffectually or, in some cases, ignored altogether. The selective application of internal risk models in this area could fill an important void in the current RBC framework for those instruments that, by virtue of their being at the forefront of financial innovation, are the most difficult to address effectively through existing prudential techniques.

One particular application is suggested by the November 1997 Notice of Proposed Rulemaking on Recourse and Direct Credit Substitutes (NPR) put forth by the U.S. banking agencies. The NPR discusses numerous anomalies regarding the current RBC treatment of recourse and other credit enhancements supporting banks' securitization activities. In this area, the Basle Accord often produces dramatically divergent RBC requirements for essentially equivalent credit risks, depending on the specific contractual form through which the bank assumes those risks.

To address some of these inconsistencies, the NPR proposes setting RBC requirements for securitization-related credit enhancements on the basis of credit ratings for these positions obtained from one or more accredited rating agencies. One concern with this proposal is that it may be costly for banks to obtain formal credit ratings for credit enhancements that currently are not publicly rated. In addition, many large banks already produce internal credit ratings for such instruments, which, given the quality of their internal control systems, may be at least as accurate as the ratings that would be produced by accredited rating agencies. A natural extension of the agencies' proposal would permit a bank to use its *internal* credit ratings (in lieu of having to

obtain *external* ratings from accredited rating agencies), provided they were judged to be "reliable" by supervisors.

A further extension of the agency proposal might involve the direct use of internal credit risk models in setting formal RBC requirements for *selected* classes of securitization-related credit enhancements. Many current securitization structures were not contemplated when the Accord was drafted, and cannot be addressed effectively within the current RBC framework. Market acceptance of securitization programs, however, is based heavily on the ability of issuers to quantify (or place reasonable upper bounds on) the credit risks of the underlying pools of securitized assets. The application of internal credit risk models, if deemed "reliable" by supervisors, could provide the first practical means of assigning economically reasonable capital requirements against such instruments. The development of an internal models approach to RBC requirements—on a limited scale for selected instruments—also would provide a useful test bed for enhancing supervisors' understanding of and confidence in such models, and for considering possible expanded regulatory capital applications over time.

B. Improved Examination Guidance

As noted above, most large U.S. banks today have highly disciplined systems for grading the credit quality of individual financial instruments within major portions of their credit portfolios (such as large business customers). In combination with other information from banks' internal risk models, these internal grades could provide a basis for developing specific and practical examination guidance to aid examiners in conducting independent assessments of the capital adequacy of large, complex banking organizations.

To give one example, in contrast to the one-size-fits-all Basle standard, a bank's internal capital allocation against a fully funded, unsecured commercial loan will generally vary with the loan's internal credit rating. Typical internal capital allocations often range from 1 percent or less for a grade-1 loan, to 14 percent or more for a grade-6 loan (in a credit rating system with six "pass" grades). Internal economic capital allocations against classified, but not-yet-charged-off, loans may approach 40 percent—*not* counting any reserves for expected future charge-offs. Examiners could usefully compare a particular bank's actual capital levels (or its allocated capital levels) with the capital levels implied by such a grade-by-grade analysis (using as benchmarks the internal capital allocation ratios, by grade, of peer institutions). At a minimum, such a comparison could initiate discussions with the bank on the reliability of its internal approaches to risk measurement and capital allocation. Over time, examination guidance might evolve to encompass additional elements of banks' internal risk models, including analytical tools based on stress-test methodologies. Regardless of the specific details, the development and field testing of examination guidance on the use of internal credit risk models would provide useful insights into the longer term feasibility of an internal models approach to setting formal regulatory capital standards.

More generally, both supervisors and the banking industry would benefit from the development of sound practice guidance on the design, implementation, and application of internal risk models and capital allocation systems. Although important concerns remain, this field has progressed rapidly in recent years, reflecting the growing awareness that effective risk *measurement* is a critical ingredient to effective risk *management*. As with trading account VaR models at a similar stage of development, banking supervisors are in a unique position to disseminate information on best practices in the risk measurement arena. In additional to permitting individual banks to compare their practices with those of peers, such efforts would likely stimulate constructive discussions among supervisors and bankers on ways to improve current risk modeling practices, including model validation procedures.

VI. Concluding Remarks

The above discussion provides examples by which information from internal credit risk models might be usefully incorporated into regulatory or supervisory capital policies. In view of the modeling concerns described in this summary, incorporating internal credit risk measurement and capital allocation systems into the supervisory and/or

regulatory framework will occur neither quickly nor without significant difficulties. Nevertheless, supervisors should not be dissuaded from embarking on such an endeavor. The current one-size-fits-all system of risk-based capital requirements increasingly is inadequate to the task of measuring large bank soundness. Moreover, the process of "patching" regulatory capital "leaks" as they occur appears to be less and less effective in dealing with the challenges posed by ongoing financial innovation and regulatory capital arbitrage. Finally, despite difficulties with an internal models approach to bank capital, no alternative long-term solutions have yet emerged.

ENDNOTE

The views expressed in this summary are those of the authors and do not necessarily reflect those of the Federal Reserve System or other members of its staff. This paper draws heavily upon information obtained through our participation in an ongoing Federal Reserve System task force that has been reviewing the internal credit risk modeling and capital allocation processes of major U.S. banking organizations. The paper reflects comments from other members of that task force and Federal Reserve staff, including Thomas Boemio, Raphael Bostic, Roger Cole, Edward Ettin, Michael Gordy, Diana Hancock, Beverly Hirtle, James Houpt, Myron Kwast, Mark Levonian, Chris Malloy, James Nelson, Thomas Oravez, Patrick Parkinson, and Thomas Williams. In addition, we have benefited greatly from discussions with numerous practitioners in the risk management arena, especially John Drzik of Oliver, Wyman & Company. We alone, of course, are responsible for any remaining errors.

REFERENCES

Jones, David, and John Mingo. 1998. "Industry Practices in Credit Risk Modeling and Internal Capital Allocations: Implications for a Models-Based Regulatory Capital Standard." Paper presented at the conference "Financial Services at the Crossroads: Capital Regulation in the Twenty-First Century," Federal Reserve Bank of New York, February 26-27.