

Credit Risk Models at Major U.S. Banking Institutions:
Current State of the Art and Implications for Assessments of Capital Adequacy

Federal Reserve System Task Force
on Internal Credit Risk Models*

May, 1998

* Board of Governors: Thomas Boemio, Roger Cole (Co-chairman), David Jones, and John Mingo (Co-chairman); Federal Reserve Bank of New York: Christine Cumming, Thomas Oravez, and Beverly Hirtle; Federal Reserve Bank of Chicago: James Nelson and Edward Waggoner; Federal Reserve Bank of San Francisco: Mark Levonian and Thomas Williams. Other System staff also contributed to this report, including Edward Ettin, Michael Gordy, James Houpt, Myron Kwast, and Chris Malloy (Board of Governors); Claire Desjardins and Gerald Iseda (Federal Reserve Bank of Boston); and Larry Wall (Federal Reserve Bank of Atlanta). The Task Force is also indebted to several consultants and practitioners, including Scott Aguais, John Drzik, Steve Kealhofer, Blythe Masters, and Tom Wilson.

I. Introduction

Nearly a decade has passed since the 1988 Basle Capital Accord established the basic architecture for setting minimum risk-based capital (RBC) requirements for banking organizations (“banks”). The Accord’s overriding objectives, achieved relatively quickly, were both to provide cross-border consistency in capital standards and to increase the capital cushions of the world’s largest banks. Along with this early success has come heightened reliance on capital-based regulatory and supervisory policies. Within the United States, for example, Prompt Corrective Action and other provisions of the FDIC Improvement Act of 1991 now link supervisory and regulatory policies explicitly to banks’ regulatory capital ratios. But, even as the formal RBC ratios have assumed great prominence, ongoing technological and financial innovations have exposed shortcomings in the Basle framework that, if not redressed, could undermine the future role of bank capital standards.

The most significant flaws in the RBC standards have long been recognized. First, the measures of “capital” embodied in the numerators of these ratios may not represent accurately a bank’s capacity to absorb either expected or unexpected losses. Loan loss reserves, for example, often tend to exceed expected credit losses during good times and to understate expected credit losses during times of stress. Second, the denominator of these ratios, total risk-weighted assets, is not an accurate measure of total risk. The regulatory risk-weights do not reflect certain risks, such as interest rate and operating risk. More importantly, they ignore critical differences in credit risk among financial instruments (e.g., all commercial credits incur the same 100 percent risk-weight or, equivalently, an eight percent total RBC requirement), as well as differences across banks in portfolio diversification, hedging activities, and the quality of internal risk management systems.

These anomalies have created substantial opportunities for regulatory capital arbitrage that are rendering the formal RBC ratios increasingly less meaningful for the largest, most sophisticated banks.¹ Through securitization and other financial innovations, many of these

¹ The basic arbitrage techniques involve (a) re-engineering financial contracts to convert a bank’s on-balance sheet credit risk into a nearly equivalent off-balance sheet exposure having a lower capital requirement; and/or (b) removing from the banking book financial instruments for which the 8 percent Basle capital standard is too high, relative to the underlying economic risks, while retaining instruments for which the Basle standard is too low (termed “cherry-picking”).

institutions have lowered their RBC requirements substantially without reducing materially their overall credit risk exposures. More recently, the Market Risk Amendment to the Basle Accord has created additional capital arbitrage opportunities by permitting banks to use their Value-at-Risk (VaR) models for calculating RBC requirements against specific risks within their trading portfolios. Under this Amendment, a bank can reduce its RBC requirement against high-grade credits from 8 percent to much smaller amounts by shifting the risks associated with these assets from its banking book to its trading account, if accounting principles permit.

With the formal RBC ratios rendered less useful, judgmental assessments of capital adequacy through the examination process necessarily have assumed heightened importance. Yet, this process, too, has become more problematic as regulatory capital arbitrage has made credit risk positions less transparent. While examination assessments of capital adequacy normally attempt to adjust reported capital ratios for shortfalls in loan loss reserves relative to *expected* future charge-offs, examiners' tools are limited in their ability to deal effectively with *credit risk* — measured as the *uncertainty* of future credit losses around their expected levels.²

Banks themselves have recognized a need for better methods of risk quantification. Responding to increased domestic and international competition and the greater complexity of their credit portfolios, many of the largest banks have developed sophisticated methods for measuring credit risks. Analogous to trading account VaR models, internal credit risk models are used in estimating the *economic* capital needed to support a bank's credit activities. By design, these systems create strong incentives for managers to economize on a bank's most expensive funding source: equity capital. Internal capital allocations are the basis for estimating the *risk-adjusted* profitability of various bank activities which, in turn, are used in evaluations of managerial performance and in determinations of managerial

² And, of course, no matter how improved their treatment of *credit risk*, minimum capital standards will be ineffective unless loan loss reserves accurately reflect *expected* credit losses. See David Jones and Kathleen Kuester King, "The Implementation of Prompt Corrective Action: An Assessment," *Journal of Banking and Finance*, 19, pp. 491-510 (1995).

compensation. Credit risk models and economic capital allocations also have been incorporated into risk management processes, including risk-based pricing models, the setting of portfolio concentration and exposure limits, and day-to-day credit risk management.

In principle, the inputs or outputs of a bank's internal risk measurement systems could provide valuable information for use in prudential assessments of bank capital adequacy. Potentially, such assessments could be made more incentive-compatible and risk-focused. Enhancements to banks' risk management systems might translate more quickly into improved prudential policies, perhaps reducing incentives to "game" supervisors via the channeling of credit risk through opaque, off-balance sheet transactions. The use of internal risk estimates by supervisors also might promote more rapid development of improved risk management techniques, and faster convergence toward a common risk measurement framework, which could lead to improved risk disclosures (greater transparency). Such potential benefits, of course, are predicated on the reliability of internal risk models.

The System Task Force on Internal Credit Risk Models was created in April, 1996, to assess potential uses of banks' internal credit risk and capital allocation models within the supervisory process. As instructed by the Board's Committee on Supervisory and Regulatory Affairs, the Task Force was assigned two broad objectives:

1. Review the credit risk modeling practices of large U.S. banks; and
2. Develop suggestions, if practical, for incorporating sound practice techniques into the supervisory processes for assessing capital adequacy and the quality of credit risk management (as part of the overall supervisory rating for management quality).

In fulfillment of the first objective, this report surveys the state-of-the-art in the design, implementation, and uses of internal credit risk models.³ The report also presents preliminary conclusions regarding potential regulatory and supervisory applications of credit risk models.

³ The study involved extensive discussions with twelve banking organizations, two nonbank securities firms, and several credit risk consultancies. Task Force members also conducted a literature review, attended conferences dealing with credit risk modeling, and benefitted from many informal discussions with practitioners.

Given time and resource constraints, there was a compelling practical need for the Task Force to limit the scope of the study. As will be often repeated below, the Task Force observed great diversity among the risk modeling and economic capital allocation practices of the sampled banks, especially in terms of implementation details. Even within a particular bank, practices were often quite variable across business or product lines. Since it was not possible to examine in detail all risk modeling and economic capital applications within the sampled banks, the Task Force elected to focus its attention on the measurement of credit risk for large- and middle-market business customers (the “large corporate business”). This decision reflected the Task Force’s view, shared by practitioners, that credit risk modeling and economic capital allocation processes were the most highly developed and widespread for this customer segment.

Based on its review of credit risk modeling practices at major banks, the Task Force has included in this report several preliminary conclusions regarding potential near-term applications of credit risk models within the regulatory and supervisory processes. With regard to formal regulatory capital requirements for credit risk, the Task Force believes that while improvements in credit risk modeling are occurring rapidly, a number of important challenges must be addressed before adoption of an internal models approach to RBC for the *banking book* would be possible -- as a replacement for the Basle Accord. Foremost are issues relating to (a) the appropriate conceptual framework for measuring “credit losses,” (b) the calibration of key model parameters, and (c) the need for more systematic approaches to model validation, including explicit treatment of model uncertainty/instability. While similar issues are relevant to market and specific risk models for the *trading account*, the magnitude of these concerns is much greater with respect to credit risk models for the *banking book*. In addition to model reliability issues, the adoption of an internal models approach to RBC for credit risk may require a more explicit treatment of operating risk within the regulatory capital framework.

In light of the ongoing progress in credit risk modeling techniques, it is conceivable that further improvements in the near future could redress many, if not most, of the concerns raised by the Task Force. However, as traditional techniques for assessing capital adequacy

are rapidly becoming outmoded, improved *supervisory* methods are needed if capital-based prudential policies are to remain viable even over the shorter term. Because the most accurate information regarding a bank's risks is likely to reside within its own internal risk measurement and management systems, clearly supervisors should utilize this information to the extent possible.

To this end, this report outlines several possible near-term uses of internal credit risk models that, while not a full replacement for the Accord, could nevertheless enhance current prudential policies. Specifically, the Task Force believes these models may be useful in two roles: (1) the development of specific and practical *examination* guidance for assessing the capital adequacy of large, complex banks; and (2) the setting of *regulatory* capital requirements against *selected* instruments that have largely evolved subsequent to the adoption of the Accord, such as credit enhancements supporting securitization programs.

The remainder of this report is organized as follows. To provide a context for later discussions of credit risk models, sections II and III present general overviews of internal capital allocation and risk measurement systems at large U.S. banks. Sections IV and V then discuss the conceptual framework and empirical techniques underpinning the current generation of credit risk models. Following these descriptions of the current state-of-the-art, section VI raises issues that may be of concern to supervisors. Lastly, section VII discusses possible near-term regulatory and supervisory uses of internal risk models.

II. Overview of Economic Capital Allocation Systems

As used in this paper, the term “risk model” refers to all of the procedures employed by a bank to quantify its economic risks, whether with respect to a single transaction or to a group of transactions, customers, or products. Such estimates are used internally to allocate “economic capital” (defined below) to activities based on their estimated contributions to the bank’s total risk -- of which credit risk is one component. To provide a context for later discussions of risk measurement practices, this section describes in general terms the structure of economic capital allocation systems and their roles within the decision making processes of major U.S. banks. Risk modeling practices, per se, will be a central topic of the remaining sections.

A. Economic Capital: Conceptual Framework and Implementation

The estimated amount of capital needed to support a bank’s risk taking activities is typically termed the bank’s required or allocated “economic capital.” While internal systems for allocating economic capital typically encompass all forms of risk facing a bank (credit, market, and operating risks), our principle focus is on the allocation of economic capital for credit risk.

Systems for allocating economic capital against credit risk are based on a bank’s estimate of the “probability density function” for credit losses (“PDF”). As illustrated in Exhibit 1, an important property of PDFs is that, for a hypothetical level of losses denoted by X, the estimated probability of actual credit losses exceeding this level is equal to the (shaded) area under the PDF to the right of X. As discussed in Section IV of this report, the precise definitions of “credit loss” tend to vary across banks depending on the conceptual frameworks underlying their risk measurement and management systems (e.g., a mark-to-market paradigm). For a particular definition, a risky portfolio, loosely speaking, is one whose PDF has a relatively long, fat tail -- that is, where there is a greater likelihood that losses will be substantially higher than expected losses (shown as the left-most dotted line). Although in this section we treat the PDF as given, later sections of the paper discuss in detail the credit risk modeling techniques used in estimating PDFs.

Risk measurement systems generally “collapse” the estimated PDF into a single metric, termed the “economic capital” allocation for credit risk. This process is analogous to the VaR methods used in allocating economic capital against market risks. Specifically, the economic capital allocation is determined in theory so that the probability of unexpected credit losses exhausting economic capital (i.e., the probability of insolvency) is less than some targeted level. For instance, the level of economic capital may be set to achieve a 0.03 percent estimated probability that unexpected credit losses would exceed this level, thereby causing insolvency. The target insolvency rate (in this example, tantamount to a 99.97% confidence interval) usually is chosen to be consistent with the bank’s desired credit rating for its liabilities: if the desired credit rating is AA, the target insolvency rate might be set at the historical one-year default rate for AA-rated corporate bonds (about 3 basis points). Economic capital, therefore, is a way of representing relative degrees of risk; a portfolio is relatively “risky” if its required or allocated economic capital is high per dollar of asset value.

Within economic capital allocation systems, a critical distinction is made between *expected* credit losses and the *uncertainty* of credit losses (i.e., *credit risk*). These 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. In Exhibit 1, therefore, the area under the PDF to the left of expected losses should be covered by the loan loss reserve, while the bank’s required economic capital is the amount of equity over and above expected losses necessary to achieve the target insolvency rate. Under this framework, a bank would consider itself to be undercapitalized if its required economic capital exceeded its actual tangible shareholder equity, adjusted for any estimated surplus in the bank’s reported loan loss reserve.⁴

⁴ For analytical purposes, practitioners generally apportion the reported loan loss reserve among “specific” and “general” components. Specific reserves represent adjustments to reported book values of loans to bring the net carrying amount of the loan portfolio into alignment with its estimated or expected true underlying value. Since specific reserves are *not* viewed as available to absorb unexpected credit losses, they are *not* considered a form of equity or economic capital. In contrast, general reserves *are* available to absorb unexpected losses and, therefore, *are* treated as equity or economic capital for risk measurement and management purposes.

Most banks allocate economic capital not only against their entire portfolios, but also against specific activities, ranging from individual transactions to entire business lines. Conceptually, the economic capital allocated to a particular activity should be measured as that activity's marginal contribution to the portfolio's overall economic capital requirement, taking into consideration diversification effects between that activity and the rest of the bank. Operationally, this marginal contribution would be calculated as the economic capital allocation for the entire portfolio *less* the hypothetical capital allocation for the portfolio *excluding* the activity of interest.

Another noteworthy feature of most banks' internal capital allocation systems is the practice of allocating economic capital for different types of risk (e.g., credit, operating, and market risks) more or less independently of one another. That is, a separate PDF is estimated for each type of risk, against which economic capital is allocated. The total economic capital allocation for the bank as a whole is then computed as the summation of the individual capital allocations for each risk. Banks are aware that this piecemeal approach is not strictly consistent with their underlying portfolio approach to risk measurement -- unless credit, market, and operating risks are perfectly correlated. Nevertheless, given the current infeasibility of estimating inter-relationships (e.g., cross-correlations) among different types of risk, this approach is viewed as a practical necessity. Moreover, since the separate risks are certainly less than perfectly correlated in practice, the resultant capital allocations are generally believed to be conservative estimates of the overall capital needed to achieve the bank's target insolvency rate.⁵

B. Internal Uses of Allocated Economic Capital

The banks reviewed by the Task Force utilize internal economic capital allocations for two broad purposes: (a) measuring *risk-adjusted* profitability, and (b) portfolio risk

⁵ Some banks allocate capital by estimating separate risk models for different lines of business (i.e., credit cards, other consumer lending, small business loans, and large corporate loans) and then allocating capital to each business on the basis of its stand-alone risk. This approach, too, is viewed as conservative.

management.

1. Measuring risk-adjusted profitability. At most banks, the primary motivation for developing an internal capital allocation system was a perceived need for internal control mechanisms to ensure that shareholders' equity was used as efficiently as possible. Allocations of economic capital are viewed as an acceptable method, within the context of the bank's existing corporate and credit cultures, for imputing to each activity an internal transfer cost or "usage charge" for the amount of shareholder equity required to support the underlying risks. In practice, economic capital tends to be allocated to activities at several levels of aggregation, ranging from discrete transactions (e.g., a specific loan or credit facility), to individual customer-relationships, to entire business divisions.⁶

Typically, risk-adjusted profits are measured by adjusting traditional cost-accounting measures of net income to reflect the opportunity cost of the equity needed to support the activity (i.e., its economic capital allocation). Specifically, risk-adjusted profits would be calculated as revenues allocated to the activity, *less* the cost of allocated debt, *less* allocated non-interest expense (including an expense for expected credit losses), *less* the cost of allocated economic capital.⁷ The cost of economic capital is defined as the activity's allocated economic capital *times* the bank's ROE target or hurdle rate.⁸

Most of the sampled banks use this approach, or some variant, to represent the amount of shareholder value created by an activity over the time period of analysis -- that is, the earnings generated over and above the opportunity rate of return on shareholder equity (allocated economic capital). Frequently, these measures of risk-adjusted profits are termed

⁶ To some degree, all the sampled banks use economic capital allocations for purposes of measuring risk-adjusted profitability. However, at one organization, internal capital allocations were used only within its large corporate business unit.

⁷ For purposes of calculating risk-adjusted profits, economic capital reflects all forms of risk -- credit, market, and operating risk. As noted above, in practice, the total economic capital for an activity is generally calculated as the simple summation of the separate allocations for each type of risk.

⁸ Although inconsistent with standard finance theory, a single ROE target generally is employed for all businesses lines.

“shareholder value added (SVA)” or “economic value added (EVA).” In this fashion, risk-adjusted profits for various activities can be placed on an apples-to-apples basis. Armed with this information, managers can make informed strategic decisions about how to allocate scarce resources -- that is, in determining which activities to increase in size or scope, which to cut back, which to eliminate. Increasingly, measurements of SVA are being used by banks in gauging managerial or product performance and in determining managerial compensation.⁹

2. Portfolio risk management. All the sampled banks employed their credit risk models and capital allocation processes for various risk management purposes. Internal credit ratings for individual customers (see below) are used in setting customer lending limits and in summarizing portfolio credit quality in reports to senior managers and boards of directors. Some institutions express customer and industry concentration limits in terms of the economic capital at risk, as opposed to notional amounts of loans and undrawn lines of credit. Also, the total economic capital allocated across all of a bank’s business lines was cited by some banks as one of many inputs into their processes for determining the adequacy of their overall capital structures.

Credit risk models and economic capital allocation processes also are used in active risk management, both at the level of individual transactions and at the level of the overall portfolio. When setting the price on a proposed new loan facility, it is now fairly common for a banker first to determine the break-even interest rate needed to cover the loan’s *expected losses* and an appropriate margin for *credit risk* or *unexpected losses* -- determined so that the expected rate of return on the capital allocated to the loan (the Risk-Adjusted Return on Capital, or RAROC) achieves the bank’s hurdle rate.¹⁰ If other market participants are

⁹ The SVA concept described above measures an activity’s “excess cash flows” above and beyond the bank’s relevant opportunity costs for a *single* period. To mitigate potential distortions created by such a myopic measure of performance, banks typically use multi-year, rolling averages of SVA when evaluating *ex post* profitability or managerial performance.

¹⁰ For loan pricing purposes, the RAROC would be calculated as the loan’s estimated *gross SVA* (before the cost allocation for equity capital) *divided by* the loan’s allocated economic capital.

charging a lower interest rate on such loans than is necessary to meet the bank's RAROC hurdle rate, the banker may decline and send the loan business elsewhere. Or, the banker may treat the loan as a "loss-leader" and hope to make up the difference via other, non-loan business with that customer.¹¹

Quite apart from determining appropriate risk-based pricing on individual loans, some banks use credit risk models in active portfolio management. To give one example, credit risk models may be used to estimate an efficient portfolio "frontier," defined as all feasible combinations of the mean and variance of portfolio rate-of-return showing, for a given mean, the lowest achievable variance. By comparing this frontier with the mean and variance of the actual portfolio, risk managers are able to develop strategies for altering the current portfolio to achieve a more preferred risk-return profile. This might be accomplished, for example, by modifying the pattern of new loan originations, by buying/selling loans in secondary markets, or by undertaking credit derivative transactions to lay off (or acquire) various credit exposures. As the markets for credit derivatives and secondary loan trading continue to become more liquid, practitioners expect the use of credit risk models in active portfolio risk management to increase over time.

III. Approaches to Risk Measurement: Aggregative vs. Structural Models

Among the largest U.S. banks, there is great diversity in risk modeling practices. To provide a taxonomy for discussion purposes, we have divided risk measurement approaches into two broad categories: "aggregative" models and "structural" models, illustrated in Exhibit 2. Aggregative approaches to risk measurement attempt to infer *total risk* (i.e., the

¹¹ Several respondents indicated that RAROC pricing models tend to generate required risk premiums on commercial loans that are above those observed in the market. That is, on a risk-adjusted basis, stand-alone commercial lending appeared unprofitable, reflecting competitors' rates that were lower than implied by the bank's internal RAROC model. When this happens, the front-line loan officer or relationship manager, after discussions with senior officers, usually is permitted to make an "exception" to the model's implied pricing in order to accommodate the customer-relationship -- provided the overall relationship generates a RAROC exceeding the bank's hurdle rate.

sum of credit, market, and operating risks) directly from the capital ratios of competitors or from the historical volatility of the cash flows associated with an activity. Structural approaches, on the other hand, estimate total risk through a multi-step process encompassing separate models for credit, market, and operating risk. This section presents an overview of these alternative methodologies.

A. Aggregative Models

Aggregative models typically are “top-down” approaches that attempt to infer the *total risk* of a broadly defined business or product line using peer analysis or historical cash flow analysis. Peer group or “market comparables” analysis attempts to estimate the capital that would be needed to achieve a hypothetical “target” credit rating for a given activity (as if operated on a stand-alone basis) from the capitalization rates of competitors engaged in that activity. Typically, this approach is applied only to complete lines of business or broad product groupings (e.g., credit cards), for which data on publicly traded competitors are readily available. To better ensure apples-to-apples comparisons, capital ratios within the peer group usually are adjusted to remove the estimated effects of (a) accounting distortions, such as securitization, (b) disparities between the bank’s desired credit rating (e.g., AA) and the actual credit ratings of peers, and (c) broad differences in portfolio composition (e.g., variations in the relative sizes of consumer versus C&I lending).

The other major aggregative technique, historical cash flow analysis, attempts to estimate an activity’s total risk from the volatility of its historical cash flows. Implicitly, historical cash flow volatility (per dollar of notional size) is assumed to equal future volatility. To minimize implementation costs to the bank, the underlying cash flow estimates generally are constructed from raw data already generated within the bank’s management information systems, again, usually for broad product groupings. Adjustments normally are applied to these raw data so that the cash flow for a period (e.g., a quarter or a year) can be interpreted as an approximation to the activity’s economic earnings, sometimes termed Net Operating Profit After Tax or “NOPAT.” Given a time series of historical NOPAT, the total risk of an activity (per dollar of notional size) is generally represented by the standard

deviation of the historical ratios of NOPAT to notional size.

While aggregative models for allocating economic capital are quite common among nonfinancial firms for which operating risks predominate, they are less prevalent among banks, which are affected more significantly by credit and market risks. Among banks, aggregative models tend to be used mainly for assessing the performance of broad business or product lines, for making large-scale strategic business decisions (such as acquisitions or divestitures), or for validating structural risk models, rather than for day-to-day investment and risk management purposes. Only one of the sampled banks used aggregative techniques as its sole approach to measuring risk across all business lines.

This pattern of usage reflects two perceived limitations of aggregative models. First, as noted above, data availability often makes it difficult to apply these models at the level of individual transactions or customer relationships (e.g., in product pricing decisions). A second drawback is these models' relative insensitivity to variations in portfolio composition within the business lines that are separately analyzed. Peer analysis, for example, may be misleading if the credit quality of a bank's portfolio differs significantly from that of its competitors. Similarly, the historical cash flow approach may be inappropriate if the current composition of the bank's portfolio (e.g., its sectoral make-up or the credit quality of the underlying customers) is substantially different from that historically.

B. Structural Models

In contrast to aggregative models, structural modeling approaches estimate total risk through the separate modeling of credit, market, and operating risks. With respect to the modeling of credit risk, most banks use multiple modeling approaches within the organization. Quite often, a bank may employ "top-down" approaches in certain lines of business (e.g., consumer or small business lending), and "bottom-up" approaches in others (e.g., large corporate customers).

Top-down models are frequently used for estimating credit risk in consumer or small business loan portfolios. Within a broad sub-portfolio (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 losses for that sub-portfolio taken as a whole. In some cases, to obtain more accurate estimates, a bank may pool its own historical credit loss experience with those of peers (derived from public financial statements and Call Reports). Top-down credit risk models generally are vulnerable to the same concerns as top-down aggregative models -- namely, the portfolio's current quality and composition may differ substantially from those historically.

Where changes in portfolio composition are a significant concern, banks appear to be evolving toward "bottom-up" approaches to credit risk modeling. This is already the predominant method for measuring the credit risks of large and middle-market customers. Unlike top-down methods, bottom-up models explicitly consider variations in credit quality and other compositional effects. A bottom-up modeling process attempts to quantify credit risk at the level of each individual credit facility (e.g., a loan or a line of credit) based on an explicit evaluation of the financial condition of the underlying customer and the structure of the credit facility. As described in the next section, for a simple term loan, this evaluation is often 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 attempt to estimate the probability distribution of the loan's loss-rate in the event of default, based on the instrument's seniority, collateral (if any), and other factors. To measure credit risk for the portfolio as a whole, the risks of individual loans are aggregated, taking into account diversification/correlation effects.

The remainder of this report focuses primarily on the design and implementation of bottom-up credit risk models. It is within this arena that the banking industry is expending significant effort and is making the greatest conceptual advances in credit risk modeling. Before discussing the modeling techniques employed by the sampled banks, below we highlight several common infrastructure issues that generally must be addressed before implementing a (bottom-up) credit risk modeling system.

1. Management information systems. All the sampled banks using structural credit risk models emphasized that substantial development work first had to be undertaken to

create the supporting management information systems (MIS). At an early stage, significant enhancements to, and investments in, the bank's MIS were required to collect, process, and disseminate in a timely manner the information needed to calculate the risk exposures associated with each individual credit facility, cutting across all of the bank's product, geographic, and business lines. A crucial enabling factor has been the dramatic improvements in computing and communications technologies over the past decade, without which the required enhancements to MIS would have been technically or economically infeasible.

Improved MIS, quite apart from its critical role in developing effective credit risk models, has had significant stand-alone benefits. Chief among these at many banks, according to practitioners, is the ability for the first time to measure aggregate credit exposures. Previously, information systems typically were designed around the accounting needs of the banks, rather than around the requirements for an effective risk measurement and management system.

As with all aspects of the credit risk modeling process, banks differed considerably in the implementation details regarding what specific information was used in their credit risk measurement systems; how this information was retrieved, transmitted, and processed; and where computations were performed (e.g., within individual business units or within centralized credit risk management departments). Often, legacy MIS and the bank's pre-existing corporate and credit cultures were the overriding factors driving these decisions.

2. Internal credit rating systems. Regarding the key components of MIS needed to support bottom-up credit risk models, all of the sampled banks emphasized that a reliable internal credit rating system was absolutely critical. Although an in-depth review of internal credit rating systems is beyond the scope of this paper, some knowledge of these systems is essential for understanding the structure of credit risk models.¹²

¹² See also *Winning the Credit Cycle Game: A Road Map for Adding Shareholder Value through Credit Portfolio Management*, a best-practice survey of credit portfolio management techniques conducted by the Robert Morris Associates in conjunction with the First Manhattan

Bankers have long recognized that “knowing your customer” is the first-line of defense against credit losses. Not surprisingly, therefore, all of the banks reviewed by the Task Force assign “credit ratings” to each large- and middle-market customer, as well as to each customer’s separate credit facilities — defined to include all on- and off-balance sheet credit exposures. Internal credit rating systems are designed to differentiate the credit quality of borrowers much more finely than under the five-point grading scale used by bank examiners (i.e., pass, special mention, substandard, doubtful, and loss). All of the internal credit rating systems reviewed by the Task Force used at least four pass grades, in addition to the four supervisory categories of “criticized” credits. The most typical configuration included six “pass” grades plus the four criticized grades, while the largest number of pass grades was 18.¹³

In general, the process of arriving at a credit rating for a customer or facility can be described as containing one or more of the following three elements: (a) the traditional “spreading of numbers” in which financial and other characteristics of the customer (e.g., country and SIC code) and specific features of the facility (e.g., maturity) are incorporated into a relatively subjective approach to determining grades; (b) the use of vendor-supplied commercial credit scoring models;¹⁴ or (c) the use of internally developed credit scoring models. None of the banks relied solely on formal credit scoring models, whether developed internally or externally. Indeed, some banks used no scoring models whatsoever (for commercial credits), relying essentially on subjective guidelines established by their credit

Consulting Group, September 1997.

¹³ A prominent feature at many banks was a “clustering” of loans within just two grades (i.e., 70% of the bank’s credit exposures to large corporates might occur within grades 4 and 5 in a 10-grade system). At most of these institutions, this clustering has led senior management to question the “resolution power” of their credit rating systems and to begin to develop more finely tuned “buckets” within those categories housing the bulk of their credits.

¹⁴ In addition to traditional credit scoring models, for firms whose shares are publicly-traded, several banks in the sample purchase from KMV Corp. direct predictions of firms’ default probabilities. Unlike traditional credit scoring techniques, the KMV methodology infers default probabilities from publicly-traded equity prices using an option-theoretic approach.

departments. To ensure consistency and discipline within the rating process, however, nearly all the sampled banks required credit ratings to be reviewed by internal auditing departments (or their equivalent) that were independent of the banks' business units.

For risk measurement purposes, the importance of the credit rating process derives from the fact that, within most credit risk models, the internal credit risk grade is treated as a "sufficient statistic" for summarizing a facility's probability of defaulting within the relevant planning horizon. This translation is often accomplished through a two-step process involving, first, the construction of a concordance table relating the bank's internal credit grades to some external rating standard, usually S&P's or Moody's ratings for corporate bonds. A grade-1 loan may be deemed roughly equivalent to an S&P bond rating from AA to AAA, a grade-2 loan equivalent to a bond rating of single-A, and so on. Given this concordance, the probability of a customer defaulting on its obligations (or migrating to another credit risk grade) is usually inferred from (a) published tables of the historical default frequencies of similarly-rated corporate bonds, (b) any available internal data on the historical default rates of loans originated by the bank itself, and/or (c) consultants' knowledge of the default rates experienced by other banks.

Internal credit ratings play an important role not only as a "first step" in the credit risk measurement process but also as an important stand-alone risk management tool. Credit ratings within the large corporate businesses are a basis for regular risk reports to senior management and boards of directors. Modern MIS permit the "slicing and dicing" of this portfolio along many dimensions, such as by credit rating, industrial sector of the obligor, geographic location, and so forth. Credit ratings are also the basis for a continuous loan review process, under which large corporate credits generally are reviewed and re-graded at least annually in order to focus attention on deteriorating credits well before they become "criticized" by examiners or external auditors.

IV. Bottom-up Credit Risk Models: Conceptual Building Blocks

This section describes the conceptual frameworks underpinning (bottom-up) credit risk models, including (a) the choice of "planning horizon" and definition of "credit losses,"

(b) the determinants of loan values, and (c) the treatment of credit-related optionality. Following this discussion, section V then discusses the empirical methods used when implementing these models. While the discussions in both sections are somewhat technical, we believe they are useful in conveying the conceptual and practical issues faced by model-builders and the range of assumptions and judgments required throughout the modeling process. Nevertheless, this material is not essential to the policy-related discussions in the final two sections of the paper, and can be skipped if so desired.

A. Planning Horizon and Definition of Credit Losses

Credit risk modeling procedures are shaped 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 for analysis and what we shall refer to as either a Default-Mode (DM) or a Mark-To-Market (MTM) paradigm for defining and measuring credit losses. The DM paradigm embraces the notion that credit losses can arise only if a loan defaults during the planning horizon, while the MTM paradigm adopts the broader economic perspective that credit events short of default may generate declines, or increases, in economic value.

1. Default-Mode Paradigm. At present, the DM paradigm is the most common approach used by banks for defining credit losses. It is sometimes called a "two-state" model because only two outcomes are relevant: non-default and default. If the loan does not default, there is no credit loss. If the loan defaults, however, there generally *is* a credit loss, equal to the present value of the difference between the customer's contractual obligations and the loan's actual net cash flows over the workout period (e.g., recoveries less workout costs).¹⁵ The loss rate per dollar of initial value is generally represented in terms of the loan's loss-rate-given-default (LGD), which is usually treated as a random variable whose

¹⁵ Precisely what constitutes a "default" varies somewhat across banks. Normally, a default arises if the obligor becomes unable to meet its payment obligations or if the loan is placed on "non-accruing" status.

value (in the event the loan defaults) is uncertain as of the beginning of the planning horizon.¹⁶

The DM paradigm can be thought of as a representation of the traditional “buy and hold” lending business of commercial banks. Under this view, secondary loan markets are not sufficiently developed to support a full mark-to-market or trading approach to risk measurement. Note that if all credit exposures had a one-year maturity (equal to the planning horizon), the DM paradigm could, at least conceptually, account for all potential credit losses within the portfolio. For instruments having an effective maturity exceeding one year, however, the DM framework potentially ignores credit losses associated with defaults beyond the planning horizon.

Some banks attempt to attenuate this problem by adjusting an instrument’s credit rating for its maturity (tantamount to adjusting its probability of defaulting within the one-year planning horizon). That is, a longer-term loan would be assigned a lower credit rating (higher probability of default) than a short-term credit to the same customer. In practice, these adjustments tend to be *ad hoc*; for example, key parameters other than default probabilities, such as correlations among loan defaults, generally are *not* similarly adjusted for maturity. As a consequence, it is difficult to assess the overall impact and effectiveness of such adjustments.¹⁷

2. Mark-To-Market Paradigm. The MTM paradigm generalizes the DM approach by recognizing that the economic value of a credit instrument may decline even if the

¹⁶ Since the actual net cash flows on a defaulted loan are not known until the workout is completed, there is uncertainty surrounding the LGD even at the time of default.

¹⁷ In reflection of these difficulties, most banks employ different planning horizon assumptions within different applications. For purposes of performance measurement and portfolio risk management, a one-year horizon is standard. Within banks’ risk-based pricing models, RAROC calculations generally consider *expected* credit losses over the entire life of the loan, while the loan’s economic capital allocation for *unexpected* losses (the denominator in the RAROC calculation) is estimated using a one-year planning horizon.

counterparty does not formally default within the planning horizon.¹⁸ While few banks currently use the MTM framework outside their trading accounts, many practitioners believe the industry is likely to evolve from largely DM-based risk models for the banking book to the more general MTM-based models over the coming years. The MTM model is “multi-state” in that “default” is only one of several possible credit rating grades to which the instrument could migrate over the planning horizon. In effect, the credit portfolio is assumed to be marked to market, or, more accurately, “marked to model.” As discussed below, the value of a term loan, for instance, typically would employ a discounted cash flow methodology, where the discount rates used in valuing the loan would reflect the market-determined term structure of credit spreads for loans of that grade. A credit loss under the MTM paradigm is defined as an unexpected reduction in the portfolio’s value over the planning horizon due to either deteriorations in credit ratings on the underlying loans or a widening of credit risk spreads in financial markets.¹⁹

To illustrate the differences between the DM and MTM paradigms, consider a loan having an internal credit rating equivalent to BBB. Under both paradigms, the loan could lose value if it were to default during the planning horizon. In this event, the rate of credit loss would be represented by the loan’s LGD. Under the MTM paradigm, however, credit losses also could arise if the loan were to suffer a downgrade short of default (e.g., migrate from BBB to BB), or if credit risk spreads prevailing in financial markets were to increase over the planning horizon. Conversely, the value of the loan could increase (implying an

¹⁸ An MTM-type methodology for which extensive public documentation exists is J.P. Morgan’s CreditMetrics™. To provide additional background, some of the discussions below reference the CreditMetrics methodology and its related technical documentation in Greg M. Gupton, Christopher C. Finger, and Mickey Bhatia, *CreditMetrics -- Technical Document*, New York: Morgan Guaranty Trust Co., April 1997.

¹⁹ Unlike the DM model, within the MTM framework the overall portfolio’s credit loss is not a simple summation of the losses across the individual assets. Portfolio losses almost always will be less than this sum, as economic gains on some instruments due to favorable credit-related events (e.g., a rating upgrade) normally will offset at least some credit-related losses on other instruments.

economic gain, rather than a loss) if its credit rating improved or if credit spreads narrowed.

B. Valuation Models

Under both the DM and MTM loss paradigms, the model-builder is required to specify precisely how the current and future values of each credit instrument are determined at the beginning and end of the planning horizon, respectively. Below, we first discuss the loan valuation process within the more general MTM framework; the DM framework is then treated as a special-case. (The Appendix presents a more formal mathematical description of the general structure of MTM-type models.) To simplify the exposition, it is assumed that the credit portfolio consists only of fixed-rate, term loans, and that each customer has only a single loan.

1. Current values of loans (MTM-type models). The current value of a loan typically is represented as the present discounted value of its *contractual* cash flows. The interest rates used for discounting contractual cash flows reflect (a) the term structure of risk-free interest rates implied by the Treasury yield curve, and (b) for each credit rating grade, the market-determined term structure of credit spreads associated with obligors of that grade. This specification assumes that, apart from idiosyncratic random effects that wash out in the aggregate, credit risk spreads depend only on an instrument's credit rating (i.e., its probability of default).²⁰

2. Future values of non-defaulted loans (MTM-type models). Consistent with the determination of current values, the future value of a non-defaulting loan (as of the end of the planning horizon) would be calculated as the present discounted value of its remaining

²⁰ Note that this specification implies that a senior loan and subordinated loan to the same borrower would have identical market values if they called for the same *contractual* payments. In theory, the credit spreads applicable to a loan's *contractual* payments should reflect not only the loan's probability of default (i.e., its credit rating), but also the mean and volatility of its LGD. That is, financial markets would *not* apply the same credit spreads to loans having identical probabilities of default but sharply different expected loss rates in the event of default.

contractual cash flows.²¹ The discount rates can be different from those used in computing the loan's value at the beginning of the planning horizon either because the loan's credit rating may have changed or because the term structure of credit spreads on loans of a given rating may have changed.

3. Future values of defaulted loans (MTM-type models). One of the rating grades to which a loan can migrate over the planning horizon is "default." Obviously, banks do not rely on the discounting of *contractual* cash flows for modeling the present values of *defaulted* loans. Rather, the decline in the economic value of a defaulted loan is typically represented in terms of its LGD. Usually, LGDs are assumed to depend on an instrument's seniority and collateral type plus a random "risk factor."

4. Loan valuations within DM-type models. Under the DM paradigm, a loan's current value equals its book amount. Recall that only two future scenarios are relevant for valuing a loan at the end of the planning horizon: default and non-default. The future value of a non-defaulting loan also is taken to be its book amount, while the decline in value of a defaulting loan is given by the loan's book value times its LGD, as is generally the case in MTM models. (Neither changes in credit risk spreads nor downgrades short of default affect future loan values under the DM paradigm.)

C. Credit-Related Optionality

In contrast to simple loans, for many types of credit instruments a bank's credit exposure (i.e., the actual amount advanced to the customer) is not known with certainty, but rather may depend on the occurrence of future random events. One example of such "credit-related optionality" is a committed line of credit where, for a fixed period of time, a bank agrees to advance funds (up to a pre-defined credit limit) at the customer's discretion. A general characteristic of such lines is that a customer's draw-down rate (measured per dollar

²¹ For simplicity, principal and interest payments received during the planning horizon are assumed to be invested in (non-interest-bearing) cash.

of credit limit) tends to increase as the customer's credit quality deteriorates, reflecting the reduced availability or higher costs of alternative sources of funding.²²

Within the current generation of credit risk models, the credit-related optionality associated with a line of credit usually is represented by treating the draw-down rate as a known function of the customer's *end-of-period* credit rating. To illustrate, consider a one-year line of credit that is, initially, completely undrawn. Conditional on the customer's credit grade at the *end* of the planning horizon, the assumed end-of-period draw-down rate would be based on the average historical draw-down experience of customers having that future grade. Under the MTM framework, the future value of the line (conditional on the customer's *end-of-period* credit grade) would then be calculated as if the line were a loan equal to the assumed (conditional) draw-down.²³

Within the DM paradigm, since only two future credit "ratings" are relevant -- default and non-default -- a somewhat simpler approach is often employed. In effect, the undrawn credit facility is converted into a "Loan Equivalent Exposure" (LEQ) to make it comparable to a term loan. Ideally, the LEQ would be calculated as the expected draw-down under the

²² Another example of credit-related optionality would be a tendency for changes in a customer's financial condition to be reflected in changes to a facility's structure or terms. Under "grid pricing," for example, credit spreads are reset periodically based on changes in the underlying customer's credit rating or other indicators of financial condition. Similarly, prepayment options embedded in loans may generate credit-related optionality, since customers experiencing ratings upgrades may tend to exercise the prepayment option in order to refinance at lower credit risk spreads, whereas customers experiencing downgrades will not. One approach to modeling credit-related optionality within an MTM framework (KPMG Peat Marwick's Loan Analysis System™) is discussed in Scott D. Aguais, Larry Forest, Jr., Suresh Krishnamoorthy, and Tim Mueller, "Creating Value from Both Loan Structure and Price," *Commercial Lending Review* (forthcoming). Note that credit-related optionality also arises with products outside a bank's traditional lending operations. In a derivative transaction, for instance, a bank's counterparty credit risk generally will vary randomly over the life of the contract, reflecting changes in the amount by which the bank is "in the money."

²³ Within CreditMetrics, the line's future value is computed as the sum of the value of the drawn loan *plus* (for a multi-year credit line) the present value of expected future fees paid on the remaining undrawn amount. In effect, this specification assumes no draw-downs by the customer beyond the first year.

line in the event the customer were to become insolvent by the end of the period.²⁴ (Note that if the customer remains solvent, the size of the draw-down is irrelevant in DM models, since credit losses would equal zero.)

D. Summary of Credit Events and Risk Factors

The preceding discussion has highlighted three types of “credit events” that can lead to a change in the value of a loan. Under the DM and MTM paradigms, the credit loss for an individual loan reflects the combined influence of random risk factors determining (1) the loan’s LGD; (2) any change in the loan’s credit rating grade over the planning horizon (only migrations to “default” are relevant for DM models); and, within MTM models, (3) changes in credit spreads prevailing in financial markets.²⁵ (The Appendix presents a more formal characterization of the risk factors that determine each of these three types of credit events.)

In general, each risk factor is represented as a separate random variable, whose joint probability distribution with all other risk factors must be specified by the model-builder. Thus, for a portfolio consisting of N term loans, under the DM paradigm, at a minimum the joint distribution of $2N$ random risk factors would need to be modeled -- corresponding to two types of credit events per loan, representing each loan’s LGD and credit rating migration (into default) over the planning horizon. Because the MTM paradigm also includes as credit events any changes in the term structures of credit spreads, the model-builder needs to specify for each credit grade one or more additional random risk factors that are assumed to determine the term structure of credit spreads for each grade. The next section of this report describes the empirical methods used in specifying the joint probability distribution of the risk factors arising in DM and MTM models.

²⁴ For a plain-vanilla term loan, the LEQ would equal the amount of the loan.

²⁵ Changes in the risk-free yield curve are not treated as random credit events, but are normally set equal to the market expectations implied by the current risk-free term structure.

V. Bottom-up Credit Risk Models: Specification and Estimation

This section highlights what is perhaps the most challenging aspect of the credit risk modeling process, namely, the task of specifying the joint probability distribution of the risk factors affecting LGDs, credit rating migrations, and, within MTM-type models, changes in credit spreads. As will be noted repeatedly below, in general this process is difficult and imprecise. 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. Thus, the model specification process tends to involve many crucial simplifying assumptions as well as considerable judgment.

A. Correlations Among Different Types of Risk Factors

From standard portfolio theory, the overall uncertainty around a portfolio's rate of return depends on its systematic risk -- that is, co-movements in loan values arising from their dependence on common influences. Within the MTM framework described above, systematic risk may reflect four types of correlations among risk factors that potentially could contribute to co-movements in loan valuations: (1) correlations between risk factors affecting credit rating migrations, especially those corresponding to borrowers operating in related markets, such as the same geographic region or industrial sector; (2) correlations between risk factors determining LGDs; (3) correlations between risk factors driving changes in the term structures of credit risk spreads; and (4) cross-correlations among the risk factors affecting rating migrations, credit spreads, and LGDs.²⁶ Under the DM approach, of course, only three types of correlations are relevant: correlations between borrower defaults, correlations between LGDs, and cross-correlations among defaults and LGDs.

²⁶ Apart from these correlations, a change in any credit spread also would produce co-movements in loan values, since the same credit spreads are used in valuing *all* loans having a particular credit grade.

Although critically important, correlations among random variables are difficult to estimate reliably with relatively short historical sample periods. Model-builders therefore tend to impose fairly restrictive assumptions on the pattern of correlations among the risk factors. In particular, credit risk models nearly always assume *zero* correlations between risk factors of different types. That is, the risk factors affecting changes in credit ratings are assumed to be independent of those affecting changes in credit spreads, which, in turn, are assumed to be independent of those affecting LGDs. Given this assumption, model-builders typically focus on specifying the probability distribution for each type of risk factor separately from the others.

B. Parameter Estimation

Below, we summarize the modeling techniques used in determining probability distributions for each of the three types of risk factors. The discussion is intended to highlight the broad range of modeling assumptions, estimation methods, and practical judgments that permeate all aspects of the model specification process.

1. Risk factors affecting loss-rates-given-default. Outside the consumer and small business lending areas, an individual bank's historical data generally provide very limited information with which to estimate LGDs. Especially for middle-market and large corporate loans, the number of defaulted loans within a single bank's historical database is typically too small to permit the probability distribution of LGDs for any particular type of loan to be estimated accurately. This data problem is further compounded by the fact that LGDs normally will depend on the loan's seniority and collateral (if any), as well as on characteristics of the borrower, such as industrial sector and country. That is, a separate probability distribution of LGDs may be appropriate for each combination of seniority and collateral, and perhaps for each sector and country as well.

To make this estimation problem manageable, model builders generally invoke numerous simplifying assumptions. Within the current generation of credit risk models, for example, LGDs (after controlling for seniority, collateral, etc.) are usually assumed to be

independently and identically distributed over time and across all borrowers -- and in some models, even across obligations of the same borrower. Credit risk models also may assume that the probability distribution of LGDs takes a specific parametric form, such as that of a normal or beta distribution.²⁷ Given such assumptions, the underlying parameters characterizing the probability distribution of LGDs typically is inferred judgmentally by pooling data from several sources: (a) internal data on the bank's own historical loan losses, (b) consultants' proprietary data on the LGDs of their clients, and/or (c) historical LGD data culled from published articles on the subject.

2. Risk factors affecting changes in credit risk spreads. This area appears to be still in an early state of development, perhaps reflecting a lack of extensive databases on secondary market yields for lower-rated loans and bonds.²⁸ (Under the DM paradigm, of course, changes in credit risk spreads are irrelevant.) For banks that have assembled such historical data, non-parametric approaches are sometimes used to estimate the joint probability distribution of future changes in credit risk spreads. One such procedure involves constructing, for each credit rating grade, a database of historical term structures of credit risk spreads. The joint probability distribution of future credit spreads is then estimated using a within-sample Monte Carlo simulation procedure.²⁹

²⁷ Some models assume that all LGDs are known with certainty at the beginning of the planning horizon. As noted below, ignoring the uncertainty surrounding LGDs or other model parameters can contribute to significant under-estimates of a portfolio's credit risk.

²⁸ Within the current (April 1997) version of CreditMetrics, for instance, the risk factors affecting changes in credit risk spreads actually are set to *zero* for purposes of modeling future loan values.

²⁹ That is, for each Monte Carlo iteration of the credit risk model, a separate historical date is chosen randomly. All the changes in credit risk term structures used in computing future values for that iteration are then set equal to their levels as of that historical date. At the next iteration, another date is randomly selected, and the process repeated.

3. Risk factors affecting changes in credit ratings. Within most credit risk models, each customer's credit rating at the end of the planning horizon is represented in terms of the future realization of a migration risk factor (i.e., an unobservable "latent variable"). The value of that customer-specific migration risk factor in relation to various thresholds is assumed to determine the change in that customer's credit rating over the planning horizon. For example, given a customer's current credit rating (say, BBB), an extremely large positive realization of the migration risk factor might correspond to an upgrade to AAA, a somewhat smaller (but still very large) positive value might correspond to an upgrade to AA, and so on. Similarly, an extremely large negative realization might generate a downgrade to default, etc. Primarily for analytical convenience, migration risk factors usually are assumed to be jointly normally distributed.

Means and variances of migration risk factors. The stochastic properties of the migration risk factors typically are expressed in terms of a "ratings transition matrix" similar to that depicted in Exhibit 3. Given the customer's current credit rating (delineated by each row), the probability of migrating to another grade (delineated by the columns) is shown within the intersecting cell. Thus, in the exhibit, the likelihood of a BBB-rated loan migrating to single-B within one year would be 0.32 percent. (Since under the DM model only rating migrations into the default state lead to changes in the values of loans, only the last column of this matrix would be relevant.)

To help compensate for the lack of historical data on loan performance, model-builders tend to estimate ratings transition matrices through the pooling of data from various sources. Often, it is assumed that credit rating transition probabilities for loans (and correlations among the underlying risk factors) are identical to those for similarly rated corporate bonds. With this assumption, model parameters may be calibrated using public databases on the credit rating migrations of corporate bonds spanning twenty years or longer. Where possible, the migration data for corporate bonds are sometimes adjusted judgmentally to incorporate either information from a bank's own internal loan performance databases or information based on consultants' knowledge of the historical loan migrations at peer institutions.

Given the estimated ratings transition matrix, the means and variances of the underlying migration risk factors, together with the thresholds defining upgrades and downgrades, may be “reverse engineered” so as to be consistent with the individual elements of this matrix.³⁰ Under the common modeling assumption that migration risk factors are normally distributed, this process is greatly simplified by virtue of the fact that, without loss of generality, the means and variances of the migration risk factors can be set to zero and unity, respectively. Thus, only the thresholds for rating migrations need to be estimated explicitly.³¹

In theory, one would expect the rating transition probabilities applicable to a given customer at a point in time to be “conditional” on various firm-specific and macroeconomic variables, such as the cyclical volatility of the firm’s earnings (perhaps proxied by SIC code) and indicators of the current stage of the business cycle. In practice, however, there is generally insufficient data with which to estimate transition probabilities at such detail with reasonable precision. Thus, at most banks, the same rating transition matrix usually is applied to all borrowers, with no adjustment for business cycle effects.³² One potential implication of using “unconditional” transition probabilities is that estimates of expected losses and credit risk could be biased downward during the early stages of recessions, and biased upwards during the early stages of recoveries.

³⁰ See Greg M. Gupton, Christopher C. Finger, and Mickey Bhatia, *CreditMetrics -- Technical Document*, New York: Morgan Guaranty Trust Co., April 1997, p. 92.

³¹ In section IV.C, we summarize a mean/variance methodology that is often used to approximate the PDF for credit losses within the DM framework. Under this approximation approach, it is not necessary to explicitly estimate the means, variances, and correlations of the underlying migration risk factors, as defined above. Rather, the model-builder must estimate each loan’s probability of default and the correlations among loan defaults. The default probabilities would be culled directly from the ratings transition matrix, while estimates of default correlations would be based on methods similar to those described below.

³² For a description of a credit risk modeling approach that incorporates macroeconomic and business cycle effects, see Tom Wilson, “Portfolio Credit Risk (I),” *Risk Magazine*, Vol. 10/No. 9, September 1977, pp. 111-117, and “Portfolio Credit Risk (II),” *Risk Magazine*, Vol. 10/No. 10, October 1997, pp. 56-61.

Correlations between migration risk factors. With regard to the correlations between the migration risk factors affecting different customers, estimation procedures across the sampled banks are quite diverse. Within each bank, often a variety of methodologies are used to generate alternative estimates for individual correlation parameters, with each methodology involving its own restrictive and simplifying assumptions. Since these alternative estimates generally are still subject to considerable uncertainty and variation, model-builders often pool these separate estimates judgmentally in order to arrive at the specific correlation value that is used in the credit risk model.

One common estimation approach involves stratifying the credit portfolio into a relatively small number of mutually exclusive sub-portfolios or “buckets” for which annual historical default rates are available going back at least several years. Within each bucket, the loans are assumed to be statistically identical; that is, the correlation between any two risk factors is assumed to depend only on their respective buckets. Estimates of risk factor correlations for loans within the same, and across different, buckets are then inferred from the means, variances, and covariances of the historical default rates for the corresponding buckets.³³

For the most part, the stratification or bucketing schemes used in practice have tended to be based on internal credit ratings. However, the state-of-the-art in this area appears to be evolving rapidly toward more complex stratifications based on credit rating, industrial sector, *and* country. Some practitioners have begun estimating correlations between migration risk factors on a “name-by-name” basis, without any stratification whatsoever, using option-

³³ This procedure involves a two-stage process. In the first stage, the means, variances, and covariances cited in the text are used to estimate *default* correlations between loans of various types. (For loans within the same bucket, this technique is illustrated in “Appendix F: Inferring Default Correlations from Default Volatilities” in Greg M. Gupton, Christopher C. Finger, and Mickey Bhatia, *CreditMetrics -- Technical Document*, New York: Morgan Guaranty Trust Co., April 1997, pp. 159-160.) In the second stage, correlations between migration risk factors are inferred from the *default* correlations generated in the first step. The relationship between default correlations and migration risk factor correlations is developed in Chunsheng Zhou, “Default Correlation: An Analytical Result,” Board of Governors of the Federal Reserve System, May 1997.

theoretic techniques. These developments have been spurred by recent methodological advances that, under certain assumptions, permit correlations between migration risk factors to be inferred from co-movements in firms' equity prices.³⁴

C. PDF Computation Engine

Once the parameters of the credit risk model have been specified, the portfolio's PDF (as per Exhibit 1) generally is computed by one of two methods: (a) Monte Carlo simulation, or (b) approximations using a mean/variance methodology. Monte Carlo simulation techniques result in an estimated PDF whose "shape" is consistent with the parameters of the underlying credit risk model. The Monte Carlo techniques employed in credit risk modeling are essentially identical to those used within VaR models in the trading account, and are not discussed here.

Relatively few banks, however, currently use Monte Carlo methods to estimate PDFs. The vast majority use mean/variance approximations, which are viewed as computationally less burdensome.³⁵ With mean/variance approximations, the general shape of the PDF is assumed, rather than inferred from the underlying credit risk model. For example, often the PDF is assumed to take the shape of a beta or, in some instances, normal distribution having a mean and variance identical to the estimated mean and variance of the portfolio's credit losses -- even though neither distribution may be strictly consistent with the model's other assumptions and parameters. Under mean/variance approaches, the economic capital allocation process generally simplifies to setting capital at some multiple of the estimated standard deviation of the portfolio's credit losses.

1. Mean/variance approximation. The mean/variance approach is used primarily within the context of the DM paradigm for defining credit losses. A portfolio's expected

³⁴ These advances have been pioneered by KMV Corp.

³⁵ With recent advances in computing technologies, several of these institutions are considering adopting a Monte Carlo approach at some point in the future.

credit losses (μ) over the planning horizon equals the summation of the expected losses for the individual credit facilities making up the portfolio:

$$(1) \quad \mu = \sum_{i=1}^N P_i LEQ_i \overline{LGD}_i$$

where for the i^{th} facility, \overline{LGD}_i is the expected loss-rate-given-default, P_i is the probability of default, and LEQ_i is the loan equivalent exposure.

The portfolio's standard deviation of credit losses (σ) can be decomposed into the contribution from each of the individual credit facilities:

$$(2) \quad \sigma = \sum_{i=1}^N \sigma_i \rho_i,$$

where σ_i denotes the (stand-alone) standard deviation of credit losses for the i^{th} facility, and ρ_i denotes the correlation between credit losses on the i^{th} facility (per dollar of book value) and those on the overall portfolio. Typically, the economic capital allocation (for credit risk) against the i^{th} facility is set at some multiple of that facility's marginal contribution to the portfolio's overall standard deviation of credit losses (i.e., some multiple of $\sigma_i \rho_i$).

Under the further assumptions that (a) the random risk factors affecting customer defaults and LGDs are independent of one another, and (b) LGDs are independent across borrowers, the stand-alone standard deviation of credit losses for the i^{th} facility can be expressed as

$$(3) \quad \sigma_i = LEQ_i \sqrt{P_i (1 - P_i) \overline{LGD}_i^2 + P_i VOL_i^2},$$

where VOL is the standard deviation of the facility's LGD .

These equations provide a convenient way of summarizing the overall portfolio's credit risk (within the DM framework) in terms of each instrument's P , ρ , \overline{LGD} , VOL , and LEQ . They also serve to highlight those aspects of the estimation process that determine the overall reliability of a credit risk model; namely, (a) the accuracy of the above parameter

estimates as representations of the future, and (b) the validity of the underlying independence assumptions.

D. Capital Allocation Rule

Once the PDF for credit risk has been estimated, the bank must invoke a particular rule for determining how much economic capital it should hold against this risk. As indicated above, at most institutions this “capital allocation rule” is expressed as the capital necessary to achieve some target insolvency rate over the planning horizon. Among the banks sampled by the Task Force, the most widely used target insolvency rate was around 0.03 percent, the historical default rate on AA-rated corporate bonds. However, one bank in the sample used a 1.0 percent target default rate, while another used a 5.0 percent target rate. Note that the higher the target insolvency rate, the lower the allocated capital, other things the same.

In cases where the portfolio’s PDF is estimated directly via Monte Carlo simulation, the economic capital allocation against credit risk is computed directly from the estimated PDF, as shown in Exhibit 1. For banks using mean/variance approximation methods, economic capital is generally calculated as some multiple of the portfolio’s estimated standard deviation of credit losses, where this multiple is chosen to be consistent with the target insolvency rate and the assumed shape of the PDF (e.g., beta or normal). In practice, these multiples can vary widely (for example, between 3 and 7) depending on the target insolvency rate and on whether the “true” PDF is assumed to be beta- or normal-shaped. Final economic capital allocations, therefore, can differ considerably across banks owing to differences in their respective capital allocation rules.

VI. An Internal Models Approach to Setting Formal RBC Requirements

Internal risk models, and in particular credit risk models, present significant challenges as well as important opportunities for supervisors. Ongoing financial innovations, such as securitization, are rendering the formal RBC standards increasingly less relevant to the largest banks by providing ever more cost-effective means for undertaking regulatory

capital arbitrage. As credit exposures have become more complex and opaque, banks have invested heavily to improve the quality of their internal risk measurement systems. It would seem that supervisors, too, must adapt their methods for assessing credit risk and capital adequacy to the changing realities of the marketplace, or face a steady erosion in the effectiveness of capital-based regulatory/supervisory policies. The relevant question is how best to accomplish this task.

The Task Force has arrived at several tentative conclusions regarding potential near-term regulatory and/or supervisory applications of credit risk models. In this section, we examine the feasibility of replacing the Basle Accord with a full internal models approach to setting regulatory capital requirements against credit risk in the *banking book*. The final section of the report discusses near-term options for using banks' credit risk models on a more limited basis for prudential purposes.

The Task Force considered possible qualitative standards by which to assess the feasibility of an internal models approach to setting RBC requirements for credit risk in the *banking book*. The one existing framework for evaluating the current state-of-the-art in credit risk models is the set of qualitative standards for model reliability and integrity stipulated in the Market Risk Amendment, which sets forth the internal models approach to regulatory capital for market and specific risks in the *trading account*. These standards are interpreted as requiring that a risk measurement model be (a) analytically sound, (b) subject to periodic back-testing and stress testing, and (c) well integrated into the bank's management decision making process.³⁶ As discussed below, the current generation of internal credit risk models raises important concerns along each of these dimensions.

³⁶ In addition to the model-specific requirements noted in the text, a bank also must have a risk control unit that reports directly to senior management and is independent of business units. The bank also must conduct independent reviews of its risk measurement and risk management policies and procedures. (See Darryll Hendricks and Beverly Hirtle, "Bank Capital Requirements for Market Risk: the Internal Models Approach," Federal Reserve Bank of New York, September 1997.) We do *not* discuss these requirements below since, with regard to credit risk models, formal compliance with these provisions should not be a problem for any large bank satisfying the above model reliability and integrity standards.

A. Analytical Soundness

The details of credit risk model construction vary considerably across the institutions surveyed by the Task Force. Such diversity of practice is, on the one hand, desirable in that it is consistent with a continual process of improvement in best practices. On the other hand, in order to utilize information from internal credit risk models, supervisors must be in a position to evaluate the conditions under which any particular bank's model constitutes "acceptable" practice -- i.e., is analytically sound. Among the issues of potential interest to supervisors is whether the underlying conceptual framework for a bank's credit risk model (e.g., DM or MTM) encompasses essentially all of the institution's credit risks. Moreover, while the assumption of a one-year planning horizon is fairly standard, for some supervisory purposes this interval may be too short. There are also important challenges relating to the estimation of key parameters of the models -- challenges that appear to be greater than those faced in constructing market risk models. These analytical soundness issues are reviewed below.

1. Choice of planning horizon and loss paradigm. As noted above, banks typically employ a one-year planning horizon for purposes of credit risk modeling. This choice appears to be both pragmatic and, to some extent, arbitrary. On the one hand, in support of the one-year horizon, practitioners frequently suggest this interval represents a reasonable period over which -- *in the normal course of business* -- a bank could mitigate its credit exposures (at least with respect to large corporate customers), taking into account improving liquidity in secondary loan markets and the average effective maturity of most credit instruments. The vast majority of commercial lines of credit, for example, tend to have maturities of one year or less. Another, and perhaps the most important consideration, is that the estimation of many key model parameters is often viewed as infeasible for planning horizons much beyond one year due to the lack of historical data.

From a regulatory perspective, however, it is perhaps more appropriate to view the choice of planning horizon outside the *normal course of business*. Specifically, regulators tend to view capital adequacy within the context of a bank under *stress* attempting to unload the credit risk of a *significant* portfolio of deteriorating credits. Whereas the markets for

secondary loan trading and credit derivatives appear to be expanding and are becoming more liquid, they have not yet been fully tested by any large bank under severe stress. Moreover, fluctuations in economic activity and in credit losses tend to be positively serially correlated from one year to the next, implying that a bank's capital buffer may be called upon to absorb significant credit losses extending beyond a single year. Indeed, the experience of the U.S. banking agencies suggests that two or more years are typically required to resolve asset-quality problems at troubled banks and thrifts, over which time the deposit insurance funds remain exposed to potential further credit losses.

More generally, many credit instruments are subject to "adverse selection" that tends to increase the effective maturity of an instrument as its credit quality deteriorates. With regard to a nominal one-year loan commitment, for example, recall that banks typically experience greater rates of draw-down as a customer's credit rating declines. As a customer approaches insolvency, draw-downs under committed lines often approach 100 percent. When drawn, of course, the nominal one-year line of credit becomes effectively a loan, whose scheduled maturity may extend another year or longer. Furthermore, as a practical matter, the bank may continually extend or roll over the loan if it believes such actions will maximize its chances of recovery.

At least in theory, a case could be made that different planning horizons should be used for different types of credit instruments, depending on their effective maturity, the liquidity of related secondary markets, and the manner in which the underlying credit risks are managed. A "life-of-the-loan" planning horizon might be most appropriate for loans managed on a held-to-maturity basis, while a shorter horizon might be appropriate for relatively liquid portfolios that are actively managed on a mark-to-market basis. Some practitioners suggest that, in theory, it may be possible to blend multiple planning horizons within a single conceptual framework for allocating capital against credit risk.³⁷

The Task Force, however, is aware of no bank that has yet implemented a multiple-

³⁷ See for example, *CreditRisk⁺: A Credit Risk Management Framework*, Credit Suisse Financial Products, 1997; and Tom Wilson, "Portfolio Credit Risk (I)," *Risk Magazine*, Vol. 10/No. 9, September 1997, pp. 111-117.

horizon approach to credit risk measurement and internal capital allocations. Until the practical feasibility of a multi-horizon approach to credit risk modeling has established, most model-builders will likely continue to employ a single planning horizon assumption. In these circumstances, for the reasons discussed above, regulators may be skeptical of credit risk models that assume capital is needed to cover only one year of unexpected credit losses.

Since DM-type models ignore credit deteriorations short of default, economic capital allocations produced by these models may be particularly sensitive to the choice of a one-year planning horizon. As noted earlier, with respect to a three-year term loan, this assumption could mean that more than two-thirds of the loan's credit risk is potentially ignored. In contrast, under the more general MTM framework credit risk measurements encompass all potential reductions over the planning horizon in the portfolio's economic value due to credit quality deteriorations, whether to default or otherwise. Thus, the required economic capital allocations under this paradigm may be less sensitive than DM models to the assumption of a one-year horizon. Although very few large U.S. banks currently use the MTM framework within their banking books, some institutions are considering switching to such an approach down the road.

2. Estimates of loss-rates-given-default. The sophistication of methods for estimating LGDs varies considerably, especially for complex financial instruments supporting securitization activities. For example, it is not uncommon for banks to assume that the LGD for a subordinated loan functioning as a credit enhancement for publicly-issued asset-backed securities would be comparable to the LGD of a corporate loan secured by similar assets (e.g., trade receivables or consumer credit). In the event of default, however, a \$25 million subordinated loan supporting a \$1 billion pool of securitized assets will tend to exhibit a much greater expected loss rate and loss rate volatility -- corresponding to \overline{LGD} and VOL in equations (1) - (3) -- than would a typical \$25 million senior corporate loan secured by similar assets. This is because the former will generally absorb a disproportionate share, in some cases (by design) essentially all, of the credit losses on the underlying asset pool. Given the growing importance of securitization, for some banks the risk exposures arising

from credit enhancements may loom large in determining their overall capital adequacy.

3. Treatment of credit-related optionality. Methods for dealing with credit-related optionality are still evolving. The great diversity in practice frequently leads to very large differences across banks in credit risk estimates for similar instruments. With regard to virtually identical lines of credit, for instance, estimates of stand-alone credit risk can differ as much as a ten-fold. Moreover, these differences sometimes reflect modeling assumptions that are fundamentally inconsistent with a bank's own views regarding the nature of the underlying business. For example, with respect to committed lines of credit, some banks implicitly assume that future draw-down rates are *independent* of future changes in the customer's credit quality, despite evidence to the contrary, possibly leading to underestimates of the LEQs for lines of credit. No matter the process used for determining key parameters of the credit risk model, increasing (or decreasing) the LEQ for a particular instrument by a factor of X will increase (decrease) the resulting risk estimates and capital allocations for that instrument by the same factor.³⁸ Before employing internal credit risk models for setting formal regulatory capital standards, banks would need to bring greater consistency and rigor to their methods for dealing with credit-related optionality.

Issues also arise in the treatment of credit-related optionality in derivatives contracts. It is common practice to calculate the contract's LEQ as some variant of the amount by which bank is expected to be "in the money" on that contract over the planning horizon (or life-of-the-contract), based on simulations using the bank's trading account VaR models. This approach raises several potential risk measurement problems. Foremost, by setting the individual LEQs equal to the bank's *expected* exposures, the approach ignores uncertainty associated with these credit exposures, which could lead to under-estimates of overall credit risk. For example, counterparty credit risk exposures may be positively correlated across contracts (e.g., a bank having a large positive exposure with respect to future in oil prices

³⁸ This can be seen directly from inspection of equations (2) and (3), above, where the contribution of a particular instrument to the portfolio's overall standard deviation of credit losses is directly proportional to its LEQ.

could expect a significant change in oil prices to move these contracts into or out of the money together.) Similarly, in certain contracts, the extent to which a bank is “in the money” may be negatively correlated with changes in the credit quality of its counterparty.³⁹ Satisfactory treatment of such issues would seem to require a much closer integration of internal market risk and credit risk measurement systems than is typical at present.⁴⁰

4. Parameter calibration. Under both the MTM and DM frameworks, estimates of portfolio credit risk are driven largely by assumptions and parameter estimates regarding the joint probability distribution of the relevant risk factors. Because available data on the historical performance of different types of loans generally do not span sufficiently long time periods to enable precise estimation of this distribution, parameter values generally are established through a judgmental process involving considerable uncertainty.

To make the estimation process manageable, model-builders tend to invoke many critical simplifying assumptions. These often include the following:

1. Joint normality or other parametric assumptions on the probability distributions of

³⁹ Such transactions are termed “wrong-way” derivative contracts. To illustrate, consider an interest rate swap (with a cyclically-sensitive counterparty) where the bank pays a floating rate and receives a fixed rate. A large negative macro-economic shock might tend to generate a mark-to-market gain on the derivative position (as short-term interest rates fell in the economy), while at the same time tending to lower the counterparty’s credit quality. Another example would be a credit risk or foreign exchange derivative with a counterparty having substantial unhedged risk exposures that are positively correlated with its risk in this contract (e.g., a credit derivative in which the counterparty assumes the credit risk of a loan to an energy company, while already having an inordinate concentration of risk within the energy sector.)

⁴⁰ LEQ-type approaches to dealing with counterparty credit risk exposures may over-estimate a bank’s credit risk when changes in the position’s mark-to-market value and changes in the credit quality of the counterparty are negatively correlated. Over-estimates of credit risk also may arise if the bank simply sums the separate LEQs associated with multiple derivatives contracts when calculating the overall credit risk exposure with respect to each individual customer. In general, a bank should *not* sum the separate LEQs for a given customer, since at any point in the future the bank is likely to be “in the money” on some contracts and “out of the money” on others. Precisely how the separate LEQs should be aggregated would depend, in part, on the existence and enforceability of any netting agreements with the customer.

the risk factors determining credit rating migrations, or, in the case of mean/variance DM models, the assumption that portfolio credit losses have a beta (or normal) probability distribution;

2. Independence between risk factors affecting changes in credit ratings, changes in credit spreads, and LGDs;
3. Independence of LGDs across borrowers; and
4. Stability of model parameters.

In reviewing these assumptions, it should be noted that estimation of the extreme tail of a credit portfolio's PDF (the focus of credit risk models) is likely to be highly sensitive to variations in key parameters, such as correlations, or the assumption of joint normality. Surprisingly, in practice there is generally little analysis supporting the underlying modeling assumptions; indeed, model-builders generally recognize that theoretical and empirical objections can be raised concerning their plausibility in many instances. Nor is it standard practice to conduct sensitivity testing of a model's vulnerability to key parameters or assumptions. Moreover, when estimating credit risk, practitioners generally presume that all parameters and assumptions are known with certainty, thus ignoring credit risk issues arising from model uncertainty and/or instability.⁴¹

B. Model Validation: Back-testing and Stress Testing

Given the extensive judgment required in specifying credit risk models, the need for effective *model validation* procedures is clearly important. But, the same data limitations that render parameter calibration problematic also make model validation exceedingly difficult. In many ways, the task of estimating the extreme tail of the PDF is comparable to predicting the frequency at which credit losses in any year will exceed many multiples of a normal year's losses. The only entirely objective method for evaluating the statistical accuracy of a

⁴¹ Accounting for uncertainty in parameter estimates can significantly increase measured credit risk. See Gregory R. Duffee, "On Measuring Credit Risks of Derivative Instruments," *Journal of Banking and Finance*, 20, pp. 805-833 (1996).

credit risk model is to compare (over periods spanning multiple credit cycles) the model's *ex ante* estimates of PDFs against *ex post* realizations of actual credit losses. That is, only the realization of more frequent, extreme credit losses (relative to the model's predictions) can provide a purely statistical basis for concluding a model is deficient.⁴² Back-testing, therefore, is almost certain to be problematic in practice, owing to insufficient data for out-of-sample testing.⁴³ For this reason, banks generally do not conduct statistical back-testing on the PDFs predicted by their credit risk models.⁴⁴

In lieu of formal back-testing, credit risk models tend to be validated indirectly, through various market-based "reality" checks. Peer group analysis, discussed in section III, is used extensively to gauge the reasonableness of credit risk models and internal capital allocation processes. Another market-based validation technique involves comparing the bank's hurdle rate with the expected risk-adjusted rate of return (i.e., the RAROC) that could be achieved by investing in corporate bonds or syndicated loans having a particular credit rating, say, BB. An implied RAROC well below (above) the bank's hurdle rate might be interpreted as evidence that the model's capital allocation for BB-rated credits was too high (low), possibly requiring some re-calibration of the model's parameters.

Clearly, an implicit assumption underlying these techniques is that market perceptions of appropriate capital levels or appropriate credit risk spreads are "about right." If a bank elects to use such information to re-calibrate its risk model, re-calibration dates would need to be selected carefully so as to be reasonably confident that prevailing market perceptions

⁴² Discussions of various statistical issues associated with validating such models are presented in Francis X. Diebold, Todd A. Gunther, and Anthony S. Tay, "Evaluating Density Forecasts," National Bureau of Economic Research, Technical Working Paper No. 215, October 1997.

⁴³ A similar concern has been raised in the context of the validation of trading account VaR models. See Paul Kupiec, "Techniques for Verifying the Accuracy of Risk Measurement Models," Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series NO. 95-24, May 1995.

⁴⁴ The back-testing that is performed generally is limited to comparisons of *expected* and actual credit losses or default rates. Such tests do not address the accuracy of the model's predictions of *credit risk*, against which economic capital is allocated.

were economically well founded. Otherwise, the bank could cede to the vagaries of the market much of the internal pricing and risk management discipline it had hoped to achieve through implementation of an economic capital allocation system. From a supervisory perspective, the use of market-based validation methods could raise serious concerns regarding the comparability and consistency of a risk model over time.

All the issues discussed above are *qualitatively* similar to concerns that were raised in the course of developing the internal models approach to RBC for market and specific risks within the *trading account*, as set forth in the Market Risk Amendment. Within the context of internal credit risk models for the *banking book*, however, the problems are likely much more acute. At most large banks, the size of the banking book and the length of its relevant planning horizon are many times larger than those of the trading account, implying that errors in measuring risks for the banking book are more likely to affect assessments of the institution's overall financial health. Moreover, the banking book does not benefit from relatively high liquidity and a daily mark-to-market process which, in the context of the trading account, provide substantial safeguards against significant losses accumulating unnoticed and unaddressed.

In principle, *stress testing* could at least partially compensate for the data limitations, estimation problems, and shortcomings in available back-testing methods for credit risk models. Most of the uncertainty within credit risk models (and the infeasibility of back-testing) relates to estimation of the joint probability distribution of risk factors. Stress tests circumvent these difficulties by specifying, albeit arbitrarily, particular economic scenarios against which the bank's capital adequacy might be judged -- without regard to the probability of that event actually occurring. Stress testing is used routinely by the credit rating agencies, who often assign credit ratings on the basis of a security's ability to withstand various stress scenarios: to qualify for a AAA rating, the security would have to avoid defaulting under a AAA-scenario, to qualify for a AA rating, the security would have to withstand a AA-scenario, and so forth. Similarly, with respect to banks' trading activities, stress tests designed to simulate hypothetical market disturbances (e.g., the October 1987 stock market crash) provide useful checks on the reasonableness of the required capital levels

generated by banks' VaR models. Although stress testing protocols might be developed for internal credit risk models used within the booking book, the Task Force is unaware of banks actively pursuing this approach.

C. Integration of Credit Risk Models into Decision Making Processes

The extent to which the output of a risk model is incorporated into a bank's actual decision making processes is highly suggestive of management's confidence in that model. In practice, the extent of reliance on credit risk models differs greatly among banks. Much of this variation may reflect differences across institutions in the length of time over which their risk measurement/capital allocation systems have been operational. Generally speaking, the longer the period over which such systems have been in place, the greater their penetration into the bank's decision making processes.

Within those institutions having the most sophisticated systems for allocating economic capital for credit risk, the outputs of these systems frequently are embedded throughout the bank's risk management and incentive systems. At such banks, economic capital allocations are critical components of the processes for determining breakeven prices on credit instruments, for setting customer credit limits and broad portfolio concentration limits, and, in some instances, for actively managing overall portfolio credit risk on a day-to-day basis. Moreover, at some institutions, risk-adjusted measures of profitability are significant factors in assessing customer profitability and managerial performance and compensation. Thus, among some banks, the commitment to the quantification of credit and other risks appears genuine.⁴⁵ Nevertheless, most internal capital allocation systems have been implemented only within the last five years, and thus have not yet been tested under the stress of a full business cycle.

⁴⁵ Even within these institutions there are notable instances where credit risk models are not considered in situations where they might have been. For example as described in section II, economic capital allocations assume implicitly that the institution's loan loss reserve is consistent with the level of expected losses implied by the underlying credit risk model. In practice, however, actual loan loss provisions tend to be developed outside the institution's credit risk modeling and capital allocation systems.

While many banks use internal capital allocations for credit risk within a variety of decision making environments, one aspect of these processes suggests an additional caveat when contemplating the use of internal credit risk models for setting regulatory capital requirements. Credit and market risks are not the only types of risk against which banks allocated economic capital. Frequently, operating risks account for a substantial fraction (20 percent or more) of large banks' total risk and allocated capital.⁴⁶ Thus, assessing capital for credit and market risks, but not against operating risks, could substantially understate banks' overall capital needs. This problem is complicated by the fact that, while operating risks are viewed as quite important, models for quantifying these risks generally are primitive compared with those for market and credit risks. Before adopting an internal models approach to setting RBC requirements for *credit risk*, therefore, regulators would need to consider carefully whether and how *operating risks* should be incorporated into the RBC framework.

VII. 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. To be sure, some issues (such as credit-related optionality and the treatment of operating risk) are conceptually difficult. Still others (such as parameter uncertainty arising from data limitations) may be inherently problematic, at least until secondary market liquidity improves to the point where credit risk models for the banking book, like VaR models for the trading account, can be constructed directly from market price data. Nevertheless, it is important to recognize that, to be useful for prudential purposes, internal credit risk models need not be specified with absolute precision, so long as tools are available to reasonably account for the underlying modeling uncertainties. Although the development of such tools has not been a central focus of credit risk model-builders to date, it would seem that the bulk of the reliability concerns cited above could be addressed by using existing technologies for

⁴⁶ Operating risk generally is defined rather broadly to encompass all risks that are not clearly credit risks or market risks.

sensitivity testing, stress testing, lengthening the planning horizon, and recognizing parameter uncertainty explicitly within the modeling process.

Credit risk modeling practices are progressing so rapidly it is conceivable, perhaps even likely, they will become the foundation for a new approach to setting formal regulatory capital requirements. Whatever that time frame, however, if prudential capital policies are to remain an effective policy instrument even over the relatively short run, supervisors need to improve their existing methods for assessing the overall soundness and capital adequacy of large, complex banking organizations.

The need for more effective supervisory tools with which to assess banking risks and capital adequacy is heightened by ongoing financial innovation and new methods of regulatory capital arbitrage. While securitization has been the classic mode for conducting such arbitrage, new securitization techniques (such as collateralized loan obligations and securitized credit linked notes) and the Market Risk Amendment raise the possibility that such arbitrage may expand dramatically in the coming years. The key issue is not only that too little capital may be required against a bank's retained securitization risks (or against its trading account positions). Rather, a major concern is that, by removing the highest quality assets from the banking book ("cherry-picking"), such arbitrage may reduce the banking book's residual credit quality to the point where the 8 percent Basle standard is no longer sufficient. Financial innovation clearly increases pressure on supervisors to continually revise the capital treatment of banking book assets, or face ongoing erosion of the effective capital standards for large banks.

In addition, there is the related concern that the current regulatory capital framework may actually inhibit banks' attempts to engage in transactions that, on a portfolio-wide basis, *reduce* credit risk. For example, in some cases the nearly complete hedging of credit risk positions, through the use of credit derivatives, may result in little reduction in a bank's overall capital, relative to a strategy of no hedging at all.⁴⁷ Practitioners also claim that the "unreasonably" high regulatory capital requirements on high quality non-sovereign or non-

⁴⁷ See "Credit Risk and Regulatory Capital," International Swaps and Derivatives Association, March, 1998.

bank credits (AAA/AA) reduce rates of return to levels that make these assets uneconomic for banks to hold, thus depriving them of simple methods for reducing portfolio credit risk.⁴⁸

Given these concerns, it would appear that supervisors have little choice but to work aggressively to develop better analytical methods for quantifying credit and other risks,. Regardless of how formal RBC *regulation* evolves, supervisors need appropriate tools for making independent assessments of a bank's capital adequacy, based on the best available information -- which generally will be derived from a bank's own internal risk measurement and management systems. This view is consistent with the emphasis of the U.S. banking agencies on risk-focused supervision, which is inherently centered on banks' internal risk management systems and on fostering improvements to these systems over time.

The Task Force is considering several near-term possibilities for utilizing internal credit risk models within prudential capital policies. These potential applications may be divided into two main areas: (a) the setting of RBC requirements for *selected* credit instruments, and (b) the development of enhanced examination guidance on assessing the capital adequacy of large, complex banks.

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 enhancement supporting banks' securitization activities. This class of credit instruments includes, for example, the first-loss positions and other subordinated

⁴⁸ Op. cit., p. 8.

interests of most securitized asset pools. Many current securitization structures were not contemplated when the Accord was drafted. In this area, the 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. In lieu of requiring that banks obtain *external* ratings from outside rating agencies, a natural extension of the above proposal would permit a bank to use its own *internal* credit ratings, provided they were judged to be “reliable” by supervisors.

A further extension might involve the direct use of internal credit risk models in setting formal RBC requirements for securitization-related credit enhancements. Internal credit risk models deemed “reliable” by supervisors could provide the first practical means of assigning economically reasonable capital requirements for these instruments. Indeed, market acceptance of securitization programs is based heavily on the ability of issuers to quantify the credit risks of the underlying pools of securitized assets. The development of an internal models approach to RBC requirements -- on a limited scale for these selected instruments -- also would provide a useful test-bed for enhancing supervisors’ understanding and confidence in credit risk models. In addition, such an approach would encourage banks to continue to improve their internal systems for measuring and managing the risks of these rapidly expanding, yet highly complex credit activities.

B. Improved Examination Guidance

Apart from their possible use in setting formal RBC standards, the inputs and outputs of banks’ internal credit risk models could enhance assessments of bank capital adequacy

through the examination process. For instance, examiners could use a bank's own internal credit rating systems to assess the relative riskiness of a bank's credit portfolio. Provided that a concordance schedule could be developed that appropriately translated each bank's rating "buckets" into a common standard (perhaps paralleling S&P's or Moody's rating systems), examiners could assess how the credit quality of a particular bank's portfolio compared with that of its peers. This information could be used in much the same way that senior bank managers now use their own internal credit rating reports to evaluate the adequacy of the loan loss reserve and changes in a portfolio's credit quality over time.

More broadly, it may be possible for supervisors effectively to use internal credit risk models to assess the quality of a bank's risk management systems and its overall capital adequacy. To give one example, in contrast to the current RBC framework, typical internal capital allocations for unsecured term loans often range from 1 percent or less for AAA-rated loans to more than 30 percent against loans classified as "doubtful" -- *not* counting any reserves for expected future charge-offs. Examiners might 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 risk measurements or internal capital allocation ratios, by grade, of peer institutions). Over time, examination guidance might evolve to encompass additional elements of banks' internal risk models, including analytical tools based on stress test methodologies.

To have a good chance of success, an examination-based internal models approach to assessing capital adequacy on a bank-by-bank basis would require that banks themselves have in place internal review processes for evaluating their own overall capital adequacy. Unless such internal review processes are in place, it may be impractical for examiners to conduct independent assessments of capital adequacy and to engage senior management in constructive dialogues on this subject, absent clear indications of extant asset-quality problems. Arguably, internal reviews of capital adequacy should be a core element of any bank's overall risk management procedures and practices. However, often this is not the case, even at large banks that already have the basic components of such a process. To encourage greater progress in this direction, the banking agencies might consider issuing

“sound practices” guidance regarding the importance of such internal reviews, especially for the largest, most complex banks.⁴⁹

Regardless of the specific details, the development and field testing of examination guidance dealing with internal credit risk models would provide several useful benefits. Such an initiative would encourage further model development by banks, and would help ensure that supervisors remained abreast of ongoing improvements in risk modeling practices. In addition, both supervisors and the banking industry would benefit from the development of sound practice guidance on the design, implementation, and application of internal credit risk models and capital allocation systems. 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. Such efforts also would likely stimulate constructive discussions among supervisors and bankers on ways to improve credit risk measurement and management practices.

C. Concluding Remarks

The discussion above provides several possible examples by which information from internal credit risk models might be usefully incorporated into regulatory or supervisory capital policies. None of these examples constitute specific recommendations of the Task Force at this time. Rather, they are intended solely to provide a basis for continuing discussions regarding the possible roles for internal credit risk models within prudential capital policies. In view of the model reliability concerns described above, incorporating internal credit risk measurement and capital allocation systems into these policies will occur neither quickly nor without significant challenges. Nevertheless, we 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

⁴⁹ Recently, in the context of securitization and secondary loan market activities, the Federal Reserve issued similar guidance requiring that banks have in place appropriate internal risk measurement and capital allocation systems for managing the underlying credit and other risks (SR Letter 97-12, “Risk Management and Capital Adequacy of Exposures Arising from Secondary Market Credit Activities,” July 11, 1997).

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 that do not involve a models-based framework for capital adequacy assessments by the bank or the supervisor.

Appendix

The MTM Approach to Credit Risk Modeling

This appendix provides a more technical description of the stylized MTM credit risk model discussed in the text.

A. Valuation of Loans

Suppose the bank has N customers, where the current credit rating grade of the i^{th} customer is denoted g_i . The number of internal rating grades is denoted G , where grades one through $G-1$ are non-default states, and grade G represents a default. The term loan to the i^{th} customer has a contractual coupon payment of C_i dollars per period until maturity in period M_i , at which point the final payment (principal plus coupon) equals $C_i + P_i$.

1. Current values of loans. Given these assumptions, the current MTM value (at the beginning of the planning horizon) of a loan to the i^{th} customer equals the present discounted value of its contractual cash flows:

$$(A1) \quad V_i = \frac{C_i}{[1 + {}_1r + {}_1R(g_i)]} + \frac{C_i}{[1 + {}_1r + {}_1R(g_i)][1 + {}_2r + {}_2R(g_i)]} + \dots + \frac{C_i + P_i}{\prod_{k=1}^{M_i} [1 + {}_kr + {}_kR(g_i)]}$$

The discount rate for period k equals the sum of (a) the forward risk-free rate implied by Treasury term structure, denoted ${}_kr$; and (b) the market risk premium for deflating period- k contractual cash flows of g_i -rated obligors, denoted ${}_kR(g_i)$. In principle, the discount factors for the i^{th} customer could include a purely idiosyncratic component, affecting only that individual customer. However, to simplify the following exposition, this component is ignored. That is, credit risk spreads are assumed to depend only on the obligor/facility credit rating (g_i).

2. Future values of non-defaulted loans. From expression (A1), the value of a non-defaulting loan to the i^{th} customer as of the end of the planning horizon is given by

$$(A2) \quad \hat{V}_i = C_i + \frac{C_i}{[1 + 2r + 2\hat{R}(\hat{g}_i)]} + \frac{C_i}{[1 + 2r + 2\hat{R}(\hat{g}_i)][1 + 3r + 3\hat{R}(\hat{g}_i)]} + \dots + \frac{C_i + P_i}{\prod_{k=2}^{M_i-1} [1 + kr + k\hat{R}(\hat{g}_i)]}$$

where a hat (^) or tilde (~) over a variable indicates that its value is taken as of the end of the planning period. A tilde signifies that the variable is exogenous (i.e., does not depend on other variables), while a hat signifies that it is endogenous. Thus, ${}_k\hat{R}(\hat{g}_i)$ denotes the market risk premium for obligors rated \hat{g}_i , where both the risk premium and the credit rating are endogenous variables measured as of the end of the holding period.

3. Future values of defaulted loans. Banks generally do *not* rely on the valuation equation (A2) -- which discounts *contractual* cash flows -- for modeling the end-of-period values of *defaulted* loans. While the default probability on a commercial loan might be reasonably expected to behave like that on the same obligor's bond, commercial loans tend to exhibit a very different seniority and collateral status. The decline in the economic value of a defaulted loan (relative to its book value, B_i) is typically determined as the loan's book value *times* its random loss rate given default ($L\hat{G}D$):

$$(A3) \quad \hat{V}_i = B_i(1 - L\hat{G}D_i).$$

Within this simplified model, $L\hat{G}D$ s are assumed to equal the sum of a fixed average loss rate, L , and a zero-mean random error term, \tilde{l}_i :

$$(A4) \quad L\hat{G}D_i = L + \tilde{l}_i.$$

B. Credit Rating Migrations

The likelihood of a facility migrating to another credit risk grade over the planning horizon is represented through a “ratings transition matrix,” similar to that depicted in Exhibit 3. For a given customer, a rating migration from g_i to \hat{g}_i is assumed to depend on the future realization of a customer-specific latent random variable, \tilde{v}_i , representing the change in that borrower’s financial condition over the planning horizon.

Specifically, for an internal credit rating system with G grades

$$(A5) \quad \hat{g}_i = \begin{cases} 1 & \text{if } \tilde{v}_i \leq V_1(g_i) \\ 2 & \text{if } V_1(g_i) < \tilde{v}_i \leq V_2(g_i) \\ \vdots & \\ G-1 & \text{if } V_{G-1}(g_i) < \tilde{v}_i \leq V_G(g_i) \\ G & \text{Otherwise} \end{cases}$$

where for a customer having a credit rating of g_i , the $V_1(g_i), \dots, V_G(g_i)$ denote the threshold levels of \tilde{v}_i that trigger rating downgrades or upgrades. Thus, for a grade-4 facility (i.e., $g_i = 4$) a value of \tilde{v}_i less than or equal to $V_1(4)$ would imply a future credit rating of grade-1, a value greater than $V_1(4)$ but less than or equal to $V_2(4)$ would imply a grade-2, and so forth. Mathematically, the threshold levels are chosen so that the probability of any borrower migrating to another grade, given its current rating, agrees with the assumed rating transition matrix.

C. Changes in Credit Risk Spreads

It is assumed that for a given credit rating g , changes in the credit risk spread for period k are random:

$$(A6) \quad {}_k\hat{R}(g) = {}_kR(g) + \tilde{z}_k(g), \quad \text{for } k=1, 2, \dots, M$$

where M is the longest maturity of any loan and $\tilde{z}_k(g)$ denotes a random risk factor. (In


practice, the model for credit risk spreads may be expressed in terms of relative or logarithmic changes in yields, rather than absolute changes in yields.)

D. Risk Factors

The main body of this paper refers to three types of risk factors within the MTM model. In terms of the above model specification, these risk factors correspond to (1) the random variables affecting rating migrations (the \tilde{v}_i , for $i=1, 2, \dots, N$); (2) the random variable affecting credit risk spreads (the $\tilde{z}_k(g)$, for $g=1, 2, \dots, G$, and $k=1, 2, 3, \dots, M$); and (3) the random variables affecting loss-rates-given-default (the \tilde{l}_i , for $i= 1, 2, \dots, N$).

Exhibit 1

Relationship Between PDF and Allocated Economic Capital



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

Exhibit 2 Overview of Risk Measurement Systems

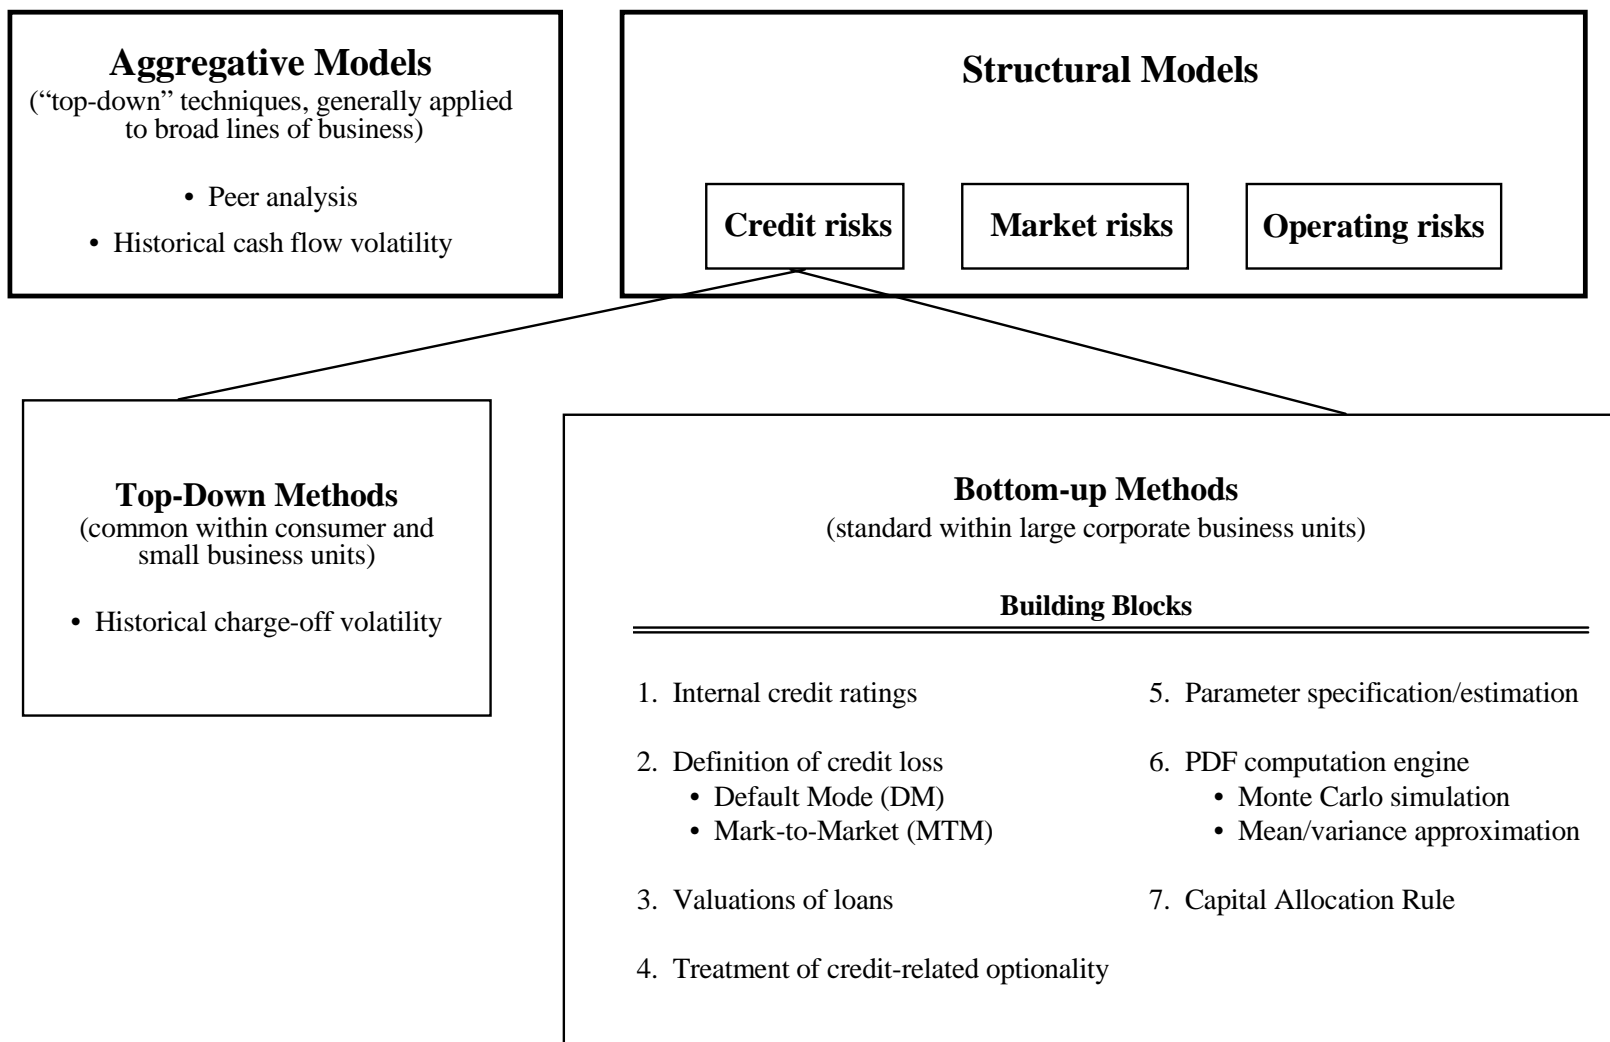


Exhibit 3
Sample Credit Rating Transition Matrix
(probability of migrating to another rating within one year, percent)

Credit Rating One-year in the Future

		AAA	AA	A	BBB	BB	B	CCC	Default
Current Credit Rating	AAA	87.74	10.93	0.45	0.63	0.12	0.10	0.02	0.02
	AA	0.84	88.23	7.47	2.16	1.11	0.13	0.05	0.02
	A	0.27	1.59	89.05	7.40	1.48	0.13	0.06	0.03
	BBB	1.84	1.89	5.00	84.21	6.51	0.32	0.16	0.07
	BB	0.08	2.91	3.29	5.53	74.68	8.05	4.14	1.32
	B	0.21	0.36	9.25	8.29	2.31	63.89	10.13	5.58
	CCC	0.06	0.25	1.85	2.06	12.34	24.86	39.97	18.60

Source: Greg M. Gupton, Christopher C. Finger, and Mickey Bhatia, *CreditMetrics -- Technical Document*, New York: Morgan Guaranty Trust Co., April 1997, p. 76.

Note: The credit rating transition matrix is based on the historical migration frequencies of publicly rated corporate bonds. The transition probabilities in the table have been statistically “smoothed” in order to attenuate the effects of sampling variation in the actual migration patterns of corporate bonds. Nevertheless, some anomalies remain. For example, the probability of a single-B loan being upgraded to BB (2.31 percent) is less than its probability of being upgraded to single-A (9.25 percent).