Failures in Risk Management

Ralph C. Kimball

Economist, Federal Reserve Bank of Boston, and Associate Professor of Finance, Babson College. *The author wishes to thank Richard Kopcke and Peter Fortune for helpful comments.*

Risk management has received increasing attention in recent years, both from academics and from practitioners. The number of academic articles indexed in the *Journal of Economic Literature* containing the key word "risk" in the subject description increased from 545 per year in 1988 to 859 in 1998, and the membership of the Global Association of Risk Professionals (GARP) has grown to over 8,500 since its founding in late 1996.

The heightened interest in risk management is the result of a number of coincident secular trends. Globalization of trade and production have increased financial and direct investment in volatile emerging markets. In addition, in both developed and emerging economies, capital markets have become more important as a means of allocating resources. As a result, both banks and nonfinancial firms find that the number, type, and extent of their exposures have increased significantly. Finally, a spate of volatile financial innovations are simultaneously a source of risk and a means to mitigate it.

But risk management has also attracted attention as a result of the repeated and well-publicized failures associated with its implementation. Despite the increased academic and professional attention paid to risk management, frequent instances still occur when sophisticated investors or firms experience sudden, unexpected, and devastating losses. In cases such as Barings, Metallgesellschaft, and Long-Term Capital Management, losses have been in the billions of dollars. In such cases the sophistication of the risk management processes in place was clearly inadequate for the level and type of risks assumed.

What is the source of these failures in risk management? Is it just extreme bad luck, similar to being struck by lightning while out jogging? If so, it is hard to conclude that the victims were negligent; nor do such freakish outcomes say much about the desirability of either risk management or regular exercise. On the other hand, if such failures occur because of flaws in conceptual approaches or in the way these approaches are implemented, then one can expect such failures to be repeated.

This article discusses failures in risk management, why they occur, and what can be done to reduce their occurrence. Part I discusses the nature of risk and the objectives of risk management. Part II argues that intuitively attractive conceptual simplifications often create significant errors in risk measurement. Part III then describes failures in risk mitigation. Part IV discusses implications, both for managers and for regulators.

I. Risk and Risk Management

To an economist, risk is defined as the existence of uncertainty about future outcomes. Risk is a key factor in economic life because people and firms make irrevocable investments in research and product development, plant and equipment, inventory, and human capital, without knowing whether the future cash flows from these investments will be sufficient to compensate both debt and equity holders. If such real investments do not generate their required returns, then the financial claims on these returns will decline in value.

A key conceptual advance in risk analysis occurred when analysts began to describe the risk of an investment as being equivalent to the distribution of potential outcomes, where the distribution consists of all possible outcomes weighted by their relative probability of occurrence. The more extreme the distribution of outcomes, the riskier the project. Two projects could have the same expected return (the weighted average of all possible outcomes) but differ in their risk, if one project had a broader range of outcomes or a higher probability of extreme outcomes than the other.

A second key insight was that while individual outcomes were not predictable, their distribution often was. That is, distributions often could be described by mathematical models that depended on a few key parameters, such as the mean and the standard deviation for the well-known normal distribution. If the appropriate type of distribution could be established, then an analyst could use relatively sparse historical data to forecast the key parameters and thus the future distribution of returns. This ability to predict the distribution of returns was critical both to the quantification of risk premiums and to the determination of the capital structure of the firm.

The relationship between capital structure, defined as the relative proportions of debt and equity in the firm's balance sheet, and the distribution of returns is shown in Figure 1. The horizontal axis is denominated in terms of percent of total assets. The distance OA represents the ratio of debt to assets, while the distance AB represents the ratio of equity to assets. Their sum, represented by the distance OB, equals 100 percent of assets. The distance BC represents the expected pretax return on assets, approximated by the historical mean of pretax return on assets (ROA). The curve RR′ represents the potential variation in pretax ROA around the expected value.

Book insolvency occurs when operating losses exceed the firm's equity capital; it is represented by the shaded area under the curve to the left of A on the horizontal axis. This shaded area represents the risk borne by the debt holders and is equivalent to the probability of the firm's insolvency. The area under the curve to the right of A represents the risk borne by the equity holders. Clearly, the division of risk between debt and equity holders will be determined by the position of point A relative to the distribution of outcomes. The risk borne by debt holders can be decreased if the proportion of equity capital is increased, shifting point A to the left, or increased if the proportion of equity is decreased, shifting point A to the right. Similarly, an increase in risk, signified by a flatter distribution, will require a higher proportion of equity to generate the same probability of insolvency.

Figure 1 oversimplifies matters, since the distribution of outcomes actually extends to the left past that shown by the curve RR′. That is, the possibility of catastrophic but extremely low-probability outcomes always exists. Firms can never hold enough equity to

The possibility of catastrophic but extremely low-probability outcomes always exists. Firms can never hold enough equity to guarantee their solvency, but they can estimate the amount needed to reduce the probability of insolvency to a socially acceptable small number.

guarantee their solvency, but the expected distribution of outcomes can be used to estimate the amount of equity capital needed to reduce the probability of insolvency to a socially acceptable small number.[1] This number represents a compromise among different providers of capital (debt, equity, and trade credit and, for some institutions, regulators) who have different tolerances for risk.

Figure 1

Capital Structure and Bankruptcy Risk

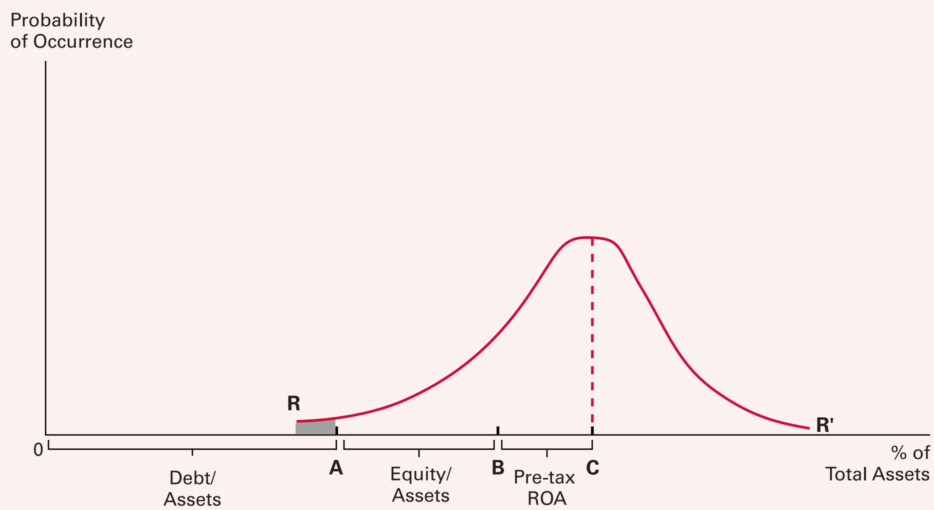

In addition to varying the proportions of equity and debt in their capital structure, firms can also affect their probability of insolvency by mitigating the risk exposures they face. That is, by purchasing insurance, hedging, screening customers, closely supervising employees and monitoring suppliers, or through diversification, firms can reduce the dispersion of potential outcomes. A popular misconception is that the objective of risk management is to eliminate risk. In fact, firms appear to pick and choose among the types and degrees of exposures, assuming those that they believe they have a competitive advantage in managing and laying others off into the capital markets, or accepting small or moderate exposures while insuring against catastrophic ones (Stulz 1996). Thus, a commercial bank may accept credit risk but avoid interest rate risk, while an investment bank does the opposite.

Other aspects of the firm's operations, such as the convexity of its tax schedules, the ability of owner/managers to diversify their portfolios, and the proportion of intangible assets, can also affect the extent to which managers attempt to mitigate risks (Tufano 1996). Increasingly, both economists and strategic planners are coming to view risk management as being related to the issue of the boundaries of the firm.[2] That is, each firm must decide which activities are to be carried out within the firm, and which ones are to be externalized to suppliers or customers. In this framework, the decision to assume or mitigate particular risks is analogous to the decision to integrate backwards or to outsource a particular function. To the extent a firm has a competitive advantage in assuming and managing a particular exposure, it can earn higher returns or, alternatively, grow faster than its competitors. Thus, risk management, like technology, distribution, or scale, is a source of competitive advantage.

If a firm fails because of "unexpected" losses, the failure is due to one of three causes. The firm may have accurately estimated the loss distribution (its exposure), but has insufficient capital to absorb the draw from the distribution, usually because the losses are catastrophic and exceed the "socially acceptable" hurdle for insolvency.[3] This can be termed the "hundred-year-storm" explanation and it has been advanced by the management of Long-Term Capital Management to explain the near-insolvency of that entity in the summer of 1998.[4]

While "hundred-year-storms" do occur, firms fail

[1] Even a firm with 100 percent equity can fail if it incurs liabilities that exceed its assets. This sometimes occurs if products, such as asbestos, prove harmful to the customer or the environment and create liabilities for damages.

[2] For a discussion of the boundaries of the firm, see Besanko, Dranove, and Shanley (1996), Chapters 2 & 5.

[3] Insolvency can also occur if owner/managers deliberately, and without disclosing it, choose a probability of insolvency that is higher than the socially acceptable small number. This situation can arise if the owners are attempting to maximize the value of the implicit call option on the assets of the firm represented by their equity.

[4] See "Investors May See 'LTCM, the Sequel,'" *Wall Street Journal*, May 20, 1999, p. C1. For an alternative explanation of LTCM's losses, see Section II, below.

far more frequently because they have estimated the distribution of outcomes incorrectly. Such failures are due to errors in risk measurement. This is sometimes called model error, but it can also represent a form of management myopia when managers fail to recognize that an exposure exists. The latter might be termed risk ignorance, while model error usually results because some parameter of the loss distribution or the covariance among different risks is misestimated. Both forms of measurement error are discussed in Part II. Finally, the firm may accurately measure its exposure, but then be ineffective in its efforts to mitigate the risks it assumes. Such errors in risk mitigation are discussed in Part III.

While "hundred-year storms" do occur, firms fail far more frequently because they have estimated the distribution of outcomes incorrectly, through model error or management myopia, or risk ignorance.

II. Issues in Risk Measurement

Are Returns Normally Distributed?

Using the distribution of potential outcomes to measure risk is a great conceptual advance, but it is a difficult one to implement. Estimation of a potential return distribution is usually based on historical data, but the availability of such data is often limited, and even when available, older data may have little forecasting value because of institutional or structural changes in the environment. In particular, estimation of the tails of the distribution, the area of special interest for risk managers, is difficult, since by definition the number of observations in the tails is limited.

Given these limitations on the availability of data, it is not surprising that the normal distribution became the paradigm for risk management. The normal distribution has one very strong advantage: The probability of any given outcome can be estimated, given the mean and standard deviation of the underlying distribution. Because good estimates of the mean and standard deviation can be obtained from relatively few observations, the use of the normal distribution allows risk managers to extrapolate the probability of extreme outcomes from relatively few data points. Moreover, while it has long been recognized that returns on most assets are not exactly normally distributed, the discrepancy is often so small that many analysts find it convenient to ignore the differences.

Unfortunately, while the assumption of normality in returns is beguiling, it is not justified by empirical research, and basing risk management on such assumptions can lead to serious risk-measurement errors. Many studies have concluded that most asset returns are not normally distributed but instead are fat-tailed and skewed to the left. (See Fama 1965, Duffie and Pan 1997.) The use of the normal distribution to estimate frequency of outcomes in such circumstances will result in estimates of the frequency of major losses that are too low.

This is illustrated in Table 1, which shows summary statistics for monthly total returns for both large stocks and long-term government bonds over the period from 1926 through 1997. The first two columns show the mean and the standard deviation, based on the monthly data. If returns on both stocks and bonds are normally distributed, then we can use these summary statistics to estimate the frequency of extreme losses. For example, because the mean and standard deviation of monthly large stock returns are 1.03 percent and 5.70 percent, respectively, and because 98 percent of the area under the normal distribution is encompassed by 2.33 standard deviations around the mean, then in only one month out of a hundred should the percentage loss on large stocks exceed 12.25 percent (1.03 percent $-$ 2.33 $\times$ 5.70 percent). Column 3 shows the expected frequency for losses at the 1 percent level under the assumption of normality, while column 4 shows the actual frequency of losses for the same range. The actual frequency of extreme losses is almost twice that expected under the assumption of normality, for both large stocks and long-term government bonds. In effect, the assumption of normality will result in a serious underestimate of the frequency of major losses. "Hundred-year storms" will occur much more frequently than once every hundred years.

If the assumption of normality results in downward-biased estimates of risk, then alternative models are needed if actual distributions of returns are to be estimated. Is there an alternative statistical model that can be used to estimate the actual distribution of

Table 1
Extreme Losses on Large Stocks and Long-Term U.S. Treasuries, 1926 to 1997

	Mean Monthly Return (%)	Standard Deviation of Mean Monthly Return (%)	Number of Expected Extreme Losses if Distribution Was Normal	Actual Number of Extreme Losses
Large Stocks	1.0291	5.7036	8.64	16
Long-Term U.S. Treasuries	.4469	2.2194	8.64	15

Source: Ibbotson Associates, *Stocks, Bonds, Bills and Inflation: 1998 Yearbook.*

returns on an ex ante basis? To date, our ability to address this issue is limited. One possibility is to use non-parametric models that do not depend upon the assumption of a particular type of distribution, but usually such non-parametric models require more data than are available. Alternatively, one could assume other distributions or return-generating functions that produce fat-tailed and skewed distributions. For example, one might assume some form of the stable distribution, which is bell-shaped, symmetrical, and fat-tailed. Or, one might assume a jump diffusion model, in which the returns usually behave as if drawn from a normal distribution but are periodically "jumped" up or down by adding an independent, normally distributed shock. Or one might assume that volatility is not constant but itself changes randomly over time. (See Simons 1997 and Fortune 1999.) But while these alternative statistical models might fit the data better in a conceptual sense, their use often requires the estimation of some specific parameter for which we have little confidence.

Serial Correlation

A second attractive simplification when analyzing risk is to assume serial independence, that is, that outcomes are not correlated over time, so that the outcome next period does not depend on the outcome this period. The assumption of serial independence has two major implications. If outcomes are serially independent, then the standard deviation of returns increases with the square root of time. That is, daily data can be used to estimate weekly, monthly, or annual volatility by multiplying the standard deviation of the daily data by the square root of the number of trading days in the longer period. This is very helpful, since it permits the use of relatively few observations to extrapolate ex ante the distribution of outcomes over a much longer period.

Second, the assumption of serial independence does not mean that runs of bad luck will not occur, but it does mean that such runs will be normally distributed. Thus the chances are only one in four that a firm will experience adverse outcomes in two consecutive periods, and only one in eight that a firm will experience adverse outcomes in three consecutive periods. The practical implication of this is that a firm can neglect the probability of adverse runs in computing its required equity capital.

Like the assumption of normality in returns, the assumption of serial independence in returns is not justified by the empirical evidence (Campbell, Lo, and McKinlay 1997). Instead, returns tend to be positively serially correlated, so that an adverse outcome in this period is likely to be followed by an adverse outcome next period. Table 2 shows estimates of serial correlation based on monthly return data for both large stocks and long-term U.S. Treasuries over the period from 1926 through 1997. Both series show a positive serial correlation of around 0.09.

If returns are assumed to be serially independent but actually are positively correlated, then estimates of long-period standard deviations extrapolated from short-period returns will be too low. Once again, the frequency of seriously adverse outcomes will be underestimated, and a firm will require additional equity capital if it is to achieve the same probability of insolvency.

Table 2
Serial Correlation for Large Stocks and Long-Term U.S. Treasuries, 1926 to 1997

	Coefficient of Serial Correlation
Large Stocks	.0864
Long-Term U.S. Treasuries	.0894

Source: Ibbotson Associates, *Stocks, Bonds, Bills and Inflation: 1998 Yearbook.*

Correlation among Outcomes

In measuring the risk exposure of a firm or financial institution, the estimation of the correlation among asset returns is as important or more important than the estimation of the distribution of the individual asset returns. This is so because the risk of a portfolio of assets depends not only on the stand-alone risks (standard deviations) of the individual assets but also on the correlation (covariance) among them. Unless the different assets are perfectly positively correlated, then the assets will act as partial natural hedges for each other, so that diversification of the portfolio among different asset types provides an inexpensive and readily available means to mitigate risk.

Indeed, the existing paradigm for the quantification of risk premiums, the capital asset pricing model (CAPM), distinguishes between two types of risks: independent and systemic. Independent risks are measured by the standard deviation of returns, while systemic risks are measured by the extent of correlation in returns. Independent risk can be diversified away through a broadly diversified portfolio, while systemic risk cannot. Indeed, because independent risk can be so easily and cheaply hedged through diversification, the CAPM implies that risk premiums depend only on systemic risks, the risks that cannot be diversified away.

The probability of a significantly adverse return on one security being accompanied by a significantly adverse return on the other is substantially greater than one would expect, based on historical data.

Calculations of correlation among asset returns usually use historical data, with the implicit assumption that correlation is stable across periods. But several recent empirical studies contradict this assumption: Correlation among asset returns is much higher during periods of extreme returns than during "normal" periods. (See Roll 1988; Longin and Solnik 1995; and Booth and Broussard 1998.) That is, the probability of a significantly adverse return on one security being accompanied by a significantly adverse return on the other is substantially greater than one would expect, based upon the coefficient of correlation calculated using data from all periods. For example, in a study of common stocks and REITs, Booth and Broussard found a positive correlation for negative returns of 0.834. However, if the 1 percent of the most negative returns are eliminated from both series, then the correlation coefficient for the remaining data pairs declines to 0.214.[5]

Like positive serial correlation and non-normality in returns, the tendency for asset returns to be more positively correlated during financial crises tends to reduce estimates of the frequency of seriously adverse outcomes. In particular, diversification may be an effective way to mitigate portfolio risk in "normal" periods, but may be ineffective in extreme periods.

Risk Ignorance

Assumptions of normality, serial independence, or non-varying return correlation are all examples of "model error." Model error occurs when the potential exposure is recognized but misestimated because some parameter of the distribution of outcomes is misestimated, or because the correlation between different risks is misestimated. Model error often results either because managers make inappropriate ex ante assumptions concerning the shape of the distribution, or because the conceptual models fail to capture some important aspect of reality.

But a second and extreme form of risk measurement error occurs when the firm fails to recognize it has any exposure whatsoever. Such a case might be termed "risk ignorance." Risk ignorance is often associated with the introduction of new products or processes where the technical aspects or environmental effects of such products or processes are not fully understood, or with changes in the legal environment through legislation or litigation. A classic example of risk ignorance was provided by the losses experienced by Chase Manhattan Bank in connection with Drysdale Securities, a government securities dealer, in the early 1980s. Chase's loans to Drysdale were executed via repurchase agreements, then a relatively new form of borrowing, using the latter's government securities as collateral. Because it was holding as collateral highly liquid and default-free securities sufficient to cover the principal of the loan, Chase believed its

[5] See Booth and Broussard (1998), Table 3.

loans to be fully secured. But Chase's exposure included not just the principal but also the accrued interest, and when Drysdale failed, Chase was left with an unsecured claim of $50 million.[6]

III. Risk Mitigation

A firm does not necessarily have to accept a particular distribution of outcomes, but often can modify the probability of adverse outcomes through its own efforts. These efforts to alter the distribution of outcomes can be termed "risk mitigation."[7] To the extent that a firm successfully mitigates its risks, then its distribution of outcomes will be less extreme, and it will require less equity capital than if it had undertaken no risk mitigation.

Risk mitigation can take a number of different forms. Perhaps the two most obvious are the purchase of insurance, where the firm pays an unrelated third party to assume the exposure, and hedging, where the firm takes an offsetting position in a security, commodity, or currency that is closely correlated with the exposure it wishes to mitigate. But firms also employ a number of other measures to mitigate exposures, including market research, geographic and product line diversification, screening and monitoring of customers, outsourcing, imposing risk premiums in pricing products, carrying inventories or slack in productive capacity, and imposing defined procedures designed to minimize operational risks.[8]

To the extent a firm can mitigate its risks economically, it can earn higher returns for shareholders. A firm with less variable returns will require a lower proportion of equity in its capital structure to obtain the same probability of insolvency. On the other hand, risk mitigation usually involves real costs. If these real costs of risk mitigation are less than the cost of the equity that would otherwise be required, the firm will earn higher returns than competitors who do not mitigate risks. Thus, risk mitigation, like scale economies or superior technology or distribution, can be viewed as a potential source of competitive advantage. Of course the benefits from risk mitigation will accrue only if such efforts are effective. If they are not, then the firm's actual distribution of returns will be more volatile than envisaged and the probability of insolvency for any given level of equity will be higher than planned. In the past few years, the increased efforts by firms to mitigate risks have led to more instances where poorly implemented risk mitigation efforts have led to major losses. (See Figlewski 1994.)

A prominent example of poorly implemented risk mitigation is the losses experienced by the German firm Metallgesellschaft. In the Metallgesellschaft case, a U.S. subsidiary contracted to sell 154 million barrels of oil on fixed-price contracts extending over a 10-year period, and then sought to hedge the resulting substantial exposure to oil price increases by taking a long position in oil futures and commodity swaps. When the spot price of oil subsequently fell, the long-term fixed-price contracts to deliver oil increased in value while Metallgesellschaft's futures and swap positions declined in value. But because its futures positions were marked to market while its delivery contracts were not, Metallgesellschaft's financial statements showed large losses and the firm experienced large cash outflows.[9] Evidently convinced they could not sustain further cash outflows, Metallgesellshaft's management chose to liquidate the hedge, leaving themselves exposed when oil prices subsequently rose again.[10] (See Mello and Parsons 1995; Culp and Miller 1995.)

Risk mitigation efforts may be ineffective for a number of reasons. Perhaps the best known is agency risk, the risk that a manager or employee, inadvertently or purposefully, will fail to follow the policies or procedures designed to mitigate risk. For example, a rogue trader whose compensation or tenure is dependent upon his trading results may fail to abide by position limits or hide cumulative losses, or mainte-

[6] In many cases, risk ignorance creates model error. For example, many of the more notorious failures in risk management involve a firm using an innovative financial instrument or hedging technique without fully appreciating the range of potential outcomes. See Figlewski (1994).

[7] Risk mitigation implies that the efforts of the firm are aimed at reducing risk, but firms can also make an explicit decision to increase risk through activities such as increasing leverage or reducing diversification.

[8] See, for example, Biddle (1999).

[9] In the futures markets capital gains and losses are computed and paid at the end of each day, giving rise to continual cash flows. This process is called "marking to market."

[10] Another factor contributing to its losses was the structure Metallgesellschaft chose for the hedge. Metallgesellschaft chose to construct a rolling hedge consisting of short-term futures contracts. In order to hedge the long-term delivery contracts, these short-term contracts needed to be rolled over at delivery. Usually this would not have been a problem, since in petroleum markets spot prices are usually higher than near-dated futures and Metallgesellschaft would actually have made money each time it rolled the hedge. Unfortunately, soon after Metallgesellschaft established its position, the usual relationship between spot and near-dated futures prices reversed itself and spot fell below the futures price. Thus, each time Metallgesellschaft rolled the hedge it experienced losses.

nance personnel may overlook an incipient equipment failure. Agency risk has been responsible for a number of notorious episodes, including the bankruptcies of Orange County and Barings, and the large losses of Sumitomo.

Orange County operated a pooled short-term investment fund for local governmental units within its jurisdiction. While the espoused objective of the pooled fund was to invest in risk-free or very low-risk securities, the Orange County official in charge of the fund sought increased returns by borrowing to invest in repurchase agreements and mortgage-backed securities. When interest rates rose abruptly, the now leveraged fund experienced massive losses. (See Jorion 1995.)

Both Barings, a British merchant bank, and Sumitomo, a Japanese trading company, were victims of rogue traders who successfully concealed huge cumulative trading losses. In both cases the traders were able to conceal such huge cumulative losses because they were responsible for reporting and settling their own trades. (See Fay 1997; *The Economist* 1996.)

Risk Migration

But risk mitigation efforts can also fail for more subtle and indirect reasons. The first is the tendency for risk to shift or change form. While an individual firm may mitigate its risks by purchasing insurance or hedging, these actions do not reduce systemic risk in the economy, but only transfer it elsewhere.[11] Moreover, in many cases hedging or purchasing insurance does not really transfer risk, but merely transforms the nature of the exposure.

For example, consider an interest rate swap in which Party A pays a fixed rate and receives a variable rate on some notional principal. The swap is a zero-sum transaction, so Counterparty B receives the fixed rate and pays a variable one. Party A's motivation for entering into the swap is to use it to hedge an imbalance between its fixed-rate assets and variable-rate liabilities. By entering into the swap, A protects itself against a potential increase in short-term rates, since any increase in the cost of its liabilities will be offset by increased revenue from the variable leg of the swap. A may believe it has obtained protection against interest rate risk, but such protection will only exist so long as B continues to perform on its side of the swap. That is, the swap does not really extinguish the exposure of A, but only transforms it from an exposure to interest rates to a credit risk exposure to B. Indeed, the greater the success of the swap in mitigating interest rate risk, the greater the credit exposure of A to B.[12]

The tendency for other types of risk to migrate and reappear as counterparty or credit risk is illustrated by the experience of U.S. banks during the Russian financial crisis of August 1998. Many U.S. banks had hedged their exposure to fluctuations in the value of the ruble by executing foreign exchange swaps with Russian banks that required the latter to exchange dollars for rubles at a fixed rate. When the ruble was effectively devalued in August 1998, the U.S. banks nevertheless experienced substantial losses when a substantial number of Russian counterparties defaulted. (See Bomfim and Nelson 1999.)

A second factor that often causes risk mitigation efforts to fail is the tendency for risk management processes to fail incrementally over a long period of time.

Any loss, whether resulting from operations or from market fluctuations, must ultimately be absorbed by someone's net worth. If Party A suffers a loss due to a market fluctuation, A's net worth will decline. If the loss exceeds A's net worth, then the excess must be absorbed by the net worth of A's creditors. If that net worth is insufficient, any excess must be absorbed by the net worth of A's creditors' creditors. In short, like

[11] Independent risk is firm specific and can be reduced through diversification. Systemic risks cannot be affected through diversification. Because hedging is a form of diversification (the hedger creates a portfolio of offsetting exposures), it only reduces independent risk, although it does transfer systemic risks elsewhere in the economy.

[12] Of course A can mitigate its credit exposure to B by enlisting the assistance of swap dealer C. Both A and B contract with C. C receives fixed payments from A and pays them to B, simultaneously receiving variable payments from B and paying them to A. Thus neither A nor B has credit exposure to the other. C has no interest rate risk exposure but assumes the credit risk that either A or B will fail to perform. But while this interpolation of a guarantor mitigates A's credit exposure to B, it does so only by transforming it to a credit exposure to C.

a row of dominos toppling in sequence, losses will migrate through the financial system until completely absorbed by net worth.[13] Thus, a firm's exposure is not limited to its own counterparties, but also includes the other exposures of these counterparties.[14]

A second factor that often causes risk mitigation efforts to fail is the tendency for risk management processes to fail incrementally over a long period of time (Grabowski and Roberts 1997). Case studies of major industrial accidents have noted that such accidents are often preceded by a long incubation period marked by a gradual degradation of the risk management processes. An initial failure to follow maintenance schedules or to test backup systems does not usually result in immediate harmful incidents, leading managers to conclude that the probability of failure has been overestimated, or that such processes are redundant. In effect, the organization becomes desensitized to the existence of risk. As a result, efforts to rectify the degradation of the risk mitigation process are actually less likely to occur as the degradation continues. Over time these incremental failures accumulate until some incident causes them to interact in a fashion that results in a major loss—a loss that could have been avoided, had the original mitigation processes been followed.

While the long incubation period associated with failures in risk mitigation processes has been most clearly identified with respect to industrial catastrophes, a similar process appears in financial crises.[15] In particular, financial crises are often preceded by a long period of gradually increasing exposures to a particular customer or asset type, with an accompanying loss of diversification. Because returns are attractive and losses minimal in the initial stages, managers often convince themselves that they have learned how to manage the risks involved, or that potential losses are less than previously believed, and that they are thus justified in increasing their exposures.

IV. Implications

Implications for Managers

The existence of non-normality in returns, positive serial correlation, and state-sensitive correlation in returns means that managers must view their ability to forecast the distribution of future outcomes with some skepticism. Use of simplifying assumptions such as normality is likely to result in significant underestimation of the probability of seriously adverse outcomes. Many institutions have recognized the danger of building their risk management processes upon assumptions such as normality, and have developed approaches that address model error.

Perhaps the most common way to address this issue of model error is to allow for a margin of error. For example, the normal distribution might be used to generate outcomes, but then excess capital is held to offset the model error that is believed to exist. This is the procedure used in the application of value-at-risk models to estimate risk exposures and set capital requirements for the trading books of commercial and investment banks. Currently, commercial banks are permitted to use the value-at-risk methodology to compute exposures, but then are required to hold three times as much capital as that indicated by the model. While the margin-of-error approach is directionally correct in addressing the danger of underestimating the frequency of seriously adverse events, we do not understand whether the extra capital required is sufficient or excessive.

As risk migrates through the system, it tends to emerge in its most basic form, as credit risk. This means that those institutions that specialize in managing and absorbing credit risks, commercial banks, play a special role.

One way to test the sufficiency of a margin of error is to "stress test" the firm's current exposures. Historical data on returns and their correlation during some crisis period, such as the November 1987 stock market crash or the 1998 Asian crisis, may be used to

[13] In severe crises where the losses exceed private domestic equity, the losses must be absorbed by the government or by providers of foreign capital. This often occurs in banking crises when the banking system must be re-capitalized.

[14] Because ex post risk migration is more likely to occur when losses are large, this may explain why asset return correlation is higher during extreme events.

[15] Some macroeconomists have minimized the significance of lax regulatory policies as a contributing factor to the 1998 Southeast Asia crisis because such regulations were in place over a sustained period of time and were not relaxed just prior to the crisis. But such an argument ignores the long incubation period that occurs as risk mitigation processes undergo degradation. See Whitt (1999).

'measure the potential losses in a worst-case scenario and to determine if the firm has sufficient equity to weather such an extreme event. For example, one large international re-insurer stress tests its portfolio under the assumption of simultaneous 8.5 Richter-scale earthquakes in Tokyo and Los Angeles.

Failures in risk mitigation most commonly arise due to agency risk. Such risk is best controlled through active monitoring of exposures and employees, to ensure compliance with limits and procedures. Often this involves specialized staff who in effect conduct real-time audits, both to confirm reported exposures and to rectify minor deterioration in risk mitigation processes before they can incubate into major losses.

Implications for Regulators

As noted above, risk tends to migrate in the financial system. In particular, hedging does not reduce systemic risk, but only transfers the exposure elsewhere or transforms the type of the exposure. Thus, risk migration has three important implications. First, because risk mitigation activities such as hedging do not reduce the amount of systemic risk in the system, they also do not reduce the aggregate amount of equity capital needed to absorb this risk.[16] That is, the amount of equity capital needed systemwide is independent of the amount of risk mitigation that is undertaken. Second, the greater the amount of risk mitigation undertaken through hedging or the purchase of insurance, the more likely that unforeseen losses will migrate quickly from one market to another, or from one country to another. That is, while hedging acts to reduce independent risk, it can enhance systemic risk. Finally, as risk migrates through the system, it tends to emerge in its most basic form, as credit risk. This tendency for errors in risk management to ultimately emerge as credit exposures means that those institutions that specialize in managing and absorbing credit risks, that is, commercial banks, play a special role.

Because commercial banks are in the business of accepting and managing credit risk, they act as shock absorbers, absorbing errors in risk management made elsewhere in the system. The capacity of the banks to act as buffers against errors in risk management depends on their ability to measure and mitigate their own exposures, as well as the sufficiency of their equity capital. A well-managed and well-capitalized banking system is thus requisite for avoiding systemic economic and financial crises.

[16] Once again, hedging can reduce independent risk, but not systemic risk. To the extent that firms in the aggregate hold equity capital against independent risk, hedging can reduce the amount of equity capital held, but not that needed against systemic risk.

References

Besanko, David, David Dranove, and Mark Shanley. 1996. *Economics of Strategy.* New York: John Wiley & Sons.

Biddle, Frederic M. 1999. "Can Kaufman & Broad Beat the Home Building Cycle?" *The Wall Street Journal*, October 27, p. B4.

Bomfim, Antulio N. and William R. Nelson. 1999. "Profits and Balance Sheet Developments at U.S. Commercial Banks in 1998." *Federal Reserve Bulletin,* June, pp. 369–95.

Booth, G. Geoffrey and John Paul Broussard. 1998. "REIT Returns, Probability of Large Losses, and Asset Allocation." Eli Broad Graduate School of Management, Michigan State University, Processed.

Campbell, John Y., Andrew W. Lo, and A. Craig McKinlay. 1997. *The Econometrics of Financial Markets*. Chapter 2. Princeton, NJ: Princeton University Press.

Culp, Christopher and Merton H. Miller. 1995. "Metallgesellschaft and the Economics of Synthetic Storage." *Journal of Applied Corporate Finance*, Winter, pp. 62–76.

Duffie, Darrel and Jun Pan. 1997. "An Overview of Value at Risk." *The Journal of Derivatives*, vol. 4, no. 3 (Spring), pp. 7–49.

The Economist. 1996. "Coming A Cropper in Copper," June 22, pp. 69–70.

Fama, Eugene F. 1965. "The Behavior of Stock Market Prices." *Journal of Business*, vol. 38 (January), pp. 34–105.

Fay, Stephen. 1997. *The Collapse of Barings*. New York: W.W. Norton & Co.

Figlewski, Stephen. 1994. "How to Lose Money in Derivatives." *The Journal of Derivatives*, vol. 2, no. 2 (Winter), pp. 75–82.

Fortune, Peter. 1999. "Are Stock Returns Different Over Weekends? A Jump Diffusion Analysis of the 'Weekend Effect.' " *New England Economic Review*, September/October, pp. 3–19.

Grabowski, Martha and Karlene Roberts. 1997. "Risk Mitigation in Large-Scale Systems: Lessons from High-Reliability Organizations." *California Management Review*, Summer, pp. 152–59.

Jorion, Phillipe. 1995. *Big Bets Gone Bad: Derivatives and Bankruptcy in Orange County.* San Diego, CA: Academic Press.

Longin, F. and B. Solnik. 1995. "Is the Correlation in International Equity Returns Constant? 1960–1990." *Journal of International Money and Finance*, vol. 14, pp. 3–26.

Mello, A. and J. E. Parsons. 1995. "Maturity Structure of a Hedge Matters: Lessons from the Metallgesellschaft Debacle." *Journal of Applied Corporate Finance*, Spring, pp. 106–20.

Roll, Richard R. 1988. "The International Stock Market Crash of 1987." *Financial Analysts Journal*, vol. 44, no. 5, pp. 19–35.

Simons, Katerina. 1997. "Model Error." *New England Economic Review*, November/December, pp. 17–25.

Stulz, René. 1996. "Rethinking Risk Management." *Journal of Applied Corporate Finance*, Fall, pp. 8–24.

Tufano, Peter. 1996. "Who Manages Risk? An Empirical Examination of the Risk Management Practices of the Gold Mining Industry." *Journal of Finance*, September, pp. 1097–1137.

Whitt, Joseph. 1999. "The Role of Shocks in the Asian Financial Crisis." *Federal Reserve Bank of Atlanta Review*, Second Quarter, pp. 18–31.