

Fondo Interbancario di Tutela dei Depositi



A Flexible Framework for
Credit-Risk-Sensitive Capital Allocation
and Sustainable Regulatory Capital Standards

Mark Carey

Federal Reserve Board

Collana Working Papers

Numero 1 / 1999

A Flexible Framework for Credit-Risk-Sensitive Capital Allocation and Sustainable Regulatory Capital Standards

Mark Carey[§]

--- Abstract ---

This paper suggests a flexible framework for quantitative capital allocation systems for internal use by banks and as the basis for more risk-sensitive regulatory capital standards for credit risk. Simple versions of the framework are feasible in the near term for a variety of banks, countries, and types of assets. Incremental evolution to more precise and complex systems is supported as well. However, a variety of research projects must be completed before any version can be implemented with confidence.

§ - Federal Reserve Board - The views expressed herein are not necessarily those of the Board of Governors, other members of its staff, or the Federal Reserve System. Thanks to Bill Treacy, John Mingo, David Jones, Ed Ettin, and Michael Gordy for useful conversations. Correspondence should be addressed to the author at Mail Stop 153, Federal Reserve Board, Washington, DC 20551. (202) 452-2784 (voice), (202) 452-5295 (fax), mcarey@frb.gov.

Formal analysis of credit risk positions is rapidly becoming an important element of commercial bank management throughout the world. Such analysis involves quantitative projections of portfolio losses caused by changes in the credit quality of counterparties. In principle, portfolio loss estimates support relatively precise analysis of a bank's loan loss reserve and capital requirements, of appropriate allocations of resources among business lines, and of product pricing. In the past, when portfolio compositions and the nature of competition were relatively stable over time, a bank could make such decisions based on broad historical experience and the judgment of its senior officers. Today, as banks move into new lines of business and new regions, and as competition intensifies, a reliance on traditional methods alone exposes a bank to insolvency risks and to lemons risks (the chance that its profitable customers and lines of business will be bid away by sophisticated competitors, leaving only those likely to go sour).

The advent of formal credit risk analysis is also important to regulators. Banks' loan loss reserves and capital are the buffers that preserve the solvency of individual banks and to a large extent the stability of banking systems. Implementation of the Basle Accord, an important regulatory achievement of the last decade, has fostered an increase in bank capital from the low levels of the 1980s. However, there is a danger of a return to inadequate capitalization. Because the Accord's risk weighting scheme applies the same marginal capital requirement to assets posing widely varying risks (and no capital requirement to other positions, especially certain off-balance sheet positions), banks are increasingly abandoning low-risk assets or moving them off the balance sheet. If this process continues, banking systems throughout the world may once again become undercapitalized and be very vulnerable to destabilization by credit risk events. One solution to the problem involves making the marginal regulatory capital requirement applied to each credit risk position correspond reasonably well to that position's contribution to the bank's portfolio credit risk. A debate aimed at producing such a more credit-sensitive revision of the Basle Accord is currently underway.

Formal analysis of credit risk is a very difficult task. Although the majority of the world's largest banks probably now have a unit devoted to credit risk modeling, very few banks have integrated such analysis into their day-to-day management decision making. Such implementation is difficult because

existing modeling systems frequently are applicable to only a portion of a bank's positions (for example, the large commercial loan portfolio), and because the proper structure and expected reliability of credit risk models are controversial subjects. Both senior managers at banks and regulators face a dilemma: although they sense the need for formal credit risk modeling as an element of bank management and bank regulation, they are reluctant to commit to incomplete, untested, difficult-to-validate technologies.

In this paper, I suggest a flexible framework for implementation of credit risk modeling in bank management and regulation. The approach can be applied to complete portfolios and supports capture of many of the benefits of existing knowledge about credit risk while allowing incremental improvements in sophistication as knowledge increases. Within the structure fall models and systems that could be implemented by banks and regulators during the next few years and that would be helpful to banks in both traditional and new lines of business and in many countries.

Two ideas form the basis for such flexibility. First, a sophisticated framework capable of incorporating many foreseeable technical advances can be collapsed to a relatively simple system by assuming that certain parameter values are constants. Second, if similar assets are grouped together in blocks, aspects of the risk modeling problem that are well-understood can be handled in a complex and sophisticated manner whereas aspects that are not well-understood can be handled through simplifying assumptions. In essence, the framework suggests that a range of operational modeling strategies can coexist, with complexity varying within and across banks and over time as appropriate.

For example, the simplest implementations of the framework need not involve operational modeling, but only the multiplication of various asset category totals by analytically developed factors, in some ways similar to the current Basle Accord. The cost of such simplifications is lesser precision relative to an ideal case, but the ideal case cannot be implemented with confidence given current knowledge, and important benefits can be captured by the simplest implementation. For example, available evidence implies that the capital needed for reasonably well-diversified portfolios is related to portfolio expected loss rates. A high-investment-grade loan portfolio demands far less capital than a below-investment-grade portfolio.

However, the simplest implementations are not adequate in all cases today and will become less

adequate as time passes because knowledge is evolving rapidly. For example, simple implementations of the framework cannot handle the impact of certain kinds of positions, such as credit derivatives, on a bank's credit risk posture and capital requirements. A virtue of the framework is that it permits simple methods to be used for traditional banking book assets while more extensive modeling addresses the more difficult positions. In general, it is important that advances be incorporated into operational systems as they are validated rather than requiring a perfection of knowledge before anything is implemented.

A primary focus of the current credit risk modeling debate is the relative merits of various existing modeling systems. Recent papers have shown that the models deliver similar capital allocations in some circumstances and different allocations in others (Gordy (1998), Koyluoglu and Hickman (1998)). In this paper, I take a practical perspective, including elements of several different models rather than trying to choose among them. Similar to CreditMetrics and KMV's PortfolioManager, this paper advocates a Monte Carlo aggregation strategy (although in the simple case no individual bank would need to operate Monte Carlo programs). Similar to CreditRisk+, assets are grouped according to certain characteristics, simplifying structure and computations. Similar to CreditPortfolioView, the impact of macroeconomic fluctuations on credit losses is an essential consideration, but in this paper I assume such fluctuations are accounted for in estimating parameters.

Banks could improve their decision making in the short run by implementing simple versions of the framework and gradually increase the sophistication of their analyses in the long run. Regulators could use the framework to implement a more risk-sensitive capital standard that would immediately remove many of the largest economic distortions inherent in the current Basle Accord, that would not be too burdensome on either banks or regulators, and that would permit advanced banks to employ individually sophisticated systems according to their needs and as the required technology becomes available and verifiable.

In this paper, considerable attention is paid to the credit risk modeling needs of regulators. Such attention is not solely because of the identity of my employer. As a practical matter, banks' own progress in modeling and managing credit risk will be most rapid if an appropriately enhanced

international regulatory capital standard is adopted quickly, for two reasons. First, a major barrier to more sophisticated capital allocation in many banks' day-to-day decision making is the resistance of line managers to measurement of the risk postures they oversee. Implementation of more risk-sensitive regulatory capital requirements will legitimize risk management units' efforts and provide an impetus to development of better systems.

Second, at this time, credit risk model development has a substantial public good component. The required databases and analysis are expensive, and even where an individual bank chooses to bear such costs the resulting databases may be too small to support confident estimates of model parameters. Regulators are the natural agent to organize industry-wide cooperative projects, but in most countries the budgetary and political realities that constrain regulators make it necessary that such projects accompany (rather than precede) development and implementation of revised regulatory capital requirements.

Another current focus of the debate is the relative merits of so called «full models» or «internal models» approaches to capital regulation versus an approach variously referred to as «ratings-based,» «rating-sensitive,» «modified Basle,» or «more buckets» (see Mingo (1998)). I believe the two approaches are fundamentally rather similar in their technical demands except for their handling of portfolio diversification. However, the approaches differ in that regulators would strongly constrain choices of model structure and parameters in the second case, whereas under the full models approach regulators would validate but in practice allow banks much more leeway (similar to their current internal models approach to market risk). In a technical sense, this paper's framework incorporates both rating-based and full models approaches as special cases. At least initially, I advocate strong regulatory discipline of modeling choices as a practical necessity and as an element of efficient production of the needed public goods.

In this paper, no attention is paid to market risk, interest rate risk, or operational risk nor to the manner in which management and regulation of such risks might be integrated with credit risk. Although such integration is intellectually and practically very desirable, in my view, we must first achieve satisfactory methods of dealing with the different risks on a stand-alone basis.

The framework is sketched in this paper rather than presented in a detailed, implementable form. The paper is likely to be unsatisfying in two respects. First, although I have a view of the best way forward, this draft is the first attempt to explain it, and thus the paper is cryptic, incomplete, and perhaps unclear at some points. Second, in many ways, the paper is an organized call for research rather than a report of results or a detailed plan. In that regard, I expect that some of the strategic choices the framework embodies will be controversial, and certainly much work remains to be done to answer technical questions. I emphasize that this paper does not represent Federal Reserve policy, but only my opinions about the best route to better credit risk management and capital regulation for banks.

The remainder of the paper is in five parts. Section 1 briefly reviews some basic concepts, while Section 2 discusses basic architectural choices involving accounting «mode» and time horizon. Section 3 is the heart of the paper, describing the basic framework and the various forms it might take if implemented. Section 4 discusses some practical problems of implementation. Section 5 briefly argues the framework is especially appropriate for use in Italy and other European countries. Section 6 concludes.

1.0 Basic Concepts

Conceptually, managing credit risk and allocating capital require that a bank (or its regulator) choose a soundness standard and produce an estimate of the probability distribution of portfolio credit losses conditional on portfolio composition. Typically, the soundness standard is expressed as a target probability of insolvency for the bank (see Mingo (1998) for a discussion of alternative criteria). For simplicity, in this paper I assume both banks and regulators have a similar taste for soundness.¹ For purposes of discussion, I assume banks desire a probability of insolvency due to credit risk alone of around 0.25 percent, which is consistent with a senior public debt rating by Moody's or Standard &

¹ If regulators demand a much smaller insolvency probability than banks desire for their own reasons, a regulatory capital standard may be unsustainable in the sense discussed below. In evaluating the relative tastes for soundness of banks and regulators, it is important to remember that most banks prefer to be better-capitalized than the regulatory minimum. The determinants of the size of such cushions is an important research question.

Poor's of about Baa or BBB. Readers unfamiliar with the soundness and loss distribution concepts may find useful the discussion in Appendix A.

Quantitative credit risk modeling and management is technically challenging because accurate estimates of portfolio loss distributions are very difficult to achieve. The difficulty arises because loss distributions are skewed and fat-tailed, because any individual asset in the portfolio can contribute materially to bad-tail losses, and because available credit loss experience data at the individual asset level are limited. In spite of these difficulties, available evidence implies that in reasonably well diversified portfolios, assets that pose little risk individually also contribute little to bad-tail risk, and thus that capital allocations to such assets are appropriately much smaller on a percentage basis than allocations to individually riskier assets. The fact that the current Basle Accord assesses the same 8 percent capital requirement introduces an important economic distortion.

Both regulatory capital requirements and a bank's internal capital allocations for various classes of assets are *unsustainable* if they depart materially from the economic requirements implied by reasonable soundness standards and available estimates of portfolio loss distributions. For example, if the regulatory capital requirement for a given class of assets is higher than the economic capital requirement, such assets will leave banks' balance sheets either through securitization or because banks will abandon the relevant business lines. In the limit, only those assets with economic capital requirements as large or larger than the regulatory requirements will remain. In such a limit, many banks will choose to have more capital than the regulatory minimum (because their portfolio risk postures will be relatively high) and thus will remain sound. However, over time, those banks that have only the minimum level of capital will become insolvent at a higher rate than implied by the soundness standard.

Such a «cherry picking» process is well underway in the United States, and is beginning to occur among European banks. The cherry picking process has been slowed by the fixed costs of regulatory capital arbitrage and by the fact that economic capital allocations for many classes of assets are not yet known with much precision (because portfolio loss distributions are so difficult to estimate). For example, as noted, loans to investment-grade borrowers require much less economic capital than do loans to below-investment-grade borrowers and thus investment-grade loans are rapidly leaving

banks' balance sheets. However, differences in the capital required to back loans to firms in one industry versus another industry are not yet well understood, and thus banks show much less tendency to move particular industries' loans off the balance sheet. In the short run, to be sustainable, a revised regulatory capital standard must at least apply different capital requirements to investment-grade and below-investment-grade loans. If the effect of a borrower's industry on economic capital requirements becomes well understood, regulatory capital requirements must rapidly reflect such understanding. To summarize, in order to be sustainable, both regulatory and banks' internal capital allocation schemes must adapt quickly as knowledge evolves, but should not outrun knowledge.²

2.0 Accounting Mode and Time Horizon

This paper's framework embodies a «Default Mode» (DM) rather than a «Mark-to-Market» (MTM) philosophy. That is, only credit defaults and individual asset loss severities are modeled; variations in the value of assets due to changes in credit quality short of default are assumed unimportant. Although a case can be made for MTM approaches, I chose a DM approach for its simplicity, feasibility, and because I believe it is more appropriate for traditional banking books (most credit risk borne by most banks still is in the traditional banking book, especially the loan portfolio). A DM approach simplifies the framework because most fluctuations in value need not be modeled. Similarly, implementation is made more feasible by a DM approach because the modeling burden on

² I suspect that at least in the United States, in order to be sustainable, marginal regulatory capital requirements for any given asset can deviate from banks' beliefs about economic capital requirements by no more than about 2 percentage points. For example, a regulatory requirement of 8 percent for an asset that needs to be backed by only about 6 percent of own funds would probably provide sufficient incentive for a bank to securitize the asset. This guess is based on the following intuition: if the required return on equity capital is 15 percent and the interest rate on banks' marginal debt liabilities is 5 percent, then a capital requirement error of 2 percent implies a «tax» on the asset of $0.02 \times (0.15 - 0.05) = 20$ basis points. In the United States, such a «tax» probably exceeds the costs of securitization. In countries with less developed capital markets, costs of securitization may currently be higher and thus a larger error in setting regulatory capital requirements may currently be sustainable. However, capital markets are developing very rapidly in most developed countries, especially in Europe. The size of the sustainable error is an important research question because it influences the level of precision that regulatory capital requirements must achieve.

banks is enormously reduced. This is because, in reality, MTM is a Mark-to-Model approach. The vast majority of commercial banks' assets (and off balance sheet positions) rarely trade and the prices of any trades that do occur are usually unobservable. Thus, current «market» values of assets must be approximated. Given the lack of transaction price data, market value approximation algorithms are very difficult to develop and validate and would be a major source of operational error in MTM-based capital allocation procedures.

The appropriate time horizon for portfolio loss distributions is an important question. A one-year horizon is conventional, but bad-tail loss rates for a given portfolio are likely to be much higher over a multiple-year period than a one-year period. Bad-tail events are typically associated with business cycle downturns, currency crises, or other systematic events. The effects of such events typically last for more than a year. To the extent that banks can easily replenish capital lost during a given year (or shorter period), then a short analysis period is appropriate. However, earnings are usually low following systematic events and outside equity is usually expensive and difficult to obtain. Thus, time horizon is an important topic for research. The remainder of the paper uses a generic, unspecified time horizon unless otherwise noted.

3.0 The Framework

For illustrative purposes, consider a portfolio of three assets for which individual and joint probabilities of default are known and appear in matrix A, and dollar losses in the event of default are also known and appear in matrix B. p_{11} , p_{22} , and p_{33} are the probabilities that only assets 1, 2, and 3

$$A = \begin{bmatrix} p_{11} & p_{12} & p_{13} \\ p_{12} & p_{22} & p_{23} \\ p_{13} & p_{23} & p_{33} \end{bmatrix} \quad B = \begin{bmatrix} L_1 \\ L_2 \\ L_3 \end{bmatrix}$$

default, respectively. The probabilities of joint default events involving any

two assets, which correspond to the «correlations» frequently referred to in discussions of credit risk modeling, are p_{12} , p_{13} , and p_{23} . In order to keep the matrix two-dimensional and thus notationally convenient, not shown is the probability that all three assets default, which is $p_{123} = p_{12}p_{23} = p_{12}p_{13}$.³ If A were modified by multiplying off-diagonal elements by 1/2, and I is a 3-element unit vector, the expected loss on the portfolio would be $I'(AB) + p_{123}I'B$.

Tracing out the loss distribution for the portfolio involves rank-ordering the possible loss permutations and tallying the associated probabilities. For example, assuming all elements of B are positive, the maximum loss is $I'B$ with probability p_{123} . The basic elements of this setup generalize to portfolios with any number of assets.

In the ideal case, good estimates of all elements of A and B would be available and portfolio loss distribution modeling would simply be an exercise in combinatorial computation. However, there is in general no hope of achieving the ideal case because, in a portfolio with N assets, the number of joint default probabilities to be estimated is $(N^2 - N)/2$. Given the relative infrequency of default events, even for rather small real-world portfolios there is no hope of estimating the elements of A directly from loss experience data. Existing credit risk modeling systems in effect estimate the elements of A by making assumptions about the relationship between joint default probabilities and various systematic factors. However, no consensus about the proper assumptions exists. This uncertainty is one of the primary barriers to a consensus about how to model credit risk.⁴

A primary assumption of this paper is that in most cases, reasonable approximations of subportfolio bad-tail risk will result from grouping similar assets into subportfolios («blocks») and assuming that elements of A (and perhaps B) are constants. Imagine a huge pair of matrices, similar to A and B, which reflect all credit risk positions for a large portfolio. If similar assets are grouped

³ To obtain the probability of any default for a given asset, individual and joint default probabilities must be summed. For example, the probability of any default by asset 1 is $p_1 = p_{11} + p_{12} + p_{13} + p_{123}$.

⁴ If it were possible to construct standard errors for estimates of the p_{ij} produced by factor models, it seems likely that in most cases a hypothesis that all p_{ij} have the same value could not be rejected.

together in the matrices as submatrices or blocks, they can be analyzed separately, and the results aggregated to obtain an overall portfolio loss distribution. Such a grouping strategy is similar in spirit to CreditRisk+'s breakdown of a portfolio's assets into exposure bands and sectors, but this paper's framework is more flexible because I envision that some blocks would be modeled simply and some in a more complex fashion. For example, parameters need not be assumed constant for all assets or even constant within a given block. Over time, as knowledge about joint default probabilities improves, I expect assumptions of constancy would be dropped for more and more blocks.⁵

An example helps show how groupings and an assumption that parameters are constants simplifies matters. Suppose that the example three-asset portfolio described in matrices A and B consisted only of commercial and industrial (C&I) loans rated the equivalent of BBB-, that the loans are all for amounts close to \$1 million, and that loss severities in event of default are 30 percent (so that all elements of matrix B equal \$300,000). Suppose further that the joint default probabilities for the loans (the off-diagonal elements of A) are identical at $p_{ij} = 0.1$ percent. Then assuming the probability of any default by one of the assets is 0.35 percent (roughly consistent with the BBB- ratings), a tally of permutations yields the following subportfolio loss distribution:⁶

<u>Loss</u>	<u>Probability</u>
\$900,000	$p_{ij}^2 = 0.01\%$
\$600,000	$3 * p_{ij} = 0.30\%$
\$300,000	$p_{ii} = 3(p_i - (p_{ij} + p_{ij} + p_{123})) = 3(.35 - (.1 + .1 + .01)) = 0.42\%$
\$0	99.27%

In general, in a block of N similar assets, the loss rate at any given percentile of the loss distribution will be a function of N, L, and the parameters p_{ii} and p_{ij} . Assumptions that p_{ii} and p_{ij} are constants within a given block drastically reduce the dimensionality of the estimation problem making it possible that the

⁵ Of course, the computational method used to get the loss distribution for a block will differ according to whether all parameters are constant or parameters vary.

⁶ The values of the joint default probabilities (i.e., the correlations) clearly can have a large influence on estimated capital requirements. For example, if $p_{ij} = 0.01$ rather than $p_{ij} = 0.1$, the probability of a loss of more than \$300,000 would be about 0.03 rather than 0.30.

constants p_{ii} and p_{ij} could be estimated using available volumes of individual asset loss experience data (but I have not yet worked out the details of how to do so).

Data to estimate even constant joint default probabilities may become available only over time for some types of assets. Thus, there is a need for generic joint-default probabilities applicable to any type of debt instrument with a given credit rating. The CreditRisk+ technical document suggests that p_{ij} can be approximated by a function of expected default rates, which are easier to estimate than joint default probabilities. More research on the best way to obtain generic parameter values is needed.

An estimate of a bank's overall credit loss distribution requires that subportfolio loss distributions be combined. In the simple case of independent subportfolios, the combined distribution would involve a straightforward aggregation of the subportfolio distributions. But independence cannot be assumed---for example, if losses are very high for a portfolio of \$1 million BBB-rated C&I loans, they are also likely to be high for \$500,000 BB-rated C&I loans. I have not worked out how cross-subportfolio correlations should be estimated, but the CreditRisk+ system includes methods for doing so, and Koyluoglu and Hickman (1998) compare some alternative methods. Research on the sensitivity of results to the aggregation method is needed. Data on aggregate loss experience data for each loan type is likely to be useful in developing and testing such methods.

Why compose subportfolios according to asset size, rating and type? Why not simply assume constant parameters for the portfolio as a whole? Because an assumption of constant joint default probabilities may be unrealistic when made across different types of assets and different ratings. For example, joint default probabilities for U.S. commercial real estate loans may be substantially higher than for U.S. C&I loans. Joint default probabilities are surely higher for groups of individually risky assets than individually safe assets.⁷

By initially assuming parameters are constant within each block, many banks could compute

⁷ Separating subportfolios by asset size may turn out to be merely an expositional convenience--it may turn out that the loss distribution for a given subportfolio of R-rated, type-X assets can be derived from a generic R-rated, type-X subportfolio loss distribution and a function of the number of assets in the given portfolio and their size distribution. Research leading to a convenient closed-form version would be helpful.

overall portfolio capital requirements for credit risk (and capital allocations for assets in each subportfolio) without any formal modeling at all. A consortium of banks and regulators could calculate reference tables of capital requirements for each generic combination of asset type and rating grade. Bank regulators are the natural coordinators of such a consortium because of their interest in developing uniform regulatory capital requirements and their ability to coordinate the necessary data-gathering and analysis. Indeed, a reference-table approach is currently the most common way of making operational capital allocations at large U.S. banks. A reference-table approach is also somewhat similar to the existing Basle capital standard and thus would be relatively easy to implement by regulators during a transition period, but I expect the reference tables would involve many more «buckets» than the current Basle standard.

However, it is already obvious that an approach involving a single set of reference tables would be inadequate today for some banks, and is very likely to become less useful as time passes. This is because the key assumption, that joint default probabilities are positive constants within blocks, is certainly inappropriate for one important class of assets: credit hedges. Moreover, because such hedges are constructed to be negatively correlated with other portfolio assets, it is both feasible and essential that such correlations be explicitly considered on a hedge-by-hedge basis in order to properly adjust capital requirements and allocations. Credit derivatives that pay off in event of default of a given loan are the clearest example: the payoff correlation of the loan and its associated derivative is near -1. Although such simple cases might be handled by offsets -- by reducing the value of L for the loan, for example -- offset treatment appears too difficult to manage when hedge positions involve many assets or macroeconomic variables like changes in GDP. The impact on loss distributions of such more complex hedges must be modeled.

A bank could make the transition from use of reference tables to modeling by implementing a model that simulates the constant parameter case for each block. That is, initially the bank's model would be calibrated to reproduce the reference-table capital allocations for each block. The bank could then relax the constant-parameter assumption in individual blocks or alter the boundaries of blocks as appropriate. In principle, different model structures could be used to handle different blocks.

As an arbitrary example, a bank might model blocks of domestic C&I loans using CreditRisk+ and blocks of international loans using a CreditMetrics approach. One virtue of this paper's flexible framework is that operational experimentation with different specific modeling strategies could be accommodated naturally.

4.0 Implementation Issues

In order to produce reasonably accurate capital allocations, a system like that described above must satisfy several conditions in addition to those mentioned previously. The credit ratings used to divide assets into blocks must classify assets by default probability or expected loss with sufficient accuracy; the subportfolios represented by each block must be reasonably large and well-diversified; reasonably accurate measures of the exposure of each position must be used, including effects of any credit-related optionality; and, of course, reasonably good estimates of parameters such as joint default probabilities will be needed. More research is needed before these conditions can be satisfied. In this section, I sketch the issues.

4.1 Credit Ratings

Credit ratings measure the risk of default or loss on individual assets. As described in Treacy and Carey (1998), most large U.S. banks operate internal rating systems in which bank staff assign ratings judgmentally, in much the same way that Moody's or Standard & Poor's assign senior public debt ratings. Ratings can also be assigned by statistical models, like the Zeta model or KMV's CreditMonitor. The relative merits of judgmental and statistical systems remain an open question---on the one hand, judgmental systems can produce biased ratings, especially where the rater's business volume or compensation is a function of the ratings he assigns. On the other hand, statistical systems cannot incorporate the inside knowledge that bank staff often have about borrowers.

Ratings are a key input to capital allocation models because a given asset's contribution to bad-tail loss rates is a function not only of that asset's default correlation with other assets, but also of the expected default rate or expected loss rate on the asset, that is, of its rating. Much of the gain in precision in capital allocation from implementing the simplest system described in this paper would come

from the division of differently rated assets into different blocks, with the parameters in each block depending very much on the associated rating. As a general rule, the more separate grades on the rating scale and the more accurate the rating assignments, the more accurate the estimates of capital requirements.

Table 1. Estimated Effects of Rating Assignment Errors on Capital Allocations

Simulated Rating Assignment	Actual Distribution of Loans Across Grades (percent)					99.75th Percentile Loss Rate
	A	BBB	BB	B	<B	
BBB	0	100	0	0	0	2.26
BBB	10	80	10	0	0	2.96
BBB	20	60	20	0	0	3.44
BBB	33	34	33	0	0	3.97
BB	0	0	100	0	0	6.44
BB	0	10	80	10	0	6.47
BB	0	20	60	20	0	6.18
BB	0	33	34	33	0	5.68

Conversely, if rating assignments are not very accurate at a particular bank---that is, if a given loan is often assigned a grade associated with a default or loss rate much different than the default or loss rate on the loan---then that bank's capital allocations are likely to be distorted. Table 1 provides a

sense of the impact of different magnitudes of systematic rating errors. I use the Monte Carlo method of Carey (1998) to estimate bad-tail loss rates at the 99.75th percentile of a generic loss distribution for randomly selected portfolios of actual privately placed bonds with different ratings. I assume that the ratings appearing in the data are correct, and simulate rating errors by including in a portfolio that nominally had only loans with a given grade some loans that actually had different grades. For example, in the first row of Table 1, a \$1 billion portfolio of loans all actually rated BBB has an estimated bad-tail loss rate of 2.26 percent. In the second row of the table I simulate a symmetric twenty percent rating assignment error by building portfolios that actually included 10 percent A-rated loans and 10-percent BB-rated loans, even though the portfolio nominally includes only BBB-rated loans. This error raises the bad-tail loss rate to 2.96 percent. The fourth row of the table shows the results of a much larger systematic error, reporting an exercise in which one-third of the nominally BBB portfolio is actually A and one-third is BB. This raises the bad tail loss rate to 3.97 percent. In this exercise, increases in the symmetric grading error rate increase the bad-tail loss rate because such errors increase the portfolio expected loss rate. The expected loss rate increases because the difference between BB and BBB expected loss rates is larger than the difference between A and BBB expected loss rates. Thus, adding equal amounts of A and BB rated assets to a BBB portfolio raises the expected loss rate.

In the second panel of the table, which simulates rating errors for a portfolio of BB-rated assets, the opposite trend occurs---bad tail loss rates decline as rating errors increase. Similar to the BBB case, this occurs because, in the data upon which the simulation is based, the BB vs. B difference in expected loss rates is smaller than the BB vs. BBB difference.

Overall, Table 1 gives the impression that even rather large symmetric errors in assigning ratings have only moderate effects on capital allocations. These results are very preliminary, but on the whole they indicate that extremely accurate rating assignments will not be required to support good capital allocation systems. However, systems that produce biased rating assignments could lead to materially biased capital allocations, and thus means of controlling biases will be required.

4.2 *Limits on Concentration Will Be Required*

The primary simplifying assumption advocated in this paper---that joint default probabilities can

be approximated with constants for any given subportfolio of traditional bank assets---will support accurate capital allocations only if subportfolios are reasonably large and well-diversified and if true joint default probabilities do not differ too much across reasonably well-diversified portfolios. Suppose a subportfolio consisted only of loans to African copper producers. Obviously, such a subportfolio by itself would represent a very concentrated risk and its joint default probabilities would be high, much higher than for a subportfolio of randomly selected loans.

As a practical matter, most large banks maintain diversified portfolios by limiting their aggregate exposure to borrowers in a given country, industry, etc. Such limits would be a necessary condition for the simplified capital allocation systems described in this paper to work well. Moreover, research to determine the appropriate ranges for such limits is needed. Is a subportfolio still well-diversified if 5 percent of its exposure is to firms in a single industry? What if the exposure is 20 percent? Where should the line be drawn for purposes of ensuring that the approximations discussed above work well?

Tight limits may not be a sufficient condition for the approximations to work well. If average joint default probabilities vary widely for different reasonably well-diversified subportfolios, then an assumption of constant probabilities for a given rating and asset class will introduce material errors into capital allocations. Sensitivity to diversification is a very difficult hypothesis to test properly for the same reasons that individual joint default probabilities are very difficult to estimate. However, I believe some useful sensitivity tests can and should be done by Monte Carlo analysis of bond experience data, and advocates of models like CreditMetrics and Portfolio Monitor almost surely can contribute useful insights about this question.

4.3 *Exposure Measurement*

In the example given above, the loss rates L in matrix B were assumed identical across equal-size loans. This assumption is obviously unrealistic. For example, several studies find that recovery rates on defaults of subordinated debt are small than for senior debt. Moreover, a bank's exposure at the beginning of a period may differ substantially from exposure at the time of a default during the period. Amortization of principal may reduce exposure on a term loan, or a distressed borrower may

make large drawdowns on its line of credit prior to default. Studies that illuminate how such dynamics should be handled are required.

4.4 *Parameter Estimation*

Those engaged in credit risk modeling are uniformly frustrated by the lack of large panel datasets detailing the credit risk experience of large numbers of individual loans over long time periods. Such datasets are required by even the simplest capital allocation systems discussed above. In the United States, datasets covering up to a decade or so of experience have been built by a few banks and insurance companies, but these are proprietary. A few data-gathering consortiums exist and have begun to report results (for a publicly available example, see Society of Actuaries (1996)). Data in one form or another are available from the rating agencies Moody's and Standard and Poor's. However, a common feature of such datasets is their focus on bond or commercial loan experience for U.S. obligors.

Data availability varies in other countries. It is my understanding that private firms, central banks, or other entities have built large databases for loan experience in several countries (in Japan, the Teikoku Databank is an example). Again, however, such databases appear mainly to cover domestic commercial loans. Data on real estate loans, consumer loans, cross-border loans, etc., where they exist, do not appear to have been used for analysis of capital allocations.

I suspect that capital requirements for well-diversified portfolios of any type of asset in any country will turn out to be largely a function of subportfolio expected loss rates, as in Table 1 of this paper. If that is correct, then the best databases available anywhere in the world can be used to estimate capital allocation model parameters. In order to establish this fact, however, coordinated analysis of loss experience using available databases in several countries appears necessary. Central banks and other banking agencies appear to be the natural agents to coordinate such analysis.

4.5 *Other Issues*

Implicitly, in this paper I have focused on capital allocation systems and capital regulations for large banks in developed countries. The form of capital allocation systems and capital regulation for small banks, or for banks of any size in developing countries, is not clear. The internal rating systems

that are crucial to quantitative capital allocation at large banks may involve fixed costs that are too large for the smallest banks to bear, and in addition small banks' loan portfolios may be rather homogenous in terms ex ante default probabilities and expected loss rates. Moreover, it seems likely that many small banks' and developing country banks' portfolios are rather undiversified. Research on the portfolio loss characteristics of such banks is needed.

5.0 Credit Risk Modeling for Italian Banks and Banking Agencies

The general framework suggested in this paper is appropriate for any country, but perhaps especially so for Italy and many other European countries because of the rapid changes occurring in financial markets. In Italy, the most important challenges flow from the relatively recent opening of capital markets and from the advent of the single European banking market and the Euro.

The implementation of quantitative capital allocation systems is particularly important in Italy because its large banks are likely to both enter new markets across Europe and internationally, and because competition from foreign banks entering Italy will intensify. Disciplined systems for measuring risk and return will help limit unprofitable mis-steps.

However, many existing credit risk models are difficult to implement in Italy because of data limitations. For example, data on the loss characteristics of foreign obligors will be hard to obtain in many cases. Moreover, equity market return timeseries are of short duration or do not exist at all for many Italian firms, and equity data is a primary source of parameter values for models like CreditMetrics or PortfolioManager.

Thus, a flexible framework that supports implementation of reasonably good capital allocations for most assets and more precise allocations where data and techniques are supportive seems especially desirable in the Italian case. The framework suggested here is obviously not ready for use---clearly much research remains to be done---but the prospect of a revision of the Basle Accord may cause results of such research to appear over a period of months or years rather than decades.

6.0 Concluding Remarks

This paper describes a flexible framework for internal capital allocation systems at banks and for regulatory capital standards for credit risk. The framework is flexible in that it supports both simple, low-marginal-cost methods of estimating portfolio loss distributions as well as evolution toward more complex and precise methods as such methods are developed and validated.

One benefit of the flexible framework that received little mention above is its applicability to all credit-risk positions on and off a bank's balance sheet. At most banks today, the focus of credit risk modeling is on commercial and international loan portfolios and on derivative instruments. However, consumer loans, such as credit card and mortgage loans, clearly pose credit risks that should be properly weighed by both internal and regulatory capital allocation systems.

To implement the framework proposed here, or any other system, much research and development clearly is required, but I believe the most important elements of such R&D can be completed over a relatively short period of a couple of years if sufficient resources are brought to bear in a cooperative fashion by the world's major banks and bank regulators.

Appendix A. Soundness, and Portfolio Loss Distributions

Over any given time period, such as a quarter or year, a typical large bank will experience some credit defaults and losses on its portfolio positions, but in most periods the losses will be relatively modest and within the normal range for that portfolio. Occasionally, however, even a well-managed bank will suffer very large losses. If the bank's capital is insufficient to absorb the losses (as well as losses during the same period from noncredit risks) it will be technically insolvent and is likely to close. Even if the bank survives, the value of its franchise is likely to suffer as a result of market and regulatory discipline.

Conceptually, managing the risk of insolvency requires two things of a bank (or regulator). First, a decision about the acceptable probability of its insolvency. Zero is not a feasible decision, because that would require the bank to be financed entirely by equity capital, in which case the bank would no longer be a bank. As a practical matter, most banks appear to find it worthwhile to maintain a low probability of insolvency in order to maintain access to relatively low-cost sources of debt finance. In terms of U.S. public debt ratings, most banks appear to prefer to maintain at least an investment-grade rating for their long-term liabilities, that is, a Moody's or Standard and Poor's rating of Baa3 or BBB- or better. Other choices (and criteria other than insolvency) are possible (see Mingo (1998)), but for simplicity in this paper I assume both banks and regulators desire at least investment-grade soundness.

Second, at the beginning of each period, the bank must estimate a probability distribution of credit losses for its portfolio. An example of such a distribution (actually, a density) appears in Figure 1: the area under the curve to the right of a given loss percentage on the horizontal axis is the probability that the bank will lose as much or more than the given percentage. Credit loss distributions for debt portfolios are skewed to the right and fat-tailed relative to the normal distribution because debt instruments have a limited upside but exposure the lender to the risk of total loss on the position. Moreover, the precise shape of such distributions is sensitive to factors like portfolio size, the standalone credit quality of portfolio positions as measured by expected losses, and diversification. Accurate estimates of such distributions are quite difficult to obtain empirically. It is this difficulty that makes

quantitative credit risk modeling and management so challenging from a technical standpoint.

Suppose a bank has estimated its portfolio loss distribution and wishes to have a probability of insolvency during the coming period no larger than 0.25 percent. The bank's capital requirement in percentage terms is given by that loss rate such that the chance of larger losses is 0.25 percent. If available capital is less than the required amount, the bank must either raise new capital or alter its portfolio in a manner that reduces bad-tail risk.

References

- Carey, Mark, 1998, Credit risk in private debt portfolios, *Journal of Finance* 53:4, 1363-1387.
- Credit Suisse Financial Products, 1997, CreditRisk+: A credit risk management framework.
- Gordy, Michael B., 1998, A comparative anatomy of credit risk models, *Journal of Banking and Finance*, forthcoming.
- Gupton, Greg M., Christopher C. Finger, and Mickey Bhatia, 1997, *CreditMetrics - technical document*, (New York: J.P. Morgan), www.jpmorgan.com.
- Koyluoglu, H. Ugur, and Andrew Hickman, 1998, A generalized framework for credit risk portfolio models, Working Paper, Oliver, Wyman & Company (New York).
- Mingo, John, 1998, Policy implications of the Federal Reserve study of credit risk models at major U.S. banking institutions, *Journal of Banking and Finance*, forthcoming.
- Robert Morris Associates and First Manhattan Consulting Group, 1997, *Winning the Credit Cycle Game: A Roadmap for Adding Shareholding Value Through Credit Portfolio Management*.
- Society of Actuaries, 1996, *1986-92 Credit risk loss experience study: private placement bonds* (Schaumburg, IL).
- Treacy, William F., and Mark Carey, 1998, Credit rating at large U.S. banks, *Federal Reserve Bulletin*, November.
- Wilson, Thomas C., 1997, Measuring and managing credit portfolio risk, *Journal of Lending and Credit Risk Management*, two parts, appearing in July and August.