Internal Risk-Management Models as a Basis for Capital Requirements

by Daniel A. Nuxoll*

In 1988, after extensive negotiation among the G-10 central-bank governors, the Basle Committee on Banking Regulations and Supervisory Practice agreed on uniform capital standards. The agreement, known as the Basle Accord, was an attempt to produce uniform capital standards for internationally active banks.[1] Until then, different countries had set different capital standards for their banks. In some countries, lower standards were thought to give a competitive advantage to banks headquartered in those countries.

Although the focus of the Basle Accord is on uniform capital requirements, the Accord also establishes risk-based capital requirements that supposedly reflect the actual credit risks faced by a bank. The Basle Accord assigns risk weights to all assets, weights that should reflect the relative risks of those assets. For example, commercial and industrial loans have a 100 percent weight, while home mortgages have a 50 percent weight. Weights are also assigned to off-balance-sheet activities, such as loan commitments and standby letters of credit.

Banks are required to hold 8 percent of the risk-weighted assets as capital.[2] This means, for instance, that banks are required to hold 4 percent (= 8 percent x 50 percent) capital for home mortgages, and 8 percent (= 8 percent x 100 percent) capital for commercial and industrial loans.

From the start, however, analysts have repeatedly demonstrated that the risk-based capital requirements do not accurately reflect risk. First, the risk weights themselves are not accurate: while studies generally indicate that the risk weights are not completely unreasonable, these same studies inevitably find that some category of loan has the wrong weight. Second, loans within a given category are not equally risky—yet a 90-day inventory loan to a profitable company with a solid credit record has the same risk weight as a five-year loan to develop commercial real estate. Third, the requirements ignore risk-reduction activities like diversification and hedging; thus, a portfolio of loans to borrowers in the same industry and the same area (for example, farmers in the same county) has the same capital requirement as a portfolio diversified across industries and regions of the country. Presently, therefore, almost everyone acknowledges that the Basle risk-based capital standards have very little to do with actual risk.

But although the Basle standards fail to reflect accurately the risk in a bank's portfolio, banks do have definite incentives to measure accurately the risk of their activities. In the past decade, banks have developed internal risk-management models to measure their risks systematically. These models are based on the

* Daniel A. Nuxoll is a financial economist in the FDIC's Division of Research and Statistics. Conversations with Jack Reidhill, George Hanc, Steven Seelig, Miguel Browne, and John Feid resulted in significant improvements in this paper.

[1] U.S. regulators have imposed these requirements on all banks.

[2] The bank also has a 4 percent Tier I capital requirement; Tier I capital is approximately equal to equity. Total capital includes Tier I and Tier II capital; the latter consists of loan-loss reserves and certain forms of nondeposit debt.

best statistical estimates of the particular risks being measured, and most models consider the effects of diversification and hedging. The possibility of replacing the Basle capital standards with these models has been widely discussed.[3]

This article considers three proposals to revise the Basle Accord. The first, which was adopted in 1996, permits banks to use internal models to estimate one kind of risk—the risk of trading activities. The second would permit banks to use somewhat different, but basically fairly similar, models to evaluate the risk of making loans. The third would permit banks to use any method to estimate their own risk, but—in contrast to the current systems—banks that underestimated the risk of their activities would be penalized.

Trading-Book Models (Market-Risk Models)

Regulators had long recognized that the Basle standards for market risk were inadequate, even while banks and securities firms had developed sophisticated methods of measuring the risk of their portfolios. Accordingly, the 1996 amendments to the Basle Accord permitted regulators to accept the calculations of the banks' internal risk-management models in setting capital requirements for the market risk in banks' trading portfolios.

Trading-book models have become increasingly common in banks, especially since 1994, when J. P. Morgan released its RiskMetrics model.[4] J. P. Morgan has distributed this model widely; some components are available over the Internet. Although other risk-management models are available, RiskMetrics has become the standard of comparison, and the other models use very similar methods. Thus, most of the discussion below applies directly to RiskMetrics; nevertheless, it ignores many aspects of this model and focuses on the deficiencies of the simplest configuration of the RiskMetrics models. The discussion also mentions different methods RiskMetrics can use to avoid some of those deficiencies.

Internal risk-management models generally estimate the value at risk; hence they are often called VaR models.[5] The value at risk is the amount of money that would be sufficient to cover most potential losses. Because VaR models focus on risk, they generally ignore profit.[6]

Trading-book VaR models use estimated probabilities of price movements to estimate the probability of losses for the whole portfolio. The data might show, for example, that over the past 20 years, the price of 10-year Treasury bonds decreased by more than 0.2 percent on only 5 percent of the days. The data might also indicate that the price decreased by more than 0.75 percent on only 1 percent of the days.

Given these data, if a portfolio consists solely of $100 million of 10-year Treasury bonds, then the VaR model would estimate that losses would exceed $200,000 (= $100 million x 0.2 percent) in a day less than 5 percent of the time. This number could be called the 5 percent value at risk because losses would exceed $200,000 less than 5 percent of the time. The model would forecast that losses would be greater than $750,000 (= $100 million x 0.75 percent) less than 1 percent of the time. Similarly, this number could be called the 1 percent value at risk—a bank could be 99 percent certain that this investment would never lose more than $750,000.

To estimate a Value at Risk model, one needs not only a probability level but also a time horizon. For regulatory purposes, the time horizon is ten trading days and the relevant probability is 1 percent, so the losses within ten days would exceed the value at risk less than 1 percent of the time.

The 1996 Basle amendments require that capital equal three times the value at risk.[7] If the bank's in-

[3] This topic was extensively discussed at a 1998 conference in New York City sponsored by the Federal Reserve Bank of New York, the Bank of England, the Bank of Japan, and the Board of Governors of the Federal Reserve System (Federal Reserve Bank of New York [1998]). In addition, John J. Mingo (1998) has critiqued the current standards and spelled out the benefits of a models approach to bank capital. The International Swaps and Derivatives Association has endorsed an approach that is a definite step toward using internal models to set capital requirements (summarized in Elderfield [1998]).

[4] For the definitive description of RiskMetrics methodology, see Longerstaey and Zangari (1996).

[5] The explanation here of trading-book VaR models applies generally to banking-book VaR models as well, except that the emphasis in the latter is on loans, not securities. Throughout this section, the terms "trading-book model" and "VaR model" are used interchangeably to refer to the trading-book subset of VaR models.

[6] Jorion (1997) discusses VaR models in detail. VaR models are closely related to risk adjusted of return on capital (RAROC) models, originally developed by Bankers Trust; for a summary of Bank of America's RAROC model, see Zaik, Walter, Kelling, and James (1996).

[7] This multiplication factor has been vociferously criticized as arbitrary, usually by banks that want a lower factor. However, the usual rule is that the risk increases with the square root of time because standard deviation increases with the square root of time. Consequently, if the relevant horizon is one year (250 trading days) instead of ten trading days, the multiplication factor should be five. This multiplication factor, too, is undoubtedly incorrect. *See* Danielsson, Hartmann, and deVries (1998) and Stahl (1997) for opposite sides of the debate. The whole debate demonstrates that no one really knows the "correct" multiplication factor and that the bank regulators have probably been conservative.

ternal model is discovered after-the-fact to have been inadequate, the capital requirement can be increased, reaching a maximum of four times the value at risk. If the performance of the model turns out to have been grossly inadequate, regulators can refuse to use it to set capital requirements.

VaR models can include the effects of both diversification and hedging, which are common methods of controlling risk. Diversification and hedging are possible because prices do not necessarily move together. On the one hand, if prices always moved together, losses from one investment would never be offset by profits from another investment, and neither diversification nor hedging would be possible. On the other hand, if prices always moved in opposite directions, losses from one investment would always be offset by profits from the other, and perfect hedging would be possible. The usual case is that prices sometimes move in opposite directions and sometimes in the same direction, so that losses are sometimes offset by profits from other investments.

The way VaR models incorporate the effects of diversification and hedging is by estimating the correlations between price changes. If two prices always move together, the correlation equals one; if they always move in opposite directions, the correlation is negative one. In fact, securities prices tend to move together, but they do not always move together, so almost all estimated correlations are greater than zero, but less than one. Consequently, diversification generally reduces risk.[8] Importantly, assumptions or estimates of these correlations amount to measurements of the effects of diversification and hedging, and they translate directly into estimates of the riskiness of the portfolio.

The significance of diversification and hedging was evident in the first stages of the thrift crisis (before 1983). Thrifts held a large number of very safe securities, namely fixed-interest-rate mortgages. However, thrifts were not well diversified, because changes in interest rates affect the value of all fixed-rate mortgages. The high interest rates in the late 1970s and early 1980s drove the value of these mortgages down. At the same time, thrifts had to pay higher interest rates to obtain deposits. The result was that many thrifts faced insolvency by 1983.[9] (Of course, the evolution of the thrift crisis after 1983 had little to do with interest rates.)

* * *

Because of the large number of traded securities, trading-book VaR models inevitably use a large number of simplifying assumptions. We now examine five of them, noting any evidence on whether they produce overly large or overly small estimates of value at risk.

The first simplifying assumption VaR models commonly make is that price changes are distributed normally (they make this assumption because the normal distribution is easy to handle mathematically). The actual distribution of price changes and financial data, however, is generally not normal. Specifically, the normal distribution understates the probability of large price changes. VaR models that use the assumption of normality therefore understate the probability of large price declines, thereby understating the probability of large losses.

Many fixes have been suggested to solve this problem, but for regulators the problem is nothing to worry excessively about: VaR models estimate the probability of large losses during a day, whereas regulators are undoubtedly more concerned about the long run. Though the normal distribution might be misleading for managing risk on a daily or weekly basis, in the long run even financial data are distributed normally.[10]

[8] Mathematically, diversification is possible when a correlation is positive and less than one. Hedging is possible when a correlation is negative. Hedging often depends on instruments like futures which permit a trader to sell "short," that is, sell for future delivery at a fixed price. The value of these contracts moves in the opposite direction to the current price. Short selling essentially turns a positive correlation into a negative correlation. Most descriptions of hedging assume that the correlation is negative one, though most actual hedges involve basis risk, which occurs because the correlation does not equal negative one.

[9] This is an instance of a classic problem, sometimes called interest-rate risk, sometimes called the mismatch of maturities. Because deposits have short maturities and mortgages have long maturities, an increase in interest rates drives up the cost of borrowing without affecting the return on loans. The classic solution to the problem is to fund long-term loans with long-term deposits. This solution could be considered a hedging program that protects a bank from interest-rate risk. The bank's interest costs and interest income are more closely correlated when maturities are closely matched.

[10] Jackson, Maude, and Perraudin (1997) use portfolios from actual banks to illustrate the nature of the problem. Conventional wisdom is that alternative techniques known as historical simulation or structured Monte Carlo simulation can solve the problem (Jorion, [1998]). However, these methods have problems of their own, which have motivated the development of more exotic fixes (see Danielsson and de Vries [1997]; Zangari [1997]; Hull and White [1998]; and Rubinstein [1998]). Duffie and Pan (1997) note that the central limit theorem implies that, in the long run, returns from even a fat-tailed distribution are normal.

A second simplifying assumption, one made by most VaR models, is the use of only a small number of estimated correlations between price changes. The number of actual correlations increases dramatically with the number of securities because for each pair of stocks, a different correlation must be estimated. For 2 securities, there is 1 correlation; for 4 securities, 6 correlations; and for 100 securities, 4,950 correlations.[11]

To avoid estimating a multitude of correlations, some models arbitrarily specify some correlations to be zero or one. Because most true correlations are probably somewhere in this range, using a correlation of one overstates the risk because it neglects the effects of diversification; but zero correlations understate the risk by overstating the possibilities for diversification.

Other VaR models avoid estimating a multitude of correlations by using the results of multifactor models of prices. Multifactor models assume that securities prices are driven by a limited number of factors. If the prices of two securities both respond to some factor, then the prices are correlated. Often, multifactor models assume that the prices of all firms in an industry move together, so they treat a firm's industry as a factor. The correlations between price changes are then a function of a handful of factors.

These methods of avoiding having to estimate a multitude of correlations might produce high or low estimates of correlations, so they might produce an understatement or overstatement of the risk eliminated by diversifying and hedging. Regulators must be concerned because systematic underestimation of the correlations produces an overestimation of the benefits of diversification, an underestimation of the amount of risk, and a consequent underprovision of capital.

A third simplifying assumption is the use of historical data to estimate the relationship between prices. Historical data often significantly understate some risks. Options prices, in particular, behave very differently when they are substantially out-of-the-money from when they are in-the-money. (Out-of-the-money options are options that will be exercised only if there is a big price change, in-the-money options will be exercised even if there is no price change.) The price of a call option on a stock might be very stable, if the current stock price were substantially below the strike price. However, if the stock price were to rise above the strike price, the price of that option would fluctuate much more.[12]

A more mundane example of the understatement of risk when historical data are used concerns prepayments of home mortgages and the price of mortgage-backed securities. Interest rates can affect prepayments of home mortgages, and prepayments are a major determinant of the price of mortgage-backed securities. When interest rates decline modestly, the cost of refinancing prevents most home owners from refinancing, but when rates decline significantly a wave of prepayments is almost certain to follow. A 0.5 percent decrease in interest rates would probably have a small effect on prepayments, but a decrease of 2 percent would almost certainly increase prepayments substantially. The effect of a 2 percent decrease in interest rates is not simply four times the effect of a 0.5 percent decrease. Thus, one cannot directly infer how a 2 percent decrease in interest rates will affect the price of mortgage-backed securities from the results of a 0.5 percent decrease. A period of relatively stable interest rates without any dramatic changes does not reveal the true risk of mortgage-backed securities.

What makes the problem created by the use of historical data especially acute is that many VaR models are estimated on the basis of only the most recent data. RiskMetrics can use long data series, but the most recent data receive more weight when the model is estimated—practitioners argue that only the most recent data reflect current market conditions. Recent periods, however, like most periods, tend to be relatively stable, with no dramatic changes. This complicates things for VaR models, which attempt to estimate a firm's losses as a result of unlikely events (a 1 percent VaR estimate is concerned only with events that happen less than 1 percent of the time). For many assets, the small price changes that occur in a stable market simply cannot serve as the basis for an estimation of the effect of market turmoil.

Again, there are ways of correcting for this problem. The data can be chosen to include periods when prices were extremely unstable. Such data exist, but only for

[11] With 4 securities, each can be compared with the 3 others for a total of 12 combinations, although this number must be divided by two to eliminate duplicates (the correlation between x and y is the same as between y and x), so there are 6 possible comparisons. With 100 securities, each is correlated with 99 others, for a total of 9,900 correlations; after duplicates are eliminated, the total is 4,950.

[12] This phenomenon is sometimes explained in terms of prices being "non-linear." A call option on a stock, for example, has value only if the stock price is above the strike price on the exercise date. That value changes dollar for dollar with the stock price. On the other hand, the actual stock price is completely irrelevant if it is below the strike price because the option has a value of zero. The relationship of the stock price on the exercise date to the option value is very different in the two cases.

some assets. More commonly, VaR models correct for this problem by incorporating explicit pricing models. For example, a VaR might use a Black–Sholes options-pricing model to estimate the effect of a large change in prices. Or a VarR might use a model to estimate the effect of a large interest-rate change on mortgage prepayments and the price of mortgage-backed securities.

Pricing models are also used to price securities that do not have a directly observable market price. The problem of pricing such securities is especially pronounced in the over-the-counter derivatives market, where derivatives contracts are customized to the needs of the various parties and therefore cannot be readily traded. Consequently, many derivatives have to be priced according to a model.

But pricing models have one major pitfall. There is no standard method for pricing many assets, and even for simple assets, the existing models do not always agree with observed prices. For example, the standard options-pricing model, the Black–Sholes model, tends to misprice out-of-the money options. This model assumes that price changes are distributed normally, so it systematically underestimates the probability of large price changes.[13]

A study by the Bank of England indicates the extent of the pricing problem. The Bank of England surveyed 40 institutions with major trading activities in London, asking them to price a number of standard derivatives as well as some more-exotic products. The 40 firms did not agree even on the value of a completely standard foreign-exchange option, and they disagreed on the extent to which prices would change with a change in the exchange rate.[14] As might be expected, the disagreement on the value of the more sophisticated derivatives was even greater.

So even though the pricing models are used to compensate for the fact that the risk of some securities is not always revealed in the historical data, the price models themselves disagree, and this disagreement would produce differences among the VaR models that use the pricing models. The differences in pricing models would translate directly into different capital requirements for banks, even if they held identical portfolios.

The fourth simplifying assumption is that the portfolio is fixed—does not change—during any one day, an assumption that is tenuous at best for a trading book. VaRs could, in principle, be calculated for every minute of the day, but such a calculation would be difficult for a large trading operation and of uncertain usefulness. In fact, VaRs are almost always calculated on a daily basis, so they measure the risk of the portfolio only at the end of the day. They completely ignore all risk that traders take during the course of the day.

The fifth simplifying assumption is that the numbers that go into the VaR model are known with certainty. Even if the other four simplifying assumptions are innocuous, VaR models still estimate risk using estimated probabilities of price changes. And even if these probabilities are estimated with the most sophisticated techniques, they are estimates—not known with precision. Significantly, VaR models do not allow for the uncertainty in the numbers they use. Nevertheless, taking an estimate as certain generally leads to an understatement of risk.[15]

Duffee (1996) pointed out that properly accounting for VaR models' use of estimates, not known numbers, generally increases the estimated level of risk. In his study he found the VaR which neglected this fact underestimated exposures by 33 percent.[16]

Another study, one by Marshall and Siegel (1997), examined the variation in the VaR estimates from four risk-management consultants, all of whom use the RiskMetrics model. Although the consulting firms all used the same model and were given the same data for the same portfolio, the VaR estimates for the four firms ranged between \$3.8 million and \$6.1 million.[17]

[13] Kupiec and O'Brien (1995a) stress this point.

[14] Specifically, for a European-style sterling/Deutschmark straddle, 10-month forward option at-the-money, they found a 2.7 percent standard deviation in the value, a 5.3 percent standard deviation in the delta, a 3.5 percent standard deviation in the gamma, and a 0.4 percent standard deviation in the vega. The last three terms would be used by a VaR model. For additional details see Walwyn and Byres (1997).

[15] This is an example of what is sometimes called "model risk." Model risk occurs because even if one is aware of all the possible problems in formulating a model, there is seldom an obvious solution to these problems, and sometimes a solution to one problem brings with it additional difficulties. Consequently, any model is a series of compromises based on well-informed judgment and any exercise of judgment oversimplifies or ignores potentially important aspects of reality. (The text discusses the fact that model-builders often ignore the reality that the numbers going into the model are estimates which are only more or less accurate.) At best, a model is an approximation of reality. "Model risk" is the risk of using an approximation.

[16] Dufee examined credit risk, not price risk, and he looked only at parameter uncertainty. Nonetheless, his argument applies to all VaRs because all VaRs are estimated. The 33 percent number is for the credit risk on the fixed side of a five-year U.S. dollar interest-rate swap; the estimate is specific to the model and should not be applied to other models.

[17] One firm submitted six different estimates of the VaR, all of which used methods slightly different from those specified by the authors. These estimates ranged from \$3.0 million to \$3.8 million.

This problem of taking estimates for reality is likely to be even more severe if the past turns out not to be prologue to the future, for all VaR models assume that future prices will behave as past prices did. Yet there is no reasonable alternative to using historical risks to forecast future risks.

All five of the simplifications discussed above are significant simply because statisticians have not developed the tools necessary to assess whether VaR models are reasonably accurate. These models are concerned with extreme events, namely very large losses, but by definition, extreme events do not occur frequently. Because the data are so sparse, statistical techniques have difficulty determining whether forecasts of extreme events are accurate.

The point can be made with the Basle Accord's method of evaluating the accuracy of VaRs. A portfolio should experience excess losses that exceed the 1 percent value at risk, on average 2.5 days a year (1 percent x 250 trading days). However, this is an average, and the actual number will obviously be higher or lower. The current Basle rules deem a model "acceptably accurate" if losses exceed the 1 percent VaR fewer than 4 of 250 days. A true 1 percent VaR model will meet this criterion approximately 89 percent of the time. However, a 2.5 percent VaR model will meet this criterion approximately 25 percent of the time. By definition, losses exceed the 2.5 percent VaR approximately two and a half times more often than they exceed a 1 percent VaR. Thus, a 2.5 percent VaR consistently understates the risks that interest regulators, yet it can pass the Basle tests for accuracy approximately 25 percent of the time.[18]

Work has been done to develop more powerful tests of VaR accuracy. Lopez (1998) discusses the most recent efforts. Nonetheless, these tests still are not very powerful, because by definition there are few data on extreme events.

In other words, this problem of determining the VaR forecasts' accuracy is completely independent of the method used to calculate the VaR. Unlike many of the problems discussed above, which can be avoided if slightly different methods are used (usually at the cost of creating some additional complexity and perhaps new problems), this problem is innate: whatever its methodology, a VaR model forecasts an extreme event; the accuracy of these forecasts therefore cannot be assessed without data; but by definition, the data are lacking.

Banking-Book Models (Credit-Risk Models)

Models have also been developed to estimate the credit risk of loans, and some economists and bankers have proposed using such models internally to set capital requirements for banks' banking books. J. P. Morgan also developed the standard credit model, CreditMetrics, which was released in 1997. By the end of 1998, competitors had entered the field: Credit Risk+ from Credit Suisse Financial Products, CreditPortfolioView from McKinsey, and Portfolio Manager from KMV. These products differ widely in approach; most of the discussion below relates to the approach used by CreditMetrics. Again, the discussion highlights deficiencies of the simplest configuration of the model; and this model, too, can be configured to circumvent some of the problems discussed below.[19]

Loans differ substantially from securities, so banking-book VaR models differ substantially from trading-book models.[20] For loans the primary risk is credit risk, the risk that a loan will not be repaid. Lenders know that some fraction of their loans will not be repaid; these losses are sometimes referred to as expected losses. The difference between actual losses and expected losses is unexpected losses. Although unexpected losses are on average zero, they can be quite large.

Banking-book VaRs attempt to estimate upper bounds for unexpected losses and thereby upper

[18] The potential understatement of the VaR is not proportionate to the probabilities. If losses are distributed normally with a zero mean, a 2.5 percent VaR is 13 percent lower than a 1 percent VaR (2.33 standard deviations from the mean, as opposed to 2.715). However, if losses are distributed according to a Student's t distribution with four degrees of freedom (a distribution with fat tails), then the difference is approximately 26 percent.

[19] As with fixes to the trading-book models, the fixes to specific problems with the banking-book models generally add complexity and inevitably involve their own set of problems. Jones and Mingo (1998) and Caouette, Altman, and Narayan (1998) provide useful overviews of these models. Gupta, Finger, and Batia (1997) provide the definitive statement of CreditMetrics methodology. It should be noted that because of the methodological differences with market-risk VaRs, some authors, such as Caouette, Altman, and Narayan, explicitly deny that these models are Value at Risk models. However, because credit-risk and market-risk models have the same objective, others (such as Gupta, Finger, and Batia) do refer to credit-risk models as VaRs. Koyluogo and Hickman (1998) discuss the differences between the various commercially available models.

[20] In this section the terms "bank-book model" and "VaR model" are used interchangeably to refer to the banking-book subset of VaR models.

bounds for the credit risk in the banking book. There are two basic methods of estimating credit risk, and they handle this question slightly differently. One method estimates only the probability of default; VaR models of this type are sometimes called two-state models, or default-mode models. The other method estimates not only the probability of default but also the probability of a deterioration in the borrower's credit rating; these models are called multi-state, or mark-to-market, models. The mark-to-market models are similar to trading-book VaRs in that they attempt to estimate potential losses because of changes in the value of loans the bank has already made.

The basic versions of credit-risk models use bond ratings to classify loans; for instance, a loan might be compared with a BBB bond. VaR models then assume that the probability of default and the probability of a credit downgrade for the loan are the same as those for a BBB bond. Two-state models use the probability that bonds of a particular rating would default; approximately 0.18 percent of BBB bonds were actually in default a year later. Multi-state models consider, in addition, the probability that ratings of bonds will change; approximately 5.95 percent of BBB bonds were rated A one year later, and approximately 5.30 percent were rated BB one year later. More sophisticated versions might use the bank's own rating system as well as probabilities that are based on the bank's actual experience.[21]

As in trading-book models, the correlations are critical. In banking-book models, however, the relevant correlations are between defaults or between rating downgrades. In this case, however, treating each loan individually is equivalent to assuming a zero correlation because the likely loss on one loan is completely unrelated to the likely loss on another loan. On the one hand this procedure of treating each loan individually almost certainly overstates the benefits of diversification: recessions generally increase the probability of default and adverse credit changes, so loan losses are almost certainly positively correlated.[22] On the other hand, adding the VaR from one loan to the VaR on another loan ignores the benefits from diversification: all the loans in a portfolio will almost certainly not go into default simultaneously. It is possible to calculate the correlations between defaults and between credit rating changes from bond data.[23]

The basic difference between two-state and multi-state models, as mentioned above, is their method for evaluating the potential losses to the lender. Two-state models consider only the losses that result from loans going into default; the multi-state models, in addition, consider the losses in the value of the loan because of changes in the credit rating. Any loan's value depends on the likelihood of repayment, and credit ratings are based on estimates of that likelihood. Thus, once a bank has made a loan, a deterioration in the creditworthiness of the borrower causes a decrease in the value of the loan to the bank. Consequently, a loan that was comparable to a BBB bond loses value when its rating deteriorates to BB. The multi-state model considers this lost value, implicitly doing a pseudo mark-to-market procedure.

In both kinds of models, calculation of the value at risk is completely analogous to the calculation in trading-book models, except that because loans are not repriced every day, banking-book VaRs usually use a horizon of one year. The VaR is the amount of money that would be sufficient to cover most potential losses from the bank's loan portfolio.

But although the banking-book and trading-book VaRs are conceptually similar, they use very different probability models. This methodological difference means that many of the criticisms of trading-book models discussed above do not apply to banking-book models. For example, banking-book models do not use normal distributions. However, two of the criticisms apply directly. First, banking-book VaRs, like

[21] Carey (1998) analyzed a set of privately placed bonds, which he argues are very similar to large corporate loans. He found that more poorly-rated, privately-placed bonds defaulted at a lower rate than similar publicly-traded bonds. He suggests that the holders of the former group of bonds closely monitor the activity of the issuing companies because they, the holders, will bear any losses from default. Owners of public bonds have less incentive because they typically hold a small fraction of the outstanding bonds. If this suggestion is correct, then a model that used bond default rates would tend to overstate the amount of credit risk that a bank faces.

[22] The Carey (1998) study indicates that the losses from a random portfolio of non-investment-grade bonds increased significantly during recessions. Losses for the investment-grade portfolio were much less variable. Because bank loans are generally regarded as more similar to non-investment-grade bonds, the Carey study suggests banks are especially vulnerable to large losses during recessions.

[23] CreditMetrics actually uses correlations from stock market data. It considers intra-industry correlations, inter-industry correlations, and international correlations. For example, it considers the correlation between firms within the U.S. chemical industry, the correlation between firms in the U.S. chemical and U.S. insurance industries, and the correlation between firms in the U.S. chemical and the German insurance industries. The theory is that firms will default when the value of their assets is less than their debt. Movements in stock prices reflect the movement in asset prices, and asset prices of firms in the same country or industry tend to move together.

trading-book VaRs, estimate very few correlations. Second, both banking-book and trading-book models ignore uncertainty about the model. If the model is imprecisely estimated, then the model is very likely to understate the true risk.[24]

In addition, banking-book VaRs have their own special problems. In contrast to market-risk models, credit-risk models cannot use the mass of publicly available data on securities prices. Consequently, credit-risk models are difficult to implement and must make more-questionable assumptions. Most important, to use credit-risk models banks either must have extensive internal data or must make the very dubious assumption that bond-market data adequately reflect the risk characteristics of their borrowers. Modern financial economics stresses that firms that issue bonds are very different from firms that borrow from banks. Only large firms can borrow in the bond and commercial paper market, while most borrowers from banks are smaller companies. Most large firms have long histories and both national and international operations. Many other firms have fairly short histories, and most have geographically concentrated operations.

And even if a bank chooses to use bond data, it must be able to translate its own internal underwriting standards into equivalent ratings for bonds. This is no mean task.[25]

Theoretically, banks can avoid these problems by estimating their own probabilities and correlations using internal data. But such estimation would demand data on hundreds of loans, if not thousands; more important, the data would have to extend over many years. In the past 10 years there has been one recession, and in the past 20 years there have been two (or three—some economists consider the “double-dip” recession of 1980–82 as two recessions). Defaults and credit downgrades increase during recessions, so even with 20 years of data, a bank would have only two cases to estimate the likely effects of a recession. In addition, mergers complicate matters because the internal rating systems of the two merged banks are likely to be inconsistent. Probably only a handful of banks have operated a consistent internal credit-rating system for 20 years.

Recent studies by Robert Morris Associates (1997) and Treacy and Carey (1998) examined the credit-rating systems at very large banks operating in the United States. Both studies found that these banks generally have rating systems in place, but they noted a number of problems even with these banks’ systems. One of the most severe is that many large banks apparently fail to differentiate between the riskiness of different loans.[26]

What makes these observations especially important is the virtual impossibilty of validating any banking-book VaR. As discussed in connection with trading-book VaRs, any test of a VaR model necessarily involves extreme events because VaR models are supposed to estimate potential losses under extreme circumstances. Trading-book models are hard to validate even with daily data, and of course banks do not reevaluate credit ratings on anything like a daily basis. Even if banks reevaluated credit ratings every month, there would be only 12 observations a year. With monthly reevaluations, a bank would take 20 years to gather as much data about its banking-book VaR as it could gather about its trading-book VaR in one year (approximately 250 trading days). If credit ratings were reevaluated once a quarter, a bank would need 80 years. If evaluating trading-book VaRs is difficult, assessing the accuracy of banking-book VaRs is virtually impossible.

The Precommitment Approach

One current proposal to revise the Basle Accord with respect to risk-based capital requirements would permit each bank to precommit to a maximum loss, and if actual losses exceeded the maximum predicted loss, either the bank would be fined or its capital requirement would increase in subsequent years.[27]

The maximum loss could be determined by an internal risk-management model or by some other technique. Most of the work that has been done on this approach is theoretical and has generally assumed that each banker actually knows the relevant risks. One could interpret this assumption as expressing a belief that each bank has a perfect internal model of risk. Actually, however, under this approach a bank need not have any model. Rather, a bank’s capital requirement

[24] As previously noted, Duffee (1996) made this point about credit risk, not about price risk.

[25] Treacy and Carey (1998) discuss the complications at some length. The most important difference between bond ratings and bank rating systems is that bond ratings reflect long-run creditworthiness, whereas bank rating systems focus on current financial health. Probably the reason for the difference in emphasis is that bonds typically have longer maturities than bank loans.

[26] Treacy and Carey (1998), 902, observe that 36 percent of the banks use systems that “assign half or more of their loans to a single risk grade. Such systems appear to contribute little to the understanding and monitoring of risk posture.”

[27] This approach is closely identified with the work of Kupiec and O’Brien (1995b, 1997, 1998).

would depend on that bank's own assessment of risk. Probably the most rigorous method of assessing risk is to use internal models, but the bank would be able to use any method.

Required capital would equal some multiple of the maximum predicted loss. In this respect, the precommitment approach closely parallels the current (since 1996) use of trading-book VaRs. The difference lies in the fact that fines and increased capital requirements could be imposed if a bank did not meet its commitment. The central issue of the precommitment approach is the strength and usefulness of sanctions as incentives to banks not to understate their risk.[28]

The theoretical work in this area has assumed that each bank has an incentive to decrease its capital. If a bank decreases capital, its stockholders gain in two ways. First, they receive cash either through dividends or through stock repurchases. Second, if the bank can decrease capital, shareholders have less money at risk. Moreover, raising capital is costly to stockholders inasmuch as the bank itself has to pay underwriting costs (a payment that lowers its profitability), and current stockholders have to give the new stockholders a share in future profits. On the other hand, a bank with less capital is more likely to fail. If a bank fails, its stockholders lose any future profits from the bank.

Generally economists assume that despite the risk, stockholders prefer low levels of capital. This assumption implies that banks—which act in the best interest of the stockholder—prefer to understate their risk and thereby be permitted to hold low levels of capital. However, theory cannot determine whether the incentive to maintain a higher level of capital (an incentive in the form of possible sanctions) is stronger than the incentives to decrease capital.[29] If a bank has good lending prospects and is initially well-capitalized, then the possibility of sanctions can prevent it from understating its risk. The possible loss of future profits and the possible sanction together can act as deterrents.

On the other hand, if a bank has poor lending prospects and is weakly capitalized, the threat of sanctions will be ineffective for a number of reasons. First, the bank has few future profits to lose. Second, if the bank gambles and wins, the bank will not have to pay a fine. If the bank gambles and loses enough that it fails, regulators will not be able to collect the fine. Regulators can collect the fine only if the bank loses money, but not enough money that the bank is forced into insolvency. Third, if the bank accurately reports its risk, the bank may have to raise more capital, imposing a cost on current stockholders. Under these conditions, stockholders would prefer that the bank mislead regulators.[30]

This analysis of the incentive to maintain a higher level of capital assumes that the threatened sanctions are credible. However, would regulators actually be willing to levy a fine that would materially weaken a bank that had already suffered substantial losses? Would regulators actually levy a fine that would force a bank into insolvency?

Bankers cannot be certain. They do not know that regulators will actually implement the sanctions. This uncertainty in itself weakens the incentives for banks to produce accurate assessments of their own risk. A penalty that might or might not be imposed is a much weaker deterrent than a definite, unavoidable penalty. A banker who doubts the regulators' resolve might decide to understate the risk of the bank, thereby reducing its capital requirement, whereas if the regulators were completely credible the same banker would correctly report the level of risk.

The problem of uncertainty is likely to be most severe at these banks that have central roles in the financial system. A failure of one of these banks may well cause financial chaos; and thus it is considered a "systemic risk." Would regulators in fact impose penalties on one of these banks, possibly driving it into insolvency, and risk possible systemwide financial instability? Unless the clear answer is yes, these banks have weak incentives to report their risks correctly.

And even if banks believe the sanctions are certain and the banks are always forthcoming with regulators, the precommitment approach still does not guarantee the reliability of bankers' risk assessment. As observed above, the theoretical work assumes that bankers are fully informed about the bank's level of risk and that the precommitment approach simply gives them incentives to convey this information to regulators. But the discussion of VaR models above suggests that banks do not have full information about their own risks, especially the risks in their banking books.

[28] The New York Clearinghouse is experimenting with a similar precommitment approach for its own members.

[29] Most firms also face higher borrowing costs if they decrease capital, because they are more likely to fail and less likely to repay their loans. Deposit insurance significantly reduces this incentive for banks.

[30] This is a form of moral hazard, a problem in virtually all insurance plans. The most noted advocates of the precommitment approach concede this point. *See* Kupiec and O'Brien (1995b, 1997, 1998).

Although regulatory sanctions give banks an incentive to develop this information, the costs of doing so are substantial. For bankers to be willing to incur both the costs of holding capital and the costs of developing reliable information on their risks, the potential sanctions must be sufficiently onerous.

Conclusion: Should Internal Models Be Used?

The preceding discussion has pointed out serious deficiencies in the proposals that regulators use the banks' own internal risk-management models in setting capital requirements. More important, the preceding discussion has argued that there is a substantial difference between trading- and banking-book VaR models. That difference is simply the availability of data. The large volume of data on securities prices has permitted the development of reasonable trading-book models. In contrast, data on loan performance are sparse, and the preeminent banking-book VaR models use bond data as a substitute. Furthermore, because securities prices are available daily, evaluating the reliability of trading-book VaRs is feasible, though still difficult. In contrast, the available statistical tests do not permit any sort of reasonable assessment of loan banking-book VaRs in the foreseeable future.

The precommitment approach does not solve the basic problems with internal risk-management models. Regulation cannot make inaccurate models accurate. The precommitment approach does give most banks a reason to develop such models and to report the results truthfully. But for some banks, notably undercapitalized banks and banks that do not believe the threats of regulators, this incentive is not sufficient.

Despite the problems and deficiencies, internal models produce the bank's best estimate of its possible losses, and internal models can incorporate the risk-reducing effects of diversification and hedging. These models undoubtedly measure the adequacy of a bank's capital more accurately than the current Basle standards.

Yet the Basle standards have some important advantages over internal models. Because capital standards—Basle or other—have legal standing, they must be both verifiable and uniform across time and institutions. They must therefore be based on simple, comprehensive calculations.[31] The Basle standards certainly meet all these criteria better than any approach that relies on internal risk-management models.

[31] This is the argument of Estrella (1995).

References

Caouette, John B., Edward I. Altman, and Paul Narayan. 1998. *Managing Credit Risk: The Next Great Financial Challenge*. John Wiley & Sons.

Carey, Mark. 1998. Credit Risk in Private Debt Portfolios. *Journal of Finance* 52, no. 4: 1363–1387.

Danielsson, Jon, Philipp Hartmann, and Casper de Vries. 1998. The Cost of Conservatism. *Risk* 11, no. 1:101–103.

Danielsson, Jon, and Casper G. de Vries. 1997. Extreme Returns, Tail Estimation, and Value-at-Risk. Working paper, London School of Economics.

Duffee, Gregory R. 1996. On Measuring the Credit Risks of Derivative Instruments. *Journal of Banking and Finance* 20, no. 5:805–833.

Duffie, Darrell, and Jun Pan. 1997. An Overview of Value at Risk. *Journal of Derivatives* 4, no. 3:7–49.

Elderfield, Matthew. 1998. Ripe for Reform. *Risk* 11 no. 4:25–29

Estrella, Arturo. 1995. A Prolegomenon to Future Capital Requirements. Federal Reserve Bank of New York *Policy Review* 1, no. 2:1–12.

Federal Reserve Bank of New York. 1998. *Financial Services at the Crossroads: Capital Regulation in the Twenty-First Century*. Proceedings of a Conference on February 26–27. Federal Reserve Bank of New York *Economic Policy Review* 4, no. 3.

Gupta, Greg M., Christopher C. Finger, and Mickey Batia. 1997. *CreditMetrics—Technical Document*. J. P. Morgan.

Hull, John, and Alan White. 1998. Value at Risk When Daily Changes in Market Variables Are Not Normally Distributed. *Journal of Derivatives* 5, no. 3:9–19.

Jackson, Patricia, David J. Maude, and William Perraudin. 1997. Bank Capital and Value at Risk. *Journal of Derivatives* 4, no. 3:73–89.

Jones, David, and John Mingo. 1998. Industry Practices in Credit Risk Modeling and the Internal Capital Allocations: Implications for a Models-Based Regulatory Capital Standard. Federal Reserve Bank of New York *Policy Review* 4, no. 3:53–60.

Jorion, Philippe. 1997. *Value at Risk*. Irwin.

Koyluogo, H. Ugur, and Andrew Hickman. 1998. Reconcilable Differences. *Risk* 11, no. 10:56–62.

Kupiec, Paul H., and James M. O'Brien. 1995a. The Use of Bank Trading Risk Models for Regulatory Capital Purposes. Finance and Economics Discussion Series, no. 95-11. Federal Reserve Board.

——. 1995b. A Pre-Commitment Approach to Capital Requirements for Market Risk. Finance and Economics Discussion Series, no. 95-36. Federal Reserve Board.

——. 1997. The Pre-Commitment Approach: Using Incentives to Set Market Risk Capital Requirements. Finance and Economics Discussion Series, no. 1997-14. Federal Reserve Board.

——. 1998. Deposit Insurance, Bank Incentives and the Design of Regulatory Policy. Federal Reserve Bank of New York *Policy Review* 4, no. 3:201–211.

Longerstaey, Jacques, and Peter Zangari. 1996. *RiskMetrics—Technical Document*. J. P. Morgan/Reuters.

Lopez, Jose A. 1998. Methods for Evaluating Value-at-Risk Estimates. Federal Reserve Bank of New York *Policy Review* 4, no. 3:119–124.

Marshall, Chris, and Michael Siegel. 1997. Value at Risk: Implementing a Risk Measurement Standard. *Journal of Derivatives* 4, no. 3:91–111.

Mingo, John J. 1998. Toward an "Internal Models" Capital Standard for Large Multinational Banking Companies. *Journal of Lending and Credit Risk Management* 80, no. 10:61–66, and no. 11:49–55.

Robert Morris Associates (RMA). 1997. *Winning the Credit Cycle*. RMA.

Rubinstein, Mark. 1998. Edgeworth Binomial Trees. *Journal of Derivatives* 5, no. 3:20–27.

Stahl, Gerhard. 1997. Three Cheers. *Risk* 10, no. 5:67–69.

Treacy, William F., and Mark S. Carey. 1998. Credit Risk Rating at Large U.S. Banks. *Federal Reserve Bulletin* 84, no. 11:897–920.

Walwyn, Howard, and Wayne Byres. 1997. Price Check. *Risk* 10, no. 11:18–24.

Zaik, Edward, John Walter, and Gabriela Kelling, with Christopher James. 1996. RAROC at Bank of America: From Theory to Practice. *Journal of Applied Corporate Finance* 9, no. 2:83–93.

Zangari, Peter. 1997. Catering for an Event. *Risk* 19, no. 7:34–36.