



FISHER
COLLEGE OF BUSINESS

*Fisher College of Business
Working Paper Series*

*Charles A. Dice Center for
Research in Financial Economics*

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

René M. Stulz, Department of Finance, The Ohio State University,
NBER, and ECGI

Dice Center WP 2008-18
Fisher College of Business WP 2008-03-017

October 2008

This paper can be downloaded without charge from:

<http://www.ssrn.com/abstract=1278073>

An index to the working paper in the Fisher College of
Business Working Paper Series is located at:

<http://www.ssrn.com/link/Fisher-College-of-Business.html>

Risk management failures: What are they and when do they happen?

René Stulz*

October 2008

Abstract

A large loss is not evidence of a risk management failure because a large loss can happen even if risk management is flawless. I provide a typology of risk management failures and show how various types of risk management failures occur. Because of the limitations of past data in assessing the probability and the implications of a financial crisis, I conclude that financial institutions should use scenarios for credible financial crisis threats even if they perceive the probability of such events to be extremely small.

* Reese Chair of Banking and Monetary Economics, Fisher College of Business, Ohio State University, NBER, and ECGI. I am grateful for assistance from Jérôme Taillard and Mike Anderson, and for comments from Rich Apostolik, Don Chew, Cliff Smith, and Peter Tufano.

In commentaries on the financial crisis that started during the summer of 2007, a constant refrain is that somehow risk management failed and that there were risk management failures at financial institutions across the world. For instance, an article in the Financial Times states that “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 examine what it means for risk management to fail. I show that the fact that an institution makes an extremely large loss does not imply that risk management failed or that the institution made a mistake. This article does not examine the subprime financial crisis or problems of financial institutions during that crisis directly. Rather, it is an attempt to make sure that if risk management is blamed, it is for the right reasons. Otherwise, changes in risk management that take place in response to the crisis might be counterproductive and top executives and investors could keep expecting more from risk management than what it can actually deliver. I therefore show when bad outcomes can be blamed on risk management and when they cannot. In the process of doing so, I provide a typology of risk management failures.

To examine risk management failures more concretely, I go back to the problems experienced by the hedge fund LTCM in 1998 to analyze how one might conclude that the failure of LTCM was a risk management failure or not. I then generalize from that example to describe what constitutes a risk management failure and what does not. I will show that some events considered in the financial press to be risk management failures actually are not risk management failures, but at the same time I will analyze many different ways in which risk management can fail. I then address the question of whether lessons from risk management failures can be used to help improve the practice of risk management. In the last part of this article, I discuss an approach to risk management that might enable institutions to better manage risks such as those that threatened them during the subprime financial crisis.

¹ “Wall Street dispatch: Imagination and common sense brew a safer culture,” by David Wighton, Nov. 26, 2007, FT.com.

Was the collapse of Long-Term Capital Management a risk management failure?

The story of Long-Term Capital Management (LTCM) is well-known.² In 1994, ex-Salomon Brothers traders and two future Nobel Prize winners started a hedge fund, the Long-Term Capital Fund. LTCM was the company that managed the fund. The fund performed superbly for most of its life: Investors earned 20% for ten months in 1994, 43% in 1995, 41% in 1996, and 17% in 1997. In August and September 1998, following the default of Russia on its ruble denominated debt, world capital markets were in crisis and the hedge fund LTCM lost most of its capital. Before its collapse, LTCM had capital close to \$5 billion, assets in excess of \$100 billion, and derivatives for 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 in the fund.

Does a loss of more than 70% of capital represent a risk management failure? Does a loss that requires a rescue by banks involving an injection of \$3.65 billion of new capital show that risk management failed? It turns out that it is not easy to answer these questions. To define a risk management failure, one must first define the role of risk management.

In a typical firm, the role of risk management is first to assess the risks faced by the firm, communicate these risks to those who make risk-taking decisions for the firm, and finally manage and monitor those risks to make sure that the firm only bears the risks its management and board of directors want it to bear. In general, a firm will specify a risk measure that it focuses on together with additional risk metrics. When that risk measure exceeds the firm's tolerance for risk, risk is reduced. Alternatively, when the risk measure is too low for the firm's risk tolerance, the firm increases its risk. Because firms are generally more concerned about unexpected losses, a frequently used risk measure is Value-at-Risk or VaR, a measure of downside risk. VaR is the

² The best public source for data on LTCM is a collection of four case studies by André Perold published in 1999, 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, Random House, 2000.

maximum loss at a given confidence level over a given period of time. Hence, if the 95% confidence level is used and a firm has a one-day VaR of \$150 million, the firm has a 5% chance of making a loss in excess of \$150 million over the next day if the VaR is correctly estimated. This measure might be estimated daily or over longer periods of time.

Even with our definition of the role of risk management, the returns of LTCM do not tell us anything 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, overall, had a 99% chance of producing a return for the fund before fees of 25% and a 1% chance of making a loss of 70% over the coming year. Though this example is hypothetical, it is plausible in light of the returns of LTCM and what LTCM was telling its investors. First, 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 wrote to its investors to tell them that it expected that the fund would experience a loss in excess of 20% only in one year out of 50 – here, instead, one year out of 100 can be expected to have a loss of 70%.³ Let's assume that whether the fund had the high return or not depended on the flip of a weighted coin, so that the risk of the fund would have been completely diversifiable for its investors. With this hypothetical example, the expected return on the fund would then have been 24.05%. Such an expected return would have been a great expected return for a hedge fund or for any investment as this would have been the expected return for bearing diversifiable risk, given my assumptions. Had the managers had the opportunity to keep repeating this investment, 99 years out of 100 they would have earned 25% before fees and would have been stars.

In my hypothetical example, when the managers of the funds (the partners) made their choice, they knew the true distribution of possible outcomes of the fund. Hence, they knew the distribution of gains and losses perfectly – the risk managers should have earned a gold medal for their work. Suppose, however, that the bad outcome occurs. In this case, the fund would have

³ See Lowenstein, p. 63.

made headlines for having lost \$3.5 billion. Some would argue that the risk of the fund was poorly managed. However, by construction, risk management could not have been improved in this case. The managers knew exactly the risks they faced – and they decided to take them. Therefore, there is no sense in which risk management failed. Ex post, the only argument one could make is that the managers took risks they should not have, but that is not a risk management issue as long as the risks were properly understood. Rather, it is an issue of assessing the costs of losses versus the gains from making large profits.

Deciding whether to take a known risk is not a decision for risk managers. The decision depends on the risk appetite of an institution. However, defining the risk appetite is a decision for the board and top management. That decision is at the heart of the firm's strategy and of how it creates value for its shareholders. A decision to take a known risk may turn out poorly even though, at the time it was made, the expectation was that taking the risk increased shareholder wealth and hence was in the best interest of the shareholders.

In the case of LTCM, it could be argued that the cost of losing \$3.5 billion for the investors in LTCM was just that - namely, there were no additional costs beyond the direct monetary loss. For most firms, however, large losses have deadweight costs. These deadweight costs are at the foundation of financial theories of why risk management creates shareholder wealth.⁴ If a financial institution makes a large loss, the institution may, for instance, have to scale back its investments because of being financially constrained, have to sell assets in unfavorable markets, lose valuable employees who become concerned for their bonuses, lose customers who are concerned about the institution being distracted or not having sufficient resources to help them, and face increased scrutiny from regulators. In any institution, the board and top management have to take into account these deadweight costs of large losses when making decisions that create the risk of large losses.

⁴ See René M. Stulz, *Risk Management and Derivatives*, Thompson Publishing, 2003.

Risk managers can estimate whether an action is profitable for the firm given its risk appetite because they can evaluate how much capital is required to support that action.⁵ However, an action that is not profitable for a given level of risk appetite can become profitable if the firm's risk appetite increases because less capital is required to support that action. Whether taking large risks is worthwhile for an institution ultimately depends on the firm's strategy. Risk managers do not set strategy. Suppose that a firm sets its risk appetite by choosing a target credit rating. Such an approach is well-established. Once the credit rating is chosen, there are multiple combinations of risk and capital that achieve the target rating. For a given choice of leverage, the firm does not have much choice in choosing its risk level if it wants to achieve its target rating. However, faced with good opportunities, the firm could choose to have less leverage so that it can bear more risk or it could choose to depart from its credit rating target.

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 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 roughly 36% of the capital to its investors. With less capital, LTCM could still execute the same trades. However, now, to implement them it had to borrow more and hence had to increase its leverage. By increasing its leverage, it could boost the return to its shareholders if things went well at the expense of making more losses if things went poorly. Was increasing leverage a poor risk management decision? In my example, the partners of LTCM knew the risks and the rewards from doing so. In the well-worn language of financial economics, increasing leverage was a positive NPV decision when it was made, but obviously ex post it was a costly decision as it meant that when assets fell in value, the fund's equity fell in value faster than it would have with less leverage.

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.

There has been much discussion of incentives of top management during the credit crisis, with various commentators arguing that part of the problem has been that top management had incentives to take too much risk. This may well be so, but before reaching conclusions one should not forget that financial economists have argued for decades that incentives of management become better aligned with those of shareholders when management has a large stake in the firm's equity. Top management owned hundreds of millions of dollars of equity in Bear Stearns and Lehman at the peak of the valuation of these firms. Similarly, the partners of LTCM collectively had almost \$2 billion invested in the fund at the beginning of 1998. If such equity stakes do not incentivize managers to make the right decisions for their shareholders, what would?

In summary, risk management does not prevent losses. With good risk management, large losses can occur when those making the risk-taking decisions conclude that taking large, well-understood risks creates value for their organization.

A typology of risk management failures

How can risk management go wrong? The way we describe the role of risk management suggests important ways in which risk management can go wrong. We started by saying that the first step in risk management is to measure risk. Let's assume, for now, that the right risk measure is used given the situation of the firm. This measure could be VaR or could be some other measure. Two types of mistakes can be made in measuring risk: Known risks can be mismeasured and some risks can be ignored, either because they are unknown or viewed as not material. Once risks are measured, they have to be communicated to the firm's leadership. A failure in communicating risk to management is a risk management failure as well. After management decides what kind of risks to take, risk management has to make sure that the firm takes these risks. In other words, risk managers must then manage the firm's risk, a task that may involve

identifying appropriate risk mitigating actions, hedging some risks, and rejecting some proposed trades or projects. Lastly, a firm's risk managers may fail to use appropriate risk metrics.

With this perspective, there are six types of risk management failures:

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

We discuss each one of these types of failures in turn.

Mismeasurement of known risks

In the LTCM example, risk mismeasurement could have taken a number of different forms. When measuring risk, risk managers attempt to understand the distribution of possible returns. With our simple example, the distribution was a binomial distribution – the outcome of the toss of a weighted coin. Risk managers could make a mistake in assessing the probability of a large loss or the size of the large loss if it occurs. However, in addition, they could use the wrong distribution altogether. Further, financial institutions have many positions, each position has a return from a given distribution, but these returns are related across positions, and that relation may be assessed incorrectly – a simple way to put this is that correlations may be mismeasured. Correlations are extremely important in risk management because the benefit of diversification falls as correlations increase.

With the LTCM example, it could be that the true probability of a loss of 70% was higher than 1%, say 25%. In this case, 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. In this case, the risk management mistake – assessing the probability of the bad outcome at 1% instead of 25% - would have had disastrous consequences for the fund because it would have led it to make trades that would have destroyed value.

Suppose that LTCM had made the mistake we just discussed. How would we know? We cannot identify such a mistake ex post because LTCM lost 70% only once. Having lost 70%, it could have done so whether the true probability of that loss was 1% or 25%. In fact, under the hypothetical conditions of my example, we can learn nothing from the fact that LTCM lost 70% about whether it made a risk management mistake of that type. It could have been that, as of January 1998, the probability of such a loss was infinitesimal or extremely large. It could have been a one in one hundred year event or a one in four year event for the portfolio of trades they had assembled.

Another risk management mistake would occur if the distribution is not binomial, but a different distribution altogether. For instance, it could be, keeping with the hypothetical example, that there was a 1% chance of a 70% loss and additionally a 9% chance of a 100% loss. In this case, the expected return would have been 12.8%, but there would also have been a nontrivial probability of a total wipeout.

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 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 change over time, at times abruptly. A partner of LTCM described the problem they faced in August and September as being one where correlations that they thought were extremely small suddenly became large. With this perspective, correlations would have been misestimated. It is well-known in finance that correlations increase in periods of crisis. Failure to assess correlations correctly would lead to the wrong assessment of the risk of a portfolio or of a firm. The problem of mismeasurement of correlations is more subtle, however, if correlations are

random and sometimes turn out to be unexpectedly large ex post. In this case, risk managers could not be expected to know what correlations will be, but their assessment of the risk of a portfolio or of the firm would depend on their estimates of the distribution of the correlations. In this case, it would be possible for realized correlations to be different from their expected value and yet there would be no risk management failure.

When risks are known, statistical techniques are generally brought to bear to estimate the distribution of risks. Such approaches work well when there is a lot of data and when it is reasonable to believe that the returns will have the same statistical distribution in the future as they had in the past. 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. In most cases, the risk manager will have a model of the volatility of the stock that will perform reasonably well.

Historical data is at times of little use, because a known risk has not manifested itself in the past. For instance, with the subprime crisis, there was no historical data of a downturn in the real estate market during which a large amount of securitized subprime mortgages was outstanding. In such a situation, risk measurement cannot be done by simply using historical data since there is a risk of a decrease in real estate prices that has not manifested itself in a comparable historical period. With such a case, statistical risk measurement reaches its limits and risk management goes from science to art. Proper understanding of risks involves an assessment of the likelihood of a decrease in real estate prices and of the economic impact of such a decrease on the prices of securities. Such probability assessments have a significant element of subjectivity. Different risk managers can reach very different conclusions.

There is a fundamental problem with the performance of risk measurement when assessments become subjective. Suppose that all parties agree that an established statistical model works well. There is then little room for people to disagree. However, subjective forecasts are easily questioned. Why would a risk manager have a better understanding of the probability of a drop in real estate prices than experts in real estate? If experts in real estate conclude that a sharp drop in

prices is unlikely, why would an organization then listen to a risk manager who wants to spend a large amount of money on a stress test to figure out the impact of such a large drop? As risk management moves away from established quantitative models, it becomes easily embroiled in intra-firm politics. At that point, the outcome for the firm depends much more on the firm's risk appetite and on its culture than on its risk management models.

Mismeasurement due to ignored risks

Ignored risks can take three different forms that have different implications for a firm. First, a firm may ignore a risk even though that risk is known. Second, somebody in the firm knows about a risk, but that risk is not captured by the risk models. Third, there is a realization of a truly unknown risk. We examine these possibilities in turn.

Ignored known risks

Consider again the case of LTCM. LTCM could have failed to take into account a risk that, if realized, would have led to a large loss. A good example of this possibility is as follows. Before Russia defaulted on its domestic debt in August 1998, many hedge funds took positions where they bought high-yielding Russian debt, hedged the debt against default risk, and finally hedged the debt against exchange rate risk. It was easy to believe that the resulting position had no risk. However, to hedge the currency risk, the funds had to sell rubles forward against dollars. The banks willing to stand on the other side of those trades were often Russian banks. When Russia defaulted, it imposed a moratorium on these banks and many collapsed, as a result, the hedge funds ended up having exchange rate risk because their counterparties did not honor the hedges. Had they taken into account counterparty risk properly, they would have understood that their positions had substantial risk in the event of an adverse shock to the Russian banking system.

I have no reason to believe that LTCM behaved like these other hedge funds. Further, LTCM's Russian exposures were relatively small. However, suppose that it made losses because

it did not correctly account for the risks of counterparties. Ex post, just knowing that LTCM lost 70% would not be sufficient to conclude that LTCM missed the counterparty risk in its risk models because it could have made a similar loss without missing that risk. Consequently, to assess whether LTCM made mistakes, one would have to look at the information it had when it made decisions, whether that information was flawed, and whether its use of that information was wrong.

Mistakes in information collection

The consequences of a risk management mistake are the same whether the risk was ignored because nobody in the firm knew about it or because somebody knew about it but it did not enter the relevant risk models. One of the benefits of implementing properly firm-wide risk management is that all risks are accounted for. If some risks are not accounted for when risk is measured for a firm, the risks left out are not adequately monitored and they can become large because organizations have a tendency to expand unmonitored risks. For instance, consider a trader whose risks are only partly monitored. Typically, traders have a compensation formula that involves an option payoff – they receive a significant share of the profits they generate, but they do not have to give back the losses. If only some of the risks of a trader are monitored, he can increase his expected compensation by increasing the risks that are not monitored, without suffering any of the consequences.

It is common practice in risk management to divide risks into market, credit, and operational risks. This distinction is partly artificial and driven by regulatory considerations. Typically, firms have trading books that are marked to market, while the credit book uses accrual accounting. However, this division of risk may be implemented in a way that ignores large chunks of risk. For instance, a firm has funding risks. Funding may become more expensive and/or less available precisely when the firm experiences bad market outcomes. To wit, an important factor contributing to the failure of Bear Stearns was the limitations it faced in accessing the repo

market in its last week. Similarly, while Basle II rules have a rather restricted view of operational risk, business risks are often of critical importance and have to be carefully assessed as part of the evaluation of a firm's risk even though they are not part of the regulatory definition of operational risk. These risks may be highly correlated with both credit and market risks for financial institutions. For instance, for many banks, the loss of income from securitizations was the realization of a business risk that was correlated with a market risk, namely the loss in value of securities issued through securitizations, and with credit risks, namely the inability to use securitization to lay off the risks associated with loans.

Accounting for all the risks in risk measurement is a difficult and costly task. However, not performing that task for an organization means that the firm's top executives are managing the company with blinders on – they see only part of the big picture they have to understand to manage effectively. There are well-known examples of incomplete risk aggregation leading to large losses from risks that were not accounted for. Perhaps one of the best examples is the one of a bank that no longer exists, the 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. However, this group of traders was using different computers from the rest of the bank, so that integrating their systems into the bank's systems would have required them to change computers. Eventually, the bank decided, at the top level, that it was more important to let the traders make money than disrupt what they were doing through changes of computers. Soon thereafter, this group of traders lost a large amount of money for the bank. The loss was partly responsible for the bank having to merge with another Swiss bank.⁶

Problems of aggregation were important at various stages of the subprime crisis as well. In particular, the management of UBS sent a report to its shareholders explaining why the bank had such large write-downs. In this report, UBS explains that “Efforts were made to capture Subprime

⁶ See Dirk Schütz, *La Chute de l'UBS: Les raisons du déclin de l'Union de Banques Suisses*, Bilan, 1998.

holdings by mid-February 2007, however, materials did not effectively include the Super Senior and Negative Basis positions.” (p. 39). It is interesting to note that, according to the report, the Super Senior positions were not included because they were hedged and hence were assigned no risk by the risk models – an evaluation which was consistent with past data used by many risk managers.

Unknown risks

Most unknown risks do not create risk management problems. To see this, we can go back to the statistical model of risk measurement for a stock. Suppose that a risk manager models the return of a stock using the normal distribution and that he has no reason to believe that future returns will come from a different distribution than the one that held in the past. With this model, each period, the stock return will be random. It will come from a known distribution. The risk manager does not need to know why the return of the stock in one period is 10% and in another period it is -15%. He has captured the relevant risk characteristics of the stock through his estimation of the statistical distribution of the returns of the stock. With his work, he knows that the volatility is 20% and that there is a 5% chance of a loss of say 30% or higher over a period. He does not need to be in a position to explain what events are associated with various losses.

Other unknown risks may not matter simply because they have a trivially low probability. There is some probability that a building will be hit by an asteroid. That risk does not affect any management decisions. Ignoring that risk has no implications for risk management.

The unknown risks that represent risk management failures are risks that, had the firm’s managers known about them, their actions would have been different. Risk managers have to look out for unknown risks, but once everything is said and done, some risks will remain unknown. Because of this, they have to conclude that they do not capture all risks in their models and, therefore, some capital has to be made available to cope with unknown risks.

Communication failures

Risk management is not an activity undertaken by risk managers for risk managers. Rather, it is an activity undertaken to enable the firm to maximize shareholder value by taking optimal decisions across the firm. In particular, the firm has to choose the level of risk it is exposed to and has to make sure that risks taken throughout the organization are valuable for shareholders. Therefore, risk management has to provide timely information to the board and top management that enables them to make decisions concerning the firm's risk and to factor the firm's risk in their decisions. In order for the board and the top management to understand the risk situation of the firm, this situation has to be communicated to them in a way that they can understand properly. If a firm has perfect risk systems, but the board and the top management cannot understand the output of these systems because the risk manager cannot communicate this output properly, the firm's systems may do more harm than good by inspiring false confidence in the performance of risk management. Even worse, information can arrive to top management too late or too distorted by intermediaries.

Communication failures seem to have played a role in the most recent crisis. For example, the UBS report to its shareholders explains 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). An industry commission that drew lessons from the crisis emphasized communication issues as well. It concluded that "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."⁷ Finally, a report from the Senior Supervisors Group, which includes top regulators from the U.S., England, and Germany as well as other countries, also emphasized communication issues, stating for instance that "In some cases, hierarchical

⁷ "Containing systemic risk: The road to reform," The Report of the CRMPG III, August 6, 2008.

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

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 and mitigate known risks to meet the objectives of top management.

We have already discussed the problem that a firm may be taking risks that it does not know about. When we discussed that problem, we focused on it as an inventory issue. However, there is a different perspective on this problem which is particularly relevant in financial firms. For the typical non-financial firm, risks often change slowly. Not so for financial firms. For a financial firm, risks can change sharply 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. Over time, these positions have become more complex.

The risk properties of portfolios of derivatives 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. Strikingly, it is perfectly possible with some products to see changes such that, during the same day, a security could have an exposure to interest rates so that it gains substantially if interest rates increase and later in the day have an exposure such that it loses substantially if interest rates increase. For such a product, hedges adjusted daily could end up creating large losses because the hedge that is optimal at the start of the day could end up aggravating the risk exposure at the end of the day.

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

One of the most obvious demonstrations of how risk exposures can change is the pricing of subprime derivatives. The ABX indices have been the most readily available data on the value of securities issued against subprime mortgage collateral. The indices are equally-weighted averages of credit-default swaps on securitization tranches. New indices were created every six months, reflecting new securitizations. Initially, the AAA indices, which represent the pricing of credit default swaps on AAA-rated tranches of securitizations, exhibited almost no variation, so that reasonable assessments of the risk of the AAA-rated tranches of securitizations using historical data would have been that they had little risk. Yet, suddenly, the value of these securities fell off a cliff as shown on Figure 1. Holders of AAA-rated tranches of subprime securities made 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 very rapidly, it is challenging for risk monitors to capture these changes and for risk managers to adjust hedges appropriately. This challenge is especially great when risk characteristics can change dramatically for small changes in the determinants of security prices. As a result, risk managers may fail to adequately measure risks or hedge risks simply because risk characteristics of securities may change too quickly to enable these managers to assess these characteristics properly or to put on correct hedges.

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. Lack of such plans could make it impossible for a firm to cope with unexpected difficulties. At the same time, however, when liquidity dries up in the markets, many risk-mitigating options that can be used easily outside of crisis periods can no longer be used.

Paradoxically, the introduction of mark-to-market accounting makes it even harder for risk managers to estimate risk and put on adequate hedges.⁹ 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: As mark-to-market losses become known, they start a chain reaction of adjustments at other institutions and affect prices of possible trades as the market understands the capital positions of institutions better.

In large complex organizations, it is also possible for individuals to take risks that remain hidden for a while. A trader might have constructed a complicated position that only he understands. This position might be such that under some circumstances it could lead to large losses. The position might use securities that are not incorporated in the risk management systems. At all times, organizations face tradeoffs. Risk management might be structured to know everything at all times. However, if risk management were organized that way, it would stifle innovation within the firm and hamper the competitiveness of the firm. In fast moving markets, employees have to have flexibility. However, that flexibility makes it possible for unobserved pockets of risk to emerge. When these risks manifest themselves, it is not clear that they represent a risk management failure. Risk management could have made sure that these risks were not taken, but ex ante shareholders would have been worse off. Besides eliminating flexibility within the firm, risk monitoring is costly so that at some point, tighter risk monitoring is not efficient.

The effectiveness of risk monitoring and control depends crucially on an institution's culture and incentives. If risk is everybody's business in an organization, it is harder for pockets of risk to be left unobserved. If employees' compensation is affected by how they take risks, they will take risk more judiciously. The best risk models in a firm with poor culture and poor incentives will be

⁹ 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.

much less effective than in a firm where the incentives of employees are better aligned with the risk-taking objectives of the firm.

Risk measures and risk management failures

So far, we have taken the risk metrics as given. We now show that focusing on metrics that are too narrow may make it harder for management to achieve its objectives. Specifically, risks that management would consider important can be left unmeasured and ignored.

A widely used risk measure in financial institutions is a daily VaR measure for trading activities. Large banks usually disclose data on that measure quarterly. They will generally say the number of times in a quarter the P&L had a loss that exceeds the daily VaR. For instance, UBS reported in its annual report for 2006 that it never had a loss that exceeded its daily VaR. In contrast, in 2007, it reported in its annual report that it exceeded its daily VaR 29 times. The results for 2007 show that fundamental changes were taking place in the economy that made it difficult for risk managers to track risk on a daily basis. However, such a large number of VaR exceedances provide little or no information about the implication of these exceedances for the financial health of UBS. It could be that the exceedances were really small and that there were many large gains as well because volatility increased rapidly. Alternatively, there could have been very large losses and few large gains. In the former case, the firm could be ahead at the end of the year. In the latter case, it could be in serious trouble. Consequently, focusing on the daily market VaR, though intellectually satisfying for risk managers because the most up-to-date quantitative techniques can be brought to bear on the problem, can only be one part of risk management and not the one that top management should focus on. Top management has to focus on the longer-run implications of risk.

Short-run VaR measures can be low and the firm can appear to do an extremely good job with them, yet it can fail. I have not seen monthly VaR estimates from LTCM. However, from March 1994 to December 1997, LTCM had only eight months with losses and the worst monthly

loss was 2.9%. In contrast, it had 37 months with gains.¹⁰ As a result, one would have a hard time using historical monthly returns to conclude that its risk management was flawed. 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 over 100 trading days and lost \$10 billion, all existing statistical tests of risk management performance based on VaR exceedances would indicate that the firm has excellent risk management. VaR does not capture catastrophic losses that have a small probability of occurring.

Daily VaR measures assume that assets can be sold quickly or hedged, so that a firm can limit its losses essentially within a day. However, both in 1998 and over the last 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.

To assess risk, firms have to look at longer horizons and have to take a comprehensive view of their risks. A one-year horizon is widely used in enterprise risk management for measures of firm-wide risk. Generally, financial institutions that focus on firm-wide risk at a one-year horizon aim for credit ratings that imply an extremely small probability of default. Such approaches are useful in assessing a firm's risk, in estimating the optimal amount of capital for a firm, and estimating the profitability of projects and lines of business through a careful evaluation of the cost of the capital required to bear their risks. However, at the same time, 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. If a firm aims for an AA credit rating, it effectively chooses a probability of default which is such that it would default less frequently than one year out of a

¹⁰ These monthly returns are for Long-Term Capital Management, L.P. (B), prepared by André Perold, Harvard Business School, 1999.

thousand. Crises occur much more often than that, so that the firm has to have a strategy which allows it to survive crises. Further, the probability of a crisis is difficult to estimate precisely, so that even if the estimate of the probability is very small, estimation error could be such that the true, unknown, probability is much higher. Consequently, the firm has to focus on crisis events in its risk measurement and management.

Existing risk models are generally not designed to capture risks associated with crises and to help firms manage them. These models use historical data and are most precise for short horizons – like days. With short horizons, crises are extremely rare events. Yet, when we consider years, crises are not extremely rare events. Months and years are a better horizon to evaluate risk when it comes to crises for at least two reasons. First, as evidenced since the summer of 2007, crises involve a dramatic withdrawal of liquidity from the markets. The withdrawal of liquidity means that firms are stuck with positions that they never expected to hold for a long time because price pressure costs involved in trading out of these positions are extremely high. Positions whose risk was evaluated over one day because the firm thought it could trade out of these positions suddenly became positions that had to be held for weeks or months. Second, during crisis periods, firms will make multiple losses that exceed their daily VaRs and these losses can be large enough to substantially weaken them. As a result, risk measures have to contemplate the distribution of large losses over time rather than over one day.

Crises involve complicated interactions across risks and across institutions. Statistical risk models typically take returns to be exogenous to the firm and ignore risk concentrations across institutions. Such an approach is appropriate for many institutions, but it is insufficient for institutions that, for whatever reasons, are important in specific markets and whose actions affect security prices. For instance, it is well-known that LTCM had extremely large positions in the index option market where it was short. During the crisis, it had little ability to change these positions because it was so large in that market. Further, a large institution can be exposed to predatory trading – i.e., of trades made by others designed to exploit its problems. An example of

predatory trading is a situation where traders from other institutions benefit from pushing a price down if they can because it might force a fire sale. Typical risk management models would not account for this. They would not account for the fact that if the institution is large in a market, its losses can lead to more losses. As a firm makes a loss, it may drag down prices for other institutions and make funding more costly across institutions, which can have feedback effects for the institution. Ignoring these potential feedback effects may lead to an understatement of the risk of positions in the event of a crisis.

There is little hope for statistical risk models relying on historical data to capture such complicated situations. Rather, a firm has to augment these models with scenario analysis that investigates how crises can unfold and how they will affect it under various assumptions about how it reacts to the crisis. With such scenarios in hand, top management can then understand how crises can endanger the franchise of their institution and how to manage risks before they occur so that they can survive them. Such a scenario approach requires economic and financial analysis. It cannot be done by risk management departments populated by physicists and mathematicians. Such an approach also cannot be successful unless top management believes that the scenarios considered represent legitimate threats to the institution and that the institution has to protect itself against such threats.

Conclusion

Risk management has made considerable progress since 1998. The difficulties of the last year have convinced many observers that somehow there are deep flaws in risk management and that the problems of the last year are partly explained by risk management failures. In this paper, I show that one ought to distinguish carefully between risk-taking decisions that unexpectedly lead to losses and risk management assessments of risk. There are many ways that risk management failures can occur, but not every loss reflects a risk management failure. However, risk management practice can be improved by taking into account the lessons from financial crises.

These crises happen often enough that they have to be carefully modeled and institutions have to focus on scenario analyses that assess the implications of crises for their financial health and survival. Such scenario analyses cannot be built from quantitative models using past data, but instead they must use economic analysis to evaluate the impact of the withdrawal of liquidity and the feedback effects that are common in financial crises. To successfully impact firm strategy, such analyses have to be deeply rooted in a firm's culture and in the strategic thinking of top management.

Figure 1

