

*Journal of***APPLIED CORPORATE FINANCE**

A MORGAN STANLEY PUBLICATION

In This Issue: Honoring Stewart Myers**The Contributions of Stewart Myers to the Theory and Practice of Corporate Finance**8 *Franklin Allen, University of Pennsylvania, Sudipto Bhattacharya, London School of Economics, Raghuram Rajan, University of Chicago, and Antoinette Schoar, MIT***MIT Roundtable on Corporate Risk Management**20 *Panelists: Judy Lewent, Merck; Donald Lessard and Andrew Lo, MIT; and Lakshmi Shyam-Sunder, International Finance Corporation. Moderated by Robert Merton, Harvard Business School.***Risk Management Failures: What Are They and When Do They Happen?**39 *René Stulz, Ohio State University***Brealey, Myers, and Allen on Valuation, Capital Structure, and Agency Issues**49 *Richard A. Brealey, London Business School, Stewart C. Myers, MIT, and Franklin Allen, University of Pennsylvania***Brealey, Myers, and Allen on Real Options**58 *Richard A. Brealey, London Business School, Stewart C. Myers, MIT, and Franklin Allen, University of Pennsylvania***Equity Issues and the Disappearing Rights Offer Phenomenon**72 *B. Espen Eckbo, Dartmouth College***Can Companies Use Hedging Programs to Profit from the Market? Evidence from Gold Producers**86 *Tim R. Adam, Humboldt University, and Chitru S. Fernando, University of Oklahoma***Corporate Leverage and Specialized Investments by Customers and Suppliers**98 *Jayant R. Kale, Georgia State University, and Husayn Shahrur, Bentley College***Estimating Risk-Adjusted Costs of Financial Distress**105 *Heitor Almeida, University of Illinois at Urbana-Champaign, and Thomas Philippon, New York University***Morgan Stanley**

Risk Management Failures: What Are They and When Do They Happen?

by René M. Stulz, Ohio State University*

In media accounts and popular commentaries on the current financial crisis, a constant refrain is that the risk management function in many of the world's largest financial institutions has failed to carry out its responsibilities. To cite just one example, an article in the *Financial Times* declares "it is obvious that there has been a massive failure of risk management across most of Wall Street."¹

In this article, I want to challenge, or at least qualify, this assertion by examining what it means for risk management to fail. My main aim in these pages is to show that the fact that an institution suffers an extremely large loss does not necessarily imply that risk management failed, or that the institution made a mistake. In making my case, I will not spend much time on the current crisis, but instead begin with the collapse of Long-Term Capital Management in 1998. Having accumulated a decade of hindsight on this case, we now have enough information to at least ask the right questions to determine whether risk management failed and to ensure that, if risk management is blamed, it is blamed for the right reasons.

Getting the diagnosis right is important because the changes in risk management that take place in response to the crisis could be the wrong ones. Most troubling, top executives and investors could continue to expect more from risk management than it can deliver. With this goal in mind, I show when bad outcomes can be blamed on risk management and when they cannot. And in so doing, I offer what amounts to a taxonomy of risk management failures.

I also address the question of whether and how risk management failures can be used to improve the practice of risk management. I conclude that the probabilities of large losses are measured very imprecisely and that, as a consequence, companies should rely less on estimates of such probabilities and pay more attention to the implications of large losses for their survival. Among other suggestions, I propose that greater use of scenario planning could allow institutions to do a better

job of anticipating the likely consequences of low-probability outcomes and developing effective responses to them.

Was the Collapse of LTCM a Risk Management Failure?

We now know a good deal about the rise and fall of Long-Term Capital Management, or LTCM.² In 1994, ex-Salomon Brothers traders and two future Nobel Prize winners started a hedge fund called the Long-Term Capital Fund. LTCM was the company that managed the fund. The fund performed superbly for most of its life, with investors earning 20% for ten months in 1994, 43% in 1995, 41% in 1996, and 17% in 1997. But 1998 turned out very differently. In August and September of 1998, following the default of Russia on its ruble-denominated debt, world capital markets were in crisis and LTCM lost most of its capital. Before its collapse, LTCM had capital close to \$5 billion, assets in excess of \$100 billion, and derivatives with a notional amount in excess of \$1 trillion. By mid-September, LTCM's capital had fallen by more than \$3.5 billion and the Federal Reserve Bank of New York coordinated a rescue by private financial institutions that injected \$3.65 billion into the fund.

Does a loss of more than 70% of capital, and a rescue by banks involving an injection of \$3.65 billion of new capital, necessarily represent a failure of risk management? It turns out that this is not an easy question to answer. To define a risk management failure, one first must have a clear understanding of the role and the limitations of risk management.

In a typical company, the role of risk management is to identify and evaluate the risks faced by the firm, to communicate these risks to senior management (and possibly the board of directors), and to monitor and manage those risks in a way that ensures the firm bears only the risks its management and board want exposure to. To guide them in monitoring and managing risk, most companies specify one (or more) measures of overall risk, perhaps along with other metrics or indicators. When a risk measure exceeds the company's stated tolerance for risk, risk is reduced. But when the risk measure

* I am grateful for assistance from Jérôme Taillard and Mike Anderson, and for comments by Rich Apostolik, Don Chew, Cliff Smith, Adrian Tschoegl, and Peter Tufano.

1. "Wall Street Dispatch: Imagination and Common Sense Brew a Safer Culture," by David Wighton, Nov. 26, 2007, FT.com.

2. The best public source of data on LTCM is the collection of four Harvard Business School case studies by André Perold published in 1999. See "Long-Term Capital Management (A) – (D)," available from Harvard Business School Publishing. Many books have been written on LTCM. Some of the numbers used in this article come from Roger Lowenstein, *When Genius Failed: The Rise and Fall of Long-Term Capital Management* (New York: Random House, 2000).

falls below the firm's targeted risk position, the firm will likely choose to increase risk.

For financial institutions, one commonly used risk measure is Value-at-Risk, or VaR. As a measure of downside risk, VaR is the maximum expected loss at a given confidence level over a given period of time. For example, in the case of a bank whose senior management specifies a 99% confidence level, a one-day VaR of \$150 million would mean that the firm has a 1% chance of making a loss in excess of \$150 million—provided, of course, the VaR has been correctly estimated. Although financial institutions generally report daily VaR measures, VaRs can also be estimated for longer periods of time.

Given this definition of the role of risk management, the actual returns of LTCM tell us little about whether its risk management failed. To understand why, it is helpful to consider a very simple hypothetical example. Suppose that you stood in the shoes of the managers of LTCM in January 1998 and had the opportunity to invest in trades that, viewed as a portfolio, had a 99% chance of producing a return for the fund before fees of 25% over the coming year, and a 1% chance of a loss of 70%. The expected return on the fund would be estimated to be 24.05% $[(25\% \times .99) + (-70\% \times .01)]$. Such an expected return would, of course, have been very attractive for a hedge fund, or any investor, again provided the assigned probabilities and other underlying assumptions were reasonably consistent with economic reality.

Though hypothetical, these numbers are roughly consistent with the actual returns of LTCM before 1998 and the expectations LTCM was holding out to its investors at the time. For one thing, in its two best years, the fund earned more than 50% before fees, so that a return of 25% does not sound implausible. Second, LTCM told its investors to expect losses in excess of 20% in one year out of 50—which, again, is roughly comparable to an expected loss of 70% in one year out of 100.³

Now, given this information, let's assume that whether the fund had the high return or not depended on the flip of a "heavily loaded" coin—one whose odds of turning up heads (a 25% return) was 99%. Had the fund been given the chance to carry out this investment strategy for 100 years, it would have earned 25% in 99 of them.

In this hypothetical example, we are assuming that when the partners of LTCM made their investment decisions, those decisions were based on the best available assessment of the expected distribution of possible outcomes of the fund. To the extent this assumption is correct, the fund's risk managers would have provided a sound basis for decision-making and should have earned a gold medal for their work.

But suppose the bad outcome—the one-in-a-hundred event—materializes, and the fund makes headlines for having lost \$3.5 billion. The natural inclination, of course, is to interpret a 70% loss as clear evidence of a risk management failure. But in the example I've just recounted, the job done by the firm's risk managers could not have been improved on.

As stated earlier, then, the main responsibility of risk management is to assess the firm's risks and then communicate that assessment to top management. And the final decision to take known risks rests not with the risk manager, but with top management. This decision depends on the risk appetite or tolerance of an institution, and defining that appetite is one of the most important responsibilities of management and the board. It is at the heart of the firm's strategy and how it creates value for its shareholders.

When determining the risk appetite of an organization, the general approach is to weigh the benefits of increased risk-taking against the costs, and to aim for the point where the marginal benefits equal the marginal costs. In the case of LTCM, it could be argued that the costs to the organization from losing \$3.5 billion for its investors was limited to just those losses—that is, there were no additional costs beyond the direct monetary loss. For many companies and financial institutions, however, large losses can have significant "deadweight" costs. For example, after a large loss, a financial institution may be forced to scale back its investments and sell assets in unfavorable markets, face increased scrutiny from regulators, and lose valuable employees and customers concerned about the institution's staying power. In any institution, management and the board must take account of such deadweight costs when making decisions that create the risk of large losses.⁴

So, again, it is top management's job to determine a company's risk appetite. And taking that appetite as a given, risk managers can help assess the expected profitability of a proposed investment by evaluating how much capital is required to support it.⁵

Also worth noting, an investment that is not value-adding for a given level of risk could become so if the firm's risk appetite increases—because less capital is then required to support it. But whether taking large risks is worthwhile for an institution ultimately depends on its strategy, which in turn depends heavily on its risk appetite. And risk managers, as already suggested, do not set either; both strategy and risk appetite are the prerogatives of top management.

To illustrate this point, suppose that a company follows the common practice of setting its risk tolerance by choosing a targeted credit rating. Once the target rating is chosen, there are multiple combinations of risk and capital that can

3. See Lowenstein, p. 63.

4. Such costs are at the foundation of financial theories of how risk management creates shareholder value. See René M. Stulz, *Risk Management and Derivatives*, Thompson Publishing, 2003.

5. My article with Brian Nocco, *Enterprise Risk Management: Theory and Practice*, *Journal of Applied Corporate Finance*, Fall 2006, v18(8), 8-20, describes the key principles of enterprise risk management, issues that arise in its implementation, and the role of capital allocation.

be expected to achieve the target rating. For a given choice of leverage, the firm does not have much flexibility in choosing its risk level if it wants to achieve its target rating. If faced with promising investment opportunities that would increase risk, the firm has two main choices: it can choose to reduce leverage (and financial risk) while taking on more business risk to keep its rating; or it can maintain its leverage and accept a lower credit rating. And there is also, of course, a third choice: maintain both the credit rating and existing leverage, but at the cost of passing up the investment opportunity.

LTCM provides a good example of such tradeoffs. In the fall of 1997, the managers of LTCM concluded that they did not want to manage a business earning just 17% for its investors, which is what their investors had earned for the year. Instead, they wanted the higher returns achieved in 1995 and 1996. At the end of 1997, LTCM had capital of \$7.4 billion but decided to return 36% of the capital to its investors. With less capital, LTCM could still execute the same trades; but to fund them it had to borrow more and hence increase its leverage. By increasing its leverage, it could boost the return to its shareholders if things went well—but only by raising the expected losses if things went poorly.

Was increasing leverage a bad risk management decision? In my example, I assume that the partners of LTCM knew the risks and the rewards from higher leverage. In the well-worn language of financial economics, increasing leverage appeared to be a positive NPV decision when it was made. But ex post, or after the fact, it proved to be a costly decision since it meant that when assets fell in value, the value of the fund's equity fell much faster.

There has been much discussion of top management incentives during the credit crisis, with many observers pointing to stock options as an inducement to take excessive risk. Before reaching this conclusion, however, it's useful to keep in mind that financial economists have argued for decades that management's incentives become better aligned with those of their shareholders when management has a large stake in the firm's equity. Top management owned hundreds of millions of dollars of equity in both Bear Stearns and Lehman at the peak valuation of these firms. And the partners of LTCM collectively had almost \$2 billion invested in the fund at the beginning of 1998. If such equity stakes do not provide managers with strong incentives to make the right decisions for their shareholders, what will?

In sum, effective risk management does not provide a guarantee against failure. Even in companies with the best risk management people and systems, large losses can and will occur as long as taking the risk of large losses increases expected profits sufficiently for top management to be willing to take that risk. With good risk management, such losses will be attributable to an unlucky "draw," to a one-in-a-hundred event. Ultimately, how likely such large losses are will depend on choices made by those entrusted with determining the

firm's risk appetite. The job of risk management is to ensure that top management knows and understands the probabilities associated with possible outcomes of the firm's strategy before they make decisions to commit the firm's capital.

A Taxonomy of Risk Management Failures

Having described the role of risk management, let's now consider what can go wrong.

The first step in risk management is to identify and measure risks. Risk measurement takes place in the context of risk metrics like VaR that aggregate various types of risks to help top management understand the risk position of the firm. The choice of risk metrics is the cornerstone of risk management because it determines what top management learns from risk managers about the firm's overall risk position. If the risk managers assess risk using measures that are ill-suited to a firm's strategy, risk management can fail before the computers are ever turned on.

Once a risk measure is chosen, there are two basic kinds of mistakes that can be made in measuring risk: (1) known risks can be mismeasured and (2) important risks can be ignored, either because they are undetected or wrongly viewed as immaterial.

Once the risks are identified and measured, they must be communicated to the firm's leadership. A failure in communicating risk to management is also a risk management failure.

After risks have been measured and communicated, top management decides how much and what kinds of risks to take. At this point, it is the responsibility of risk management to make sure that the firm actually takes these risks and not others. A company's risk managers must keep track of and manage the firm's risks to ensure they remain within the established guidelines, a task that could involve hedging risks and rejecting proposed trades or projects.

As this discussion suggests, then, there are five types of risk management failures:

- 1) Failure to use appropriate risk metrics.
- 2) Mismeasurement of known risks.
- 3) Failure to take known risks into account.
- 4) Failure in communicating risks to top management.
- 5) Failure in monitoring and managing risks.

I now discuss each of these failures in turn.

Risk Metrics and Risk Management Failures

The risk metric that is most closely associated with modern risk management is VaR. I use this metric to show how risk management failures can result when a risk metric does not provide the right information to top management not because it is calculated incorrectly, but because it answers the wrong question. A useful analogy is a weather forecast. If you are deciding to go on a three-day hike in the wilderness, a weather forecast for the first day is not really helpful when

making the decision to go or not—yes, you can still turn back in the middle of the first day, but you suffer the brunt of the bad weather if it occurs in the second day. Choosing the right risk metrics is similar to choosing the right weather forecast horizon.

The daily VaR measure is widely used in financial institutions for assessing the risk of trading activities. Though VaR has proved to be an extremely useful risk measure, it is useful only as long as the question it answers is well understood. VaR is the largest loss the firm expects to incur at a given confidence level, which means that VaR tells us nothing about the distribution of the losses that exceed VaR. This limitation of VaR has nothing to do with how institutions estimate their VaR. Rather it is simply a matter of how VaR is defined: VaR is an estimate of the *minimum worst loss* expected, as opposed to the *expected worst loss*. Though appropriate for some purposes, VaR is likely to be useless for others.

Large banks usually disclose data on that measure quarterly, and will generally report the number of times in a quarter the P&L had a loss that exceeded the daily VaR. For instance, in its annual report for 2006, UBS reported that it never had a loss that exceeded its daily VaR. But in 2007, it reported that it exceeded its daily VaR 29 times. The change from 2006 to 2007 may well indicate not an increase in the bank's risk position, but fundamental changes taking place in the economy that made it difficult for risk managers to track risk on a daily basis. Moreover, such a large number of VaR exceedances provides little or no information about the implication of these exceedances for the financial health of UBS. The larger number of exceedances could have been attributable mainly to rapid increases in volatility—increases that were anticipated by the traders and resulted in large trading gains. Or, there could have been some very large losses but few large gains. In the former case, the firm could be ahead at the end of the year; in the latter it could be in serious trouble.

Thus, though calculation of the daily market VaR can be intellectually satisfying for risk managers (since it involves up-to-date quantitative techniques), an exclusive focus on it could lead risk managers to ignore critical dimensions of risk. The daily market VaR should generally not be the primary focus of top management, which must pay attention to longer-run indicators and implications of risk. As I show with an example below, short-run VaR measures can continue to mislead management by suggesting low risk until a series of events produces a huge loss and the firm fails.

Consider a firm that has a one-day VaR of \$100 million for its trading book at the 1% probability level. This means that the firm has a one percent chance of losing more than \$100 million. If this firm exceeded its VaR once in 100 trading days but lost \$10 billion on that one occasion, all existing statistical tests

of risk management performance based on VaR exceedances would indicate that the firm has excellent risk management. In other words, VaR does not capture catastrophic losses that have a small probability of occurring.

I have not seen monthly VaR estimates from LTCM. What we do know is that, from March 1994 to December 1997, LTCM had only eight months with losses and the worst was 2.9%. In contrast, it had 37 months with gains.⁶ As a result, one would have a hard time using historical monthly returns to conclude that its risk management was flawed. Yet, in August and September 1998, LTCM experienced losses that were not predictable from its historical record. It was as if an earthquake struck. The best daily or monthly VaR measure would not have shown the possibility of such an event because the probability of an earthquake in any day or month is always extremely low. With our hypothetical LTCM example, a *one-year* VaR at the 99% confidence level would have been -70%. But the *one-day* or *one-month* VaR at the 99% confidence level could have been much lower because the possible loss of -70% would have been reflected in VaR only at a much higher confidence level.

Daily VaR measures also implicitly assume that assets can be sold quickly or hedged, so that a firm can limit its losses essentially within a day. But both in 1998 and in the past year, we have seen that markets can become suddenly less liquid, so that daily VaR measures lose their meaning. If a firm sits on a portfolio that cannot be traded, a daily VaR measure is not a measure of the risk of the portfolio because the firm is stuck with the portfolio for a much longer period of time.

Mismeasurement of Known Risks

Now let's turn to the case where management has chosen the right metrics, but the risks have been measured incorrectly, for a variety of reasons discussed below. In the LTCM example presented earlier, the mismeasurement of risk could have taken a number of different forms.

When measuring risk, risk managers attempt to understand the distribution of possible returns. In our simple example, we chose to use a binomial distribution—the kind of distribution that would result from the toss of an appropriately loaded coin. Risk managers could make a mistake in assessing the *probability* of a large loss, or they could be wrong about the size of the loss *given* that it takes place. Or they could use the wrong distribution altogether. Further, in the case of a financial institution with many positions, although the distribution associated with each position may be estimated properly, the correlation among the different positions may be mismeasured. Such correlations are extremely important in risk management because, as correlations increase, so does the vulnerability of a portfolio to low-probability events.

6. These monthly returns are for Long-Term Capital Management, L.P. (B), prepared by André Perold, Harvard Business School, 1999.

In the LTCM example, the true probability of a loss of 70% may have been higher than 1%, perhaps as high as 25%. In that event, the expected return of LTCM in my hypothetical example would have been a paltry 1.25%. At the time, investors could have earned a higher expected return by investing in T-bills. And such a risk management mistake—assessing the probability of the bad outcome at 1% instead of 25%—would have led the fund managers to make trades that, on an *ex ante* or probabilistic basis, were *expected* to reduce value.

But if LTCM made such a mistake, how would we ever know? The answer is that we cannot know based only on the information that it made a loss of 70%. We cannot identify such a mistake *after the fact* because LTCM lost 70% on only one occasion. *Given the assumed distribution of possible losses*, LTCM was capable of losing 70%, whether the true probability of that loss was 1% or 25%. And given the hypothetical conditions of this example, we can conclude nothing about the effectiveness of LTCM's risk management from its 70% loss.

Another possible failure of risk measurement would be to mistake some other probability distribution for a standard binomial one. For instance, in our hypothetical example, in addition to a 1% chance of a 70% loss, there may have been a 9% chance of a 100% loss. In this case, although the expected return would have been a respectable-sounding 12.8%, the expectation of a 9% chance of a total wipeout might have caused LTCM's partners to invest differently. Nevertheless, in January 1998, the probability of such a loss given the portfolio of trades they had assembled (or were planning to assemble) could well have been very small—a once-in-a 100-year event—or it might have been as high as 25% (though this seems unlikely). But, unless we can evaluate those probabilities in a way that doesn't depend mainly on hindsight, we cannot pass judgment on risk management.

As already noted, when an institution has many positions or projects, the risk of the institution depends on how the risks of the different positions or projects are related. If the correlation between the positions or projects is high, it is more likely that all the firm's activities will perform poorly at the same time, which leads to a higher probability of a large loss. These correlations can be difficult to assess—and they can change over time, at times abruptly. One of the partners of LTCM attributed much of the fund's losses in August and September of '98 to sudden sharp increases in correlations that were previously thought to be extremely small.

Partly as a result of the experience in 1998 of funds like LTCM, it is now well-known that correlations increase in periods of crisis. And to the extent risk managers underestimate the correlations between asset classes, they may understate the risk of a portfolio or a company. In most cases, however, the problem of mismeasurement of correlations is more subtle, since correlations evolve randomly over time and sometimes jump unexpectedly. In such cases—and LTCM

may well have been one of them—risk managers could not have been expected to anticipate such an abrupt increase in correlations. But, once such shifts in correlations are recognized, risk managers' assessments of the risk of a portfolio or company must take account of them.

The Limits of Risk Measurement. To sum up, then, statistical techniques are generally used to estimate the distribution of known risks. Such approaches work well when there is a lot of data and when it's reasonable to expect future returns to have the same distribution as past returns. For instance, suppose that a risk manager wants to estimate the volatility of the return of a liquid stock. She will have hundreds of data points to fit a model of the volatility of the stock; and in most cases that model will perform reasonably well.

But, in other cases, historical data will be of little use—say, because a risk has not manifested itself in the past. For instance, prior to the subprime crisis, there had been no experience of a downturn in the real estate market when large amounts of securitized subprime mortgages were outstanding. In such a situation, it was not possible to obtain a distribution of losses associated with a sharp downturn in real estate by using only historical data. To evaluate the risks of underwriting and purchasing mortgage-backed securities whose realization could not be observed in the past, a risk manager would have needed to understand both the likelihood of a decrease in real estate prices, and the expected effect of such a decrease on the prices of those securities. Such an exercise would have been complex and worth undertaking only if the likelihood of a large decrease in real estate prices was recognized as material from the start. But without the historical data to evaluate such risks, risk managers cannot model them using conventional quantitative approaches.

In such cases, if probability assessments are made by risk managers, they are bound to have significant elements of subjectivity. Different risk managers can reach very different conclusions. When that happens, statistical risk measurement reaches its limits and risk management changes from science to art. And once risk management moves away from established quantitative models, it becomes easily embroiled in intra-firm politics. At that point, the outcome depends much more on the firm's risk appetite and culture than its risk management models.

Mismeasurement Stemming from Overlooked Risks

Now let's turn to a different kind of measurement failure—failure to take account of risks. In my taxonomy, such “ignored” risks can take two different forms that have different implications for a company. First, a firm's risk managers may ignore a known risk, perhaps because of a mistaken assumption that it is immaterial, or because of the difficulty of incorporating it in the risk models. Second is the case of risks that are truly unknown, or at least completely unanticipated.

Ignored Known Risks

Let's now consider the case of risk management failure in which material risks are known (to at least some people inside the firm), but not reflected in the firm's risk models. In a well-functioning, truly enterprise-wide risk management system, all major risks would be identified, monitored, and managed on a continuous basis. In practice, of course, many corporate risks tend to be managed in decentralized "silos," which increases the possibility that known material risks are excluded from the central risk modeling process. When this happens, risks that are not captured by the system are not adequately monitored. And given the tendency of unmonitored risks to expand within large organizations, such risks can become material if they weren't before.

For example, consider a trader whose risks are only partly measured and monitored. Most traders have a compensation formula that involves an option-like payoff: they receive a significant share of the profits they generate, but do not have to pay back the losses. Such compensation creates incentives for traders to take risks. If only some of the risks of traders are monitored, they can increase their expected compensation by increasing the risks that are not monitored. Another example is the situation where new financial instruments are not yet incorporated in risk models. Some firms prohibit trading of financial instruments unless they are properly modeled—but such a choice has costs since it is often extremely valuable for traders to be active at the beginning of a new market. However, if financial instruments are not included in the models, traders are tempted to build up positions in these instruments, creating the possibility of large losses.

In financial institutions, risks are commonly divided into three categories: market, credit, and operational risks. These distinctions are partly artificial in the sense that they are driven by regulatory considerations. Although the trading books of financial firms are typically marked-to-market, their credit books use accrual accounting. But securities held in trading books can have credit risk as well as market risk. For instance, a bank might have credit default swaps in its trading book. In some credit default swaps, the bank is the purchaser of protection. But whether the bank actually receives the expected compensation in the event of a default depends on the ability of the protection seller to pay that compensation. And because of this somewhat artificial distinction between credit and market risk, a bank may fail to adequately take into account the counterparty risk of the credit default swaps it holds on its books.

Similarly, normal business risks—which can be quite different from those covered under the Basle II rules' narrow definition of operational risk—are often of critical importance and have to be carefully assessed as part of the evaluation of a

firm's risk. Complicating matters, such risks could turn out to be highly correlated with both credit and market risks. For instance, for many banks the loss of income from securitization was the realization of a business risk that turned out to be correlated with not only market risk—namely, the loss in value of the securities issued through securitizations—but also credit risk—the inability to use securitization to lay off the risks associated with loans.

Thus, accounting for all the risks in a risk measurement system is a difficult and costly undertaking. But a failure or refusal to attempt it means that the firm's top executives are managing the company with blinders on—they see only part of the big picture they need to manage effectively.

There are well-known examples of large losses from risks missed due to an incomplete inventorying of the firm's risks. One of the best involved the former Union Bank of Switzerland. In the second half of the 1990s, the bank was putting together risk management systems that would aggregate risks within its trading operations. One group of traders that focused on equity derivatives was extremely successful. But because this group was using different computers from the rest of the bank, integrating their systems with the bank's would have required the group to make changes in computers that would have involved some costs and downtime. Eventually, the bank's top management decided that it was more important to allow the traders to continue making money than disrupt their operations. Soon after this decision, this group of traders suffered a large loss—one that was partly responsible for the bank's forced merger with another Swiss bank.⁷

Unknown Risks

Most unknown risks do not create risk management problems. To see this, let's go back to the statistical model of risk measurement for an individual stock (as opposed to a portfolio of stocks). Suppose a risk manager has modeled the return of a stock using the normal distribution and has no reason to believe that future returns will come from a different distribution than the one that held in the past. According to this model, each period's stock return will be a random outcome—like the outcome of the toss of a coin. Since it will come from a known distribution, the risk manager does not need to know why the return of the stock was 10% in one period and -15% in another. He has captured the relevant risk characteristics of the stock through his estimation of the statistical distribution of the returns of the stock. His model tells him that the volatility (expressed as a standard deviation) of the stock's return is 20% per year, and that there is a 5% chance of a loss of, say, 30% or more over that period. He does not need to understand or explain why a loss took place. If the stock drops by 15% because of the manifestation of

7. See Dirk Schütz, *La Chute de l'UBS: Les raisons du déclin de l'Union de Banques Suisses*, Bilan, 1998.

some specific risk, the fact that the risk manager did not know about this specific risk is generally irrelevant in risk management as long as the risk manager has a good understanding of the distribution of the stock's return. For instance, this risk manager could not be faulted for having failed to anticipate the possibility of a fire destroying a factory and resulting in a 15% loss in that stock.

Unknown risks may not matter simply because they have a very low probability. For example, there is some probability that a building will be hit by an asteroid. But such a risk does not affect managerial decisions. The unknown risks that matter for our purposes are those that, had top management been aware of them, would have resulted in different actions. For example, the risk of a drop in housing prices of 30% from peak to trough could well have been viewed as an unknown risk in the past. But now, of course, that possibility has materialized in a number of regions. Had people been aware of that risk a few years back, many decisions would likely have been different.

In sum, part of the risk manager's job is to think hard about all major sources of uncertainty, perhaps using scenario planning (a suggestion I come back to later) as a methodical way of identifying and anticipating unknown risks. But having said that, some risks will inevitably remain unknown until they actually materialize. After all, how many corporate risk managers anticipated the events of 9/11? And acknowledging this possibility of unknown and unknowable risks, risk managers at some point must concede the limitations of their models—and perhaps recommend that additional capital be set aside for this possibility.

Communication Failures

Risk management is not an activity undertaken by risk managers for the benefit of risk managers. Nor should it ever be viewed as an operation or profit center unto itself. Instead, it should be designed to help the firm maximize value by “protecting” its ability to make the optimal investment and operating decisions.

But, as noted earlier, it is not the job of risk management to determine the firm's overall target level of risk or the kinds of risks it takes. Those decisions are the purview of top management and the board—of people who, when setting the firm's strategy, should begin by identifying the firm's competitive advantages and the risks its investors are “paying” the firm to take. (For example, should the airlines really be in the business of taking oil price risk, or should that be laid off to commodity traders?)

The role of risk management in such strategic decision-making is to provide timely information to the board and top management that allows them to assess the consequences of

retaining or laying off risks. But for managements and boards to understand such risks and their consequences for the firm's operations, risk managers must communicate effectively. Even if a firm has the best possible risk systems, if the risk manager is unable to make the top management understand their output, the systems may do more harm than good. If viewed as a “black box,” such a system could either lead to distrust and excessive conservatism on the part of top management—or excessive risk-taking stemming from an exaggerated sense of the protection risk management can provide.

Communication failures appear to have played a role in the recent crisis. For example, UBS published a report for its shareholders in which it discusses the causes of its subprime-related write-downs. In that report, it notes that

...a number of attempts were made to present subprime or housing related exposures. The reports did not, however, communicate an effective message for a number of reasons, in particular because the reports were overly complex, presented outdated data or were not made available to the right audience.” (p. 39).

And as stated in a report by the Senior Supervisors Group (which includes top regulators from the U.S., England, and Germany), “In some cases, hierarchical structures tended to serve as filters when information was sent up the management chain, leading to delays or distortions in sharing important data with senior management.”⁸ Finally, in the words of an industry commission report on the crisis, “risk monitoring and management reduces to the basis of getting the right information, at the right time, to the right people, such that those people can make the most informed judgments possible.”⁹

Failures in Monitoring and Managing Risks

Risk management is responsible for making sure that the firm takes the risks that it wants to take and not others. As a result, risk managers must constantly monitor the risks the firm is taking. Further, they have to hedge or otherwise manage known risks to meet the objectives of top management.

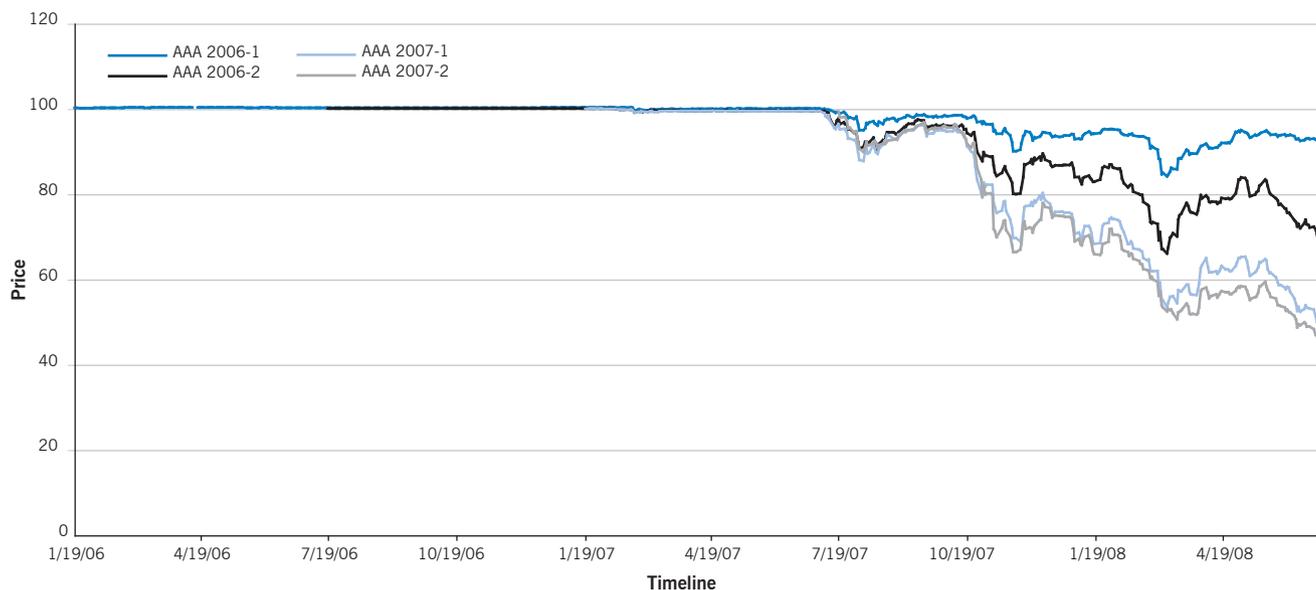
We have already discussed the problem that a firm may be taking risks it does not know about. When we discussed that problem, we focused on it as an “inventory” issue—that is, as a matter of identifying and measuring risks that are fairly stable. Such an approach is likely to be well suited to industrial companies, where risks tend to change slowly. But it may have to be modified for financial firms, where risks can change abruptly even if the firm does not take new positions.

The problem arises from the fact that financial firms have many derivatives positions and positions with embedded derivatives. The risk properties of portfolios of derivatives

8. Senior Supervisors Group, “Observations on Risk Management Practices during the Recent Market Turbulence,” March 6, 2008, p. 9.

9. “Containing systemic risk: The road to reform,” The Report of the CRMPG III, August 6, 2008.

Figure 1 Evolution of ABX AAA Tranches Based on Subprime Residential Mortgage-backed Securities (RMBS)



can change very rapidly with no trading whatsoever. This is because complex derivatives often have exposures to risk factors that are extremely sensitive to market conditions. For example, we now have products that can start the day with a significant positive exposure to an interest rate increase but, as a result of a small change in rates, end the day with a significant negative exposure. For such products, hedges adjusted daily could end up creating large losses because the hedge that was optimal at the start of the day could be amplifying risk at the end of the day.

One of the most obvious demonstrations of how risk exposures can change is the pricing of subprime derivatives. The ABX indices, which are indices based on a basket of credit-default swaps written on subprime securitization tranches, have been the most readily available source of data on the value of securities issued against subprime mortgage collateral. New indices have been created every six months that reflect new securitizations. Initially, the AAA indices, which represent the pricing of credit default swaps on AAA-rated tranches, showed almost no variation. And so reasonable assessments of the risk of the AAA-rated tranches using historical data would have indicated little risk. But, as can be seen in Figure 1, in the second-half of 2007, the value of these securities fell off a cliff. Holders of AAA-rated tranches of subprime securities reported sudden large losses if they chose to use the ABX indices as proxies for the value of their holdings.

When the risk characteristics of securities can change this rapidly, it is challenging for risk monitors to capture these changes and risk managers to adjust their hedges. This challenge is especially great when risk can change dramatically in response to small changes in the determinants of security prices. As a result, risk managers may fail to adequately monitor or hedge risks simply because the risk characteristics of securities may change too quickly to allow the managers to assess them and put on effective hedges.

Especially in these circumstances, an important component of risk management is to identify possible solutions that can be implemented quickly if a firm has to reduce its risk over a short period of time. Contingency hedging plans are therefore critical for responding to unexpected difficulties. Moreover, it's important to recognize that when liquidity dries up in the markets, many risk-reducing options that are effective in normal periods can no longer be used.

Paradoxically, the introduction of mark-to-market accounting has made it even harder for risk managers to estimate risk and put on effective hedges.¹⁰ In many ways, mark-to-market has introduced the Heisenberg Principle into financial markets: For large organizations, observing the value of a complex security affects the value of that security. The reason for this is straightforward: When mark-to-market losses become known, they set off a chain reaction of adjustments at other institutions and affect the prices of possible

10. For a discussion of some of the issues concerning mark-to-market accounting that accounts for possible feedback effects, see Guillaume Plantin, Haresh Sapra, and Hyun

Song Shin, Marking to Market: Panacea or Pandora's Box?, 2008, *Journal of Accounting Research* 46, 435-460.

trades as the market better understands the capital positions of different institutions.

In large complex organizations, it is also possible for traders to take risks that remain hidden for a while. Such positions could be constructed with securities whose risks escape detection, at least for a time, by even the best designed risk management models and systems. How can this happen? All organizations are forced to make tradeoffs, and risk management is no exception. A company's risk management function could, at least in theory, be designed to know everything at all times. But if it were organized that way, the risk management function, besides being hugely costly, would likely stifle innovation and reduce the competitiveness of the firm. In fast-moving markets, employees need flexibility.

But, as suggested, that same flexibility makes it possible for unobserved pockets of risk to emerge. And when such risks manifest themselves, it is not clear that they represent a risk management failure. Although risk management could have ensured that these risks were not taken, the firm and its shareholders would have been worse off. And even ignoring its effects on flexibility and innovation, improvements in risk monitoring are expensive enough that, at some point, they cease to be cost-effective.

One final point here: The effectiveness of risk monitoring and control depends crucially on an institution's culture and incentives. If risk is everybody's business, it is harder for major risks to go undetected and unmanaged. Moreover, if employees' compensation is affected by risks, they will be more judicious in taking them. The best risk models in a firm with limited incentives are likely to be much less effective than in a firm where the incentives of employees are better aligned with the risk-taking objectives of the firm.

Lessons from Failure (or How to Prepare for the Next Crisis)

To come up with useful assessments of risk, risk managers must look at longer horizons and take a comprehensive view of their risks. For example, one-year horizons are widely used as measures of firm-wide risk in enterprise-wide risk management (ERM) programs. Most financial institutions that focus on one-year measures of firm-wide risk aim for credit ratings that imply an extremely small yearly probability of default—say, 0.03%. Such approaches are useful not only in assessing a firm's risk, but also in estimating the optimal amount and cost of capital both for the entire firm and its various projects and lines of business.

But such approaches are not sufficient. A high target credit rating effectively means that the firm tries to avoid default in all but the most extreme circumstances. When a firm aims for an AA credit rating, it effectively chooses a probability of defaulting of roughly one in 1,000 in the coming year. But since crises are more likely than that, management needs to understand the implications of its decisions *in the event of a*

crisis and to have a strategy for responding to it.

Risk models, as we have already seen, are generally not designed to capture risks associated with crises and help companies manage them. The models use historical data and, particularly when using risk measures such as VaR, are most precise for horizons that are numbered in days; and when using such short horizons, crises appear to be highly improbable events. But, when the horizon expands to years, the probability of a crisis becomes material, something clearly worth management's attention.

And at least for purposes of evaluating the consequences of financial crises and planning for them, an extension of the horizon of risk management models is likely to be valuable for at least two reasons: First, as illustrated dramatically by events since the summer of 2007, financial crises can involve the abrupt withdrawal of liquidity from the markets. And the absence of liquidity means that firms are stuck with positions they did not expect to hold. Positions whose risk was evaluated over a single day suddenly became positions that had to be held for weeks or months. Second, during crisis periods, companies will repeatedly experience losses that exceed their daily VaRs, substantially weakening their own capital positions. For these reasons, even firms whose trading is guided mainly by daily VaR should consider complementing their use with longer-term measures.

As we have also seen in the case of LTCM, financial crises involve complicated linkages and interactions among risks and institutions. Statistical risk models typically take returns to be “exogenous” to the firm and ignore risk concentrations across institutions. But if such an approach is appropriate for many institutions, it will not work for others that, because of their size or degree of connectedness, can affect prices and volumes in certain markets. For instance, it is well-known that LTCM had extremely large short positions in the index option market. During the crisis, it had little ability to adjust these positions because they were such a large part of the market.

Further, large institutions facing distress can be exposed to “predatory” trading—that is, trades made by others that have the effect of compounding the losses from such large, illiquid positions. An example is a situation where traders from other institutions benefit from selling and pushing prices down with the aim of bringing about a fire sale (and buying back their positions). As one firm experiences large losses, it may drag down prices for other institutions and make funding more costly for all of them. Typical risk management models do not account for such possibilities, leading to a significant understatement of the risk of positions in the event of a crisis.

But, as already suggested, there is little hope for statistical risk models that rely on historical data to capture such complicated effects. And rather than attempting to introduce greater complexity and realism into their models, risk manag-

ers should consider complementing their reliance on models with tools such as scenario analysis to investigate how crises unfold, how the firm will be affected by them, and how it should best respond to them. With the help of such scenarios, top managements can anticipate the threat posed by extreme events to the franchise value of their institutions and develop strategies for limiting and responding to such events.

But contrary to some current practice, such a scenario approach must be informed by economic and financial analysis. It cannot be done by risk management departments populated only by physicists and mathematicians. Such an approach also cannot succeed unless top management believes that the scenarios considered represent serious threats to the institution—serious enough to affect the formulation of the firm's strategy and to require a plan to deal with it.

Conclusion

The difficulties of the past year have convinced some observers that there are important, and perhaps irremediable, flaws in risk management, and that such flaws were major contributors to the current financial crisis. In this paper, I show the need to distinguish between flawed assessments by risk managers and corporate risk-taking decisions that, although resulting in losses, were fundamentally reasonable at the time they were made. To help in making this distinction, the paper also identifies a number of different ways that risk management can fail. In addition to choosing the wrong risk metrics and otherwise mismeasuring risks, risk managers can fail to communicate their risk assessments and otherwise provide

effective guidance to top management and boards. And once top management has used that information to determine the firm's risk appetite and strategy, risk management can also fail to monitor risks appropriately and maintain the firm's targeted risk positions.

But if risk management has been mistakenly identified as the culprit in many cases, there is no question that risk management practice can be improved by taking into account the lessons from financial crises past and present. Such crises have occurred with enough frequency that crisis conditions can be modeled, at least to some extent. And when models reach their limits of usefulness, companies should consider using scenario planning that assesses the implications of crises for their financial health and survival. Rather than relying on past data, scenario planning must use economic analysis to evaluate the expected impact of sudden illiquidity and the associated feedback effects that are common in financial crises. But, to serve as an effective part of a firm's strategy, scenario planning and the analysis that comes out of it must be deeply rooted in a firm's culture as well as the strategic thinking of top management.

RENÉ M. STULZ holds the Reese Chair of Banking and Monetary Economics at Ohio State's Fisher College of Business. He is also a research associate at the National Bureau of Economic Research and a fellow at the European Corporate Governance Institute. He belongs to the Board of Trustees of the Global Association of Risk Professionals and its executive committee.

Journal of Applied Corporate Finance (ISSN 1078-1196 [print], ISSN 1745-6622 [online]) is published quarterly, on behalf of Morgan Stanley by Blackwell Publishing, with offices at 350 Main Street, Malden, MA 02148, USA, and PO Box 1354, 9600 Garsington Road, Oxford OX4 2XG, UK. Call US: (800) 835-6770, UK: +44 1865 778315; fax US: (781) 388-8232, UK: +44 1865 471775.

Information for Subscribers For new orders, renewals, sample copy requests, claims, changes of address, and all other subscription correspondence, please contact the Customer Service Department at your nearest Blackwell office (see above) or e-mail customerservices@blackwellpublishing.com.

Subscription Rates for Volume 19 (four issues) Institutional Premium Rate* The Americas[†] \$377, Rest of World £231; Commercial Company Premium Rate, The Americas \$504, Rest of World £307; Individual Rate, The Americas \$100, Rest of World £56, €84[‡]; Students** The Americas \$35, Rest of World £20, €30.

*The Premium institutional price includes online access to current content and all online back files to January 1st 1997, where available.

[†]Customers in Canada should add 6% GST or provide evidence of entitlement to exemption.

[‡]Customers in the UK should add VAT at 6%; customers in the EU should also add VAT at 6%, or provide a VAT registration number or evidence of entitlement to exemption.

**Students must present a copy of their student ID card to receive this rate.

For more information about Blackwell Publishing journals, including online access information, terms and conditions, and other pricing options, please visit www.blackwellpublishing.com or contact your nearest Customer Service Department.

Back Issues Back issues are available from the publisher at the current single-issue rate.

Mailing *Journal of Applied Corporate Finance* is mailed Standard Rate. Mailing to rest of world by DHL Smart & Global Mail. Canadian mail is sent by Canadian publications mail agreement number 40573520. **Postmaster** Send all address changes to *Journal of Applied Corporate Finance*, Blackwell Publishing Inc., Journals Subscription Department, 350 Main St., Malden, MA 02148-5020.

Journal of Applied Corporate Finance is available online through Synergy, Blackwell's online journal service, which allows you to:

- Browse tables of contents and abstracts from over 290 professional, science, social science, and medical journals
 - Create your own Personal Homepage from which you can access your personal subscriptions, set up e-mail table of contents alerts, and run saved searches
 - Perform detailed searches across our database of titles and save the search criteria for future use
 - Link to and from bibliographic databases such as ISI.
- Sign up for free today at <http://www.blackwell-synergy.com>.

Disclaimer The Publisher, Morgan Stanley, its affiliates, and the Editor cannot be held responsible for errors or any consequences arising from the use of information contained in this journal. The views and opinions expressed in this journal do not necessarily represent those of the Publisher, Morgan Stanley, its affiliates, and Editor, neither does the publication of advertisements constitute any endorsement by the Publisher, Morgan Stanley, its affiliates, and Editor of the products advertised. No person should purchase or sell any security or asset in reliance on any information in this journal.

Morgan Stanley is a full service financial services company active in the securities, investment management, and credit services businesses. Morgan Stanley may have and may seek to have business relationships with any person or company named in this journal.

Copyright © 2008 Morgan Stanley. All rights reserved. No part of this publication may be reproduced, stored, or transmitted in whole or part in any form or by any means without the prior permission in writing from the copyright holder. Authorization to photocopy items for internal or personal use or for the internal or personal use of specific clients is granted by the copyright holder for libraries and other users of the Copyright Clearance Center (CCC), 222 Rosewood Drive, Danvers, MA 01923, USA (www.copyright.com), provided the appropriate fee is paid directly to the CCC. This consent does not extend to other kinds of copying, such as copying for general distribution for advertising or promotional purposes, for creating new collective works, or for resale. Institutions with a paid subscription to this journal may make photocopies for teaching purposes and academic course-packs free of charge provided such copies are not resold. Special requests should be addressed to Blackwell Publishing at: journalsrights@oxon.blackwellpublishing.com.