



Six Ways Companies Mismanage Risk

by René M. Stulz

Six Ways Companies Mismanage Risk

by René M. Stulz

As investors tot up their losses from the financial crisis, many will be asking themselves, How did Wall Street mess up so badly? What went wrong with all those complicated models? Even back in November 2007, before the crisis had really hit the stock markets, one commentator in the *Financial Times* wrote, “It is obvious there has been a massive failure of risk management across most of Wall Street.”

Of course, financial institutions can suffer spectacular losses even when their risk management is first-rate. They are, after all, in the business of taking risks. When risk management does fail, however, it is in one of six basic ways, nearly all of them exemplified in the current crisis. Sometimes the problem lies with the data or measures that risk managers rely on. Sometimes it relates to how they identify and communicate the risks a company is exposed to. Financial risk management is hard to get right in the best of times. In the following pages I’ll explore the six paths to failure in detail.

1. Relying on Historical Data

Risk-management modeling usually involves extrapolating from the past to forecast the probability that a given risk will materialize. Let’s assume you were a bank’s risk manager in 2006, and you were worried about the chances that real estate prices would plunge over the coming year. Your bank’s top executives needed to know how likely such a plunge was and what losses it would cause in order to decide how much exposure the bank should have to real estate prices.

You would have begun by examining the historical volatility of house prices and then calculating the one-year mean price change and its standard deviation, on the assumption that they would provide a good approximation going forward. The data might have led you to conclude that house price movements are random, like the outcome of a coin toss; that increases and decreases of the same size have the same likelihood of occurring; and that small changes are much more likely than larger ones. If all this were true, then the probability

distribution of house price changes would graphically map a bell curve around the mean change. The horizontal axis of the graph would measure price changes, and the vertical axis would measure the probability that each change would occur. The shape of the bell would be determined by volatility: The greater the volatility, the flatter and wider the curve.

Your modeling exercise could have gone wrong in two ways. First, if house price volatility was higher in the future than in the past (which proved to be the case), you would have substantially underestimated the probability of a plunge in prices. Your bell curve should have been flatter and wider, because the probabilities of larger price movements (up or down) were greater than you thought they would be. And if you had overestimated the mean change, your curve would be too far to the right, generating faulty estimates of the probabilities of all possible changes.

Second, you might have wrongly assumed that the distribution of future house price changes was described by the bell curve. But if changes in price were not normally distributed, then a graph of the distribution might be skewed—just as the probability distribution of a coin's falling heads up would be skewed if the coin were slightly bent.

Both kinds of error illustrate how difficult it is to use past data to predict the future. Suppose you had looked at 30 years' worth of house price changes to estimate your mean and standard deviation. Perhaps that time frame wasn't long enough—especially if the two decades preceding it were more volatile. A bell curve projected from 50 years' worth of data would have been flatter than the one you plotted, or perhaps less far to the right. Or suppose a 50-year analysis showed that large changes were more likely than with the normal distribution, but a 30-year analysis did not. In that case the 30-year analysis would understate the risk of a plunge in prices compared with the 50-year analysis, because with the latter the distribution would have fat tails relative to the normal distribution.

The difficulties with historical data intensify when you start thinking about the impact of a drop in house prices on the value of assets you hold. None of the data you had would have been of much use in estimating that impact, because they didn't cover a period during which the real estate market

faltered while a large number of subprime mortgages were outstanding. When many homeowners in a neighborhood have little home equity (almost always the case for subprime borrowers) and difficulty making their mortgage payments, the neighborhood enters a foreclosure spiral: Foreclosures lower prices, leading to more foreclosures because homeowners have negative equity. Your historical data could not have predicted that. Furthermore, if your institution had been holding complex securities that didn't exist in the period from which you drew your data, such as tranches of collateralized debt obligations (CDOs) backed in part by subprime mortgages, you couldn't have predicted the effect on their value of a drop in house prices.

Even if you could have captured the full impact on your balance sheet of a significant downturn in real estate, you would have been vulnerable to another kind of error in estimating the indirect effects. Your bank, like any other financial institution, would have had positions in many asset classes whose price changes correlated with changes in real estate. Price movements in all asset classes correlate with one another more or less, and the correlations have to be estimated correctly if you are to determine your bank's full exposure to a given risk. The higher the correlations, the riskier your total portfolio of assets and therefore the more equity capital you will need to protect your institution against losses.

Obviously, in estimating correlations from historical data you would have been vulnerable to all the problems we have just described. To complicate matters further, correlations are not constant; in fact, they are known to increase during periods of crisis. Thus, in calculating the overall riskiness of your bank's positions, you would also have had to model the probability distributions of all the correlations across your various portfolios as well as the price distributions in each asset class, further increasing the likelihood of error. All this illustrates the inconvenient truth that rapid financial innovation over recent decades has made history an imperfect guide.

2. Focusing on Narrow Measures

The data you use to create your models are only part of the problem. The actual measures with which you police your risk—especially in

René M. Stulz (stulz_1@fisher.osu.edu) is the Everett D. Reese Chair of Banking and Monetary Economics at The Ohio State University's Fisher College of Business in Columbus.

securities trading—can also lead you to ignore risks that you should be taking into account.

A daily Value-at-Risk (VaR) measure is the most common way of assessing the riskiness of securities trading in financial institutions. VaR essentially measures the maximum amount of money you might lose at a given probability level. For instance, a VaR of \$100 million at the 1% level means that you have only a 1% chance of losing more than that amount over the next day. A bank assigns an upper limit to the VaR it is willing to accept—say, \$125 million, in which case it would be comfortable with a 1% probability of losing \$100 million. Large banks typically declare on a quarterly basis the number of times in the previous quarter that the P&L showed a loss higher than the daily VaR. If a bank measures VaR at the 1% level, it should exceed its VaR roughly 1% of the time. In its annual report for 2006, the Swiss bank UBS stated that it had never had a loss that exceeded its daily VaR. In 2007, however, the bank exceeded its daily VaR 23 times—demonstrating that risk measurement had been unable to keep up with the dramatic changes in market conditions. If a bank uses VaR to protect itself from losses, it may have insufficient capital to support the risks it is taking.

There are several other problems with this approach. First, the simple fact that UBS experienced all those VaR overruns says little about the company's actual financial health. The overruns could have been small, and a rapid increase in volatility could have created many large gains as well. Alternatively, UBS could have realized many large losses and only a few large gains. In the former case it might actually be ahead at the end of the year; in the latter case it might be in serious trouble.

Second, VaR does not capture catastrophic losses that have a small probability of occurring. Consider a firm with a daily VaR of \$100 million at the 1% probability level. If this firm exceeded its VaR on only one day out of 100, while executing thousands of trades, it would appear to have an excellent risk-management record: It exceeded the VaR exactly 1% of the time. This would be cold comfort, however, if the amount lost above the VaR—which VaR is not designed to assess—threatened the institution. Though it is hard to get data on the dollar amounts of VaR

overruns, the markets have been full of rumors of extremely large ones in 2008.

Finally, a daily measure doesn't capture the risk of a portfolio when the firm is stuck with the portfolio for a much longer period. Daily VaR measures assume that assets can be sold quickly or hedged, so a firm can limit its losses within a day. But as we have seen in 2008 and as we saw in other crises (such as the one in 1998), a dramatic withdrawal of liquidity from the markets leaves firms exposed for weeks or months on positions they cannot easily unwind. The market for many CDOs backed in part by securitized subprime mortgages essentially vanished, so banks with such CDOs on their balance sheets couldn't unload them except at fire-sale prices. The risk of those CDOs during that period was unrelated to a daily VaR; rather, it lasted as long as the market for them was moribund. And risk increases over the time horizon for which it is computed.

3. Overlooking Knowable Risks

Let's suppose for a moment that you can accurately measure the risks you identify. Your next challenge is to make sure you have accounted for all the risks you know—or should know—you're exposed to. I've found that risk managers can all too easily overlook four kinds of risks: those outside the class of risks normally associated with particular units; those related to the hedging strategies used to manage risks already identified and assessed; those that arise when a market is dominated by one or two large institutions; and those that pertain to changes in normal trading behavior due to doubts about the value and liquidity of assets.

Risks outside the normal risk class. Risk managers often distinguish among market, credit, and operational risks, which they measure differently and in isolation. Companies that fail to assess risk firmwide do not go beyond these measures, effectively assuming that the three types of risk aren't correlated. But when you put risks into boxes, you've ignored the fact that business units strongly identified with a particular class of risk may be exposed to risks of other types that are associated with other units. Furthermore, you are ignoring risks not included in the definitions of the boxes.

This tendency is driven in part by regulatory considerations. The Basel II rules, for

Idea in Brief

- How did the current global financial crisis get so bad? In part through the failure of conventional risk-management approaches.
- Failure in financial risk management takes essentially six forms, most of which are exemplified in the current crisis. For example, risk assessments are typically based on historical data, such as changes in house prices over time. But rapid financial innovation, including securitized subprime mortgages, has made such data unreliable. And some risks are missed because they're hidden in overly complex reports that decision makers can't understand.
- To prepare for the next crisis, take a leaf from the disaster-management handbook: Use scenario analyses to understand the various ways a crisis might unfold—and plan how your company would respond to each.

example, have an extremely narrow definition of operational risk, so banks that follow that definition overlook certain strategic and business risks. The market risk attached to securitized subprime mortgages was closely connected to a very large business risk at many banks for which securitizing subprime mortgages was a significant source of income. So the banks lost out twice: The value of their portfolios of subprime mortgages, securitized or otherwise, fell, and they lost the fee income from securitizing subprime mortgages. Firms using the Basel II definition of operational risk would have largely ignored this relationship. Thus it is crucial to measure risk in ways that cut across organizational silos and include all the material risks to which a firm is exposed.

Risks incurred by hedging. In analyzing risk-management failures, I've often seen cases in which an organization completed a risk assessment and implemented steps to protect itself but then failed to assess all the risks of the instruments used for risk mitigation. Before Russia defaulted on its domestic debt in August 1998, many hedge funds

bought high-yielding Russian debt and then hedged it against default risk as well as exchange-rate risk. It was easy to believe that the resulting position was risk-free. To hedge the currency risk, however, the funds had to sell rubles forward against dollars. The banks that were willing to stand on the other side of those trades were often Russian. When Russia defaulted, many banks collapsed, and the hedge funds ended up realizing exchange-rate risk because their counterparties did not honor the hedges. Had the fund managers properly accounted for counterparty risk, they would have understood that their positions were still exposed to substantial risk in the event of a shock to the Russian banking system. But risk managers who focus on exchange-rate risk are typically not responsible for credit risk, so they would have ignored the counterparty risk. By the same token, credit risk managers may overlook risks associated with exchange-rate contracts because exchange-rate risk is not credit risk.

Market-concentration risks. Much of the finance theory underlying statistical risk models makes the critical assumption that markets

Idea in Practice

To manage risk effectively, you have to choose the right data and metrics and have a clear sense of how all the moving parts work together. Risk managers routinely make six fundamental mistakes:

1. Relying on historical data. Risk-management modeling involves extrapolating from the past, but rapid financial innovation in recent decades has made history an imperfect guide.

Example: Historical data were of little use in estimating the impact of the recent fall in house prices, because those data didn't cover a period during which the market saw a downturn while a large number of subprime mortgages were outstanding.

2. Focusing on narrow measures. Many financial institutions use daily measures to track risk. These underestimate a firm's exposure, because they assume that assets can be sold quickly, limiting the firm's losses within a day.

Example: Financial crises involve a dramatic withdrawal of liquidity from securities markets, leaving firms exposed for weeks or months on positions they cannot easily unwind.

3. Overlooking knowable risks. Risk managers simply overlook many types of risk and sometimes even create them.

Example: Investors in Russia tried to hedge the risk of a collapse in the ruble by taking currency positions with Russian banks. But they failed to recognize that a shock to the banking system would threaten those banks' ability to meet their commitments.

4. Overlooking concealed risks. People responsible for incurring risk often don't report it—sometimes deliberately, but often unintentionally. Organizations have a tendency to expand unreported risks.

Example: If traders receive a share of the profits they generate but do not have to defray the losses, they have an incentive to take risks, which is easier to do if the risks are unmonitored.

5. Failing to communicate. Risk-management systems will provide little protection if risk managers don't communicate clearly.

Example: The Swiss bank UBS attempted to explain its subprime and housing exposures in an overly complex way and to the wrong audience.

6. Not managing in real time. Risks can change sharply and quickly with daily fluctuations in the stock market.

Example: A manager holding a barrier call option who doesn't check the risk throughout the day may fail to put appropriate hedges in place.

are largely frictionless—the economist’s way of saying that they work smoothly and cheaply. In a frictionless market, financial institutions take prices as given rather than changing them during transactions. But this can lead risk managers to ignore risks arising from market frictions, which can be introduced when a single institution accounts for a very large chunk of a market’s transactions. For instance, it is well known that when the hedge fund Long-Term Capital Management (LTCM) collapsed, in 1998, it was holding extremely large positions in the index option market. During the crisis it was unable to change its positions because other players wanted substantial discounts in order to trade on the scale LTCM required. What’s more, if a big firm realizes losses that force it to sell assets at discounted prices, it may drag down prices for other institutions, making them a little less creditworthy because their assets are worth less. That, in turn, may create funding problems that limit everyone’s ability to trade, making the market less liquid and pushing prices down further. Additional complications can come from predatory trading: Other traders try to push prices down so that the big player is forced to unwind its positions for next to nothing.

Value-assumption risks. Another important source of market friction is doubt about the true value of traded assets when markets are illiquid. This causes market participants to stop taking prices as given because transactions are too infrequent to provide clear price signals. Traders often use pricing services to mark securities to market. Recently, pricing services were observed to differ by 20% or more on the same subprime securitized bonds, so participants did not know which price to rely on. This uncertainty played a role in the collapse of Bear Stearns, which at the time was financing about 25% of its assets through repurchase agreements, or repos. With a repo, a bank sells and agrees to repurchase securities, effectively borrowing money for the period between the two transactions—usually a day. This represents a major source of funding for many banks, especially for their trading operations. The risk associated with this type of funding is typically deemed to be low, because the borrower has to provide collateral whose value exceeds the cash received by a set amount. So the primary risk for the provider

of repo financing is that it may be left with the collateral (which it can sell to a third party) in the event that the borrower defaults. But when the financial crisis struck, Bear Stearns’s major repo counterparties—unsure of the value and the liquidity of the securities Bear Stearns was using as collateral—were unwilling to roll over its short-term repo agreements. Risk managers who work with models that implicitly or explicitly assume frictionless markets would have ignored these possibilities.

4. Overlooking Concealed Risks

The risk manager may do everything in his or her power to measure and capture risk and still end up failing because the people responsible for incurring risk simply don’t report it. This can be very dangerous, since unreported risks have a tendency to expand in financial organizations. Suppose a securities-trading desk in your bank has risks that are not fully reported and therefore not thoroughly monitored. The desk traders will almost certainly receive a significant share of any profits they generate, but they won’t have to defray any of their losses. They therefore have an incentive to assume risks, which is easier to do if those risks are unmonitored.

Obviously, risk is underreported when risk takers deliberately hide their risks, as appears to have happened at the French bank Société Générale in 2007. But not all unreported risks are deliberately concealed; some of them involve positions that use securities not yet established in the markets or positions held for short periods.

Senior management decisions, too, can lead to reporting failures. Take the case of Union Bank of Switzerland. In the second half of the 1990s the bank was establishing risk-management systems to aggregate risks across all its securities-trading operations. One group of traders who focused on equity derivatives were extremely successful—but their computers were different from those in the rest of the bank, and integrating their systems with the bank’s would have required that they change computers. The bank decided, at the top level, that it was more important to let the traders make money than to disrupt what they were doing in order to make the change. As a result, the risks this group assumed were not fully accounted for

Six Sources of Failure in Managing Risk

Lack of appropriate data

The rapid financial innovation of recent decades has made historical data less useful.

Narrow measures of risk

Traditional daily measures of risk can’t capture a company’s full exposure when market fundamentals are shifting.

Overlooked risks

Hedge funds that bought high-yielding Russian debt in the 1990s failed to properly account for counterparty risk.

Hidden risks

Unreported risks have a tendency to expand in financial institutions.

Poor communication

Complex and expensive risk-management systems can induce a false sense of security when their output is poorly communicated to top management.

Rate of change

The risk characteristics of securities may change too quickly to enable managers to properly assess and hedge risks.

in the risk-management system. Soon thereafter, the group lost a large amount of money, forcing the undercapitalized bank to merge with another Swiss bank to create UBS.

Clearly, organizations face trade-offs. Risk management might conceivably be structured to keep track of everything at all times—but it would probably be too costly to implement and, worse, would stifle innovation within the firm. In fast-moving markets, employees need to have flexibility in their trading. Often, the largest profits are made in the newest securities. Effective risk reporting ultimately depends on an institution's culture and incentives. If risk is everybody's business in an organization, pockets of it aren't likely to go unobserved. And employees will take risks more judiciously if their compensation is affected. The best risk models will be much less effective in a firm with poorly devised incentives than in one where incentives align with the firm's risk-taking objectives.

5. Failing to Communicate

Ironically, the risk manager who has most scrupulously modeled, measured, and captured knowable risk is perhaps most likely to trigger the fifth type of risk-management failure: poor communication to the board and the CEO, who are ultimately responsible for

making decisions about risk. If a firm has state-of-the-art risk-management systems but the board and the CEO don't understand them because the (technically very savvy) risk manager cannot properly explain the complex reports to nonexperts, the systems may do more harm than good by inspiring unwarranted confidence in their capabilities. Even worse, information may reach top management too late or be distorted by intermediaries.

Communication failures have certainly played a role in the most recent crisis. For example, the UBS report to its shareholders explains, "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." An industry commission that drew lessons from the crisis emphasized communication issues as well. It concluded that "risk monitoring and management reduces to the basics of getting the right information, at the right time, to the right people, such that those people can make the most informed judgments possible." Finally, a report from the Senior Supervisors Group (which includes top regulators from the United States, England, and Germany, among others) said, "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."

It is always tempting for a risk manager who makes a presentation to the board or the CEO to overstate the company's ability to measure risk. Risk-management systems are extremely costly, and a CEO may be nonplussed to learn that all that money pays for imprecise estimates. Paradoxically, by developing a risk culture that accepts and understands the limitations involved, a firm can increase the value of risk management.

6. Not Managing in Real Time

So far we have looked at risk management in terms of capturing a risk profile at a given point in time. But it is a dynamic process: Risk managers are responsible for making sure that the firm takes only the risks that it wants to take. As a result, they must constantly

Don't Be Afraid of the Unknown

Contrary to what many people may believe, you can manage risks without knowing exactly what they are—meaning that most of what you'd call unknown risks can in fact be captured in statistical risk-management models.

Think about how you measure stock price risk. Suppose you model price returns using the normal distribution and you have no reason to believe that future returns will come from a different distribution. As long as the historical volatility and mean are a good proxy for the future behavior of stock returns, you will capture the relevant risk characteristics of the stock through your estimation of the statistical distribution of its returns. You do not need to know why the stock return is +10% in one period and -15% in another.

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

The unknown risks that lead to risk-management failures are those that would have provoked materially different actions had the firm's managers known about them. A few of these will always exist, of course, and the only way to account for them is to make available some capital over and above what your models predict you need. But here you have to enter the realm of intuition; by definition, a formal risk-management system will provide little guidance.

monitor, hedge, and mitigate the firm's known risks.

This responsibility is more onerous for financial firms than for most others. Elsewhere risks change more slowly and usually involve a new exposure assumed through operations, such as sales or purchases denominated in foreign currencies. But financial firms have many derivatives positions and positions with embedded derivatives; the associated risks can change sharply even if the firms take no new positions. These changes can be dizzying in periods of turmoil. Figuring out the right hedge when markets are moving rapidly is like trying to change an insurance policy on a house while it is burning. In an extreme example, in just one day a security might have an exposure to a stock price such that it gains substantially if the price rises but later have an exposure such that it loses substantially if the price rises. Suppose you hold what is called a barrier call option—an option that pays off only if the stock price stays below a certain level. Its value will start to fall as the stock price approaches the threshold. For a product like this, hedges that are adjusted only daily could create large losses: A hedge that is optimal at the start of the day—say, a short position in the stock—might increase risk exposure at the end of the day if the stock price has risen.

When the risk characteristics of securities can change quickly, it is challenging for risk monitors to capture changes and for risk managers to adjust hedges accordingly. The challenge is especially great when risk characteristics can change dramatically as a result of small changes in the determinants of the securities' price.

The introduction of mark-to-market accounting actually makes it even harder for risk managers to estimate and adequately hedge risk. In a way, marking to market has brought what is known as the observer effect into financial markets: For large organizations, observing the value of a complex security affects the value of that security. As losses become

known through the process of marking to market, they start a chain reaction of adjustments at other institutions and affect the prices of possible trades as the market comes to better understand the capital positions of the institutions involved.

•••

As the foregoing makes clear, conventional approaches to risk management present many pitfalls. Even in the best of times, if you are to manage risk effectively, you must make extremely good judgment calls involving data and metrics, have a clear sense of how all the moving parts work together, and communicate that well. In the worst of times, risk management can fall apart. Historical models can fail, liquidity can dry up, and correlations can become stronger without warning. It's doubtful, therefore, that tinkering with existing systems will be sufficient to prevent future risk-management failures. Solutions from outside the traditional framework will be required. Don't focus only on investing more in order to better estimate and track risks. Instead, augment the models you have with scenario analyses of how a financial crisis might unfold depending on how your firm and other large companies react to the crisis. In other words, take a leaf from the disaster-management handbook. If you live in Florida or Louisiana, you shouldn't spend a lot of time thinking about how likely it is that you'll be hit by a hurricane. Rather, you should think about what would happen to your organization if it was hit by one and how you would deal with the situation. Instead of focusing on the fact that the probabilities of catastrophic risks are extremely small, risk managers should build scenarios for such risks, and the organization should design strategies for surviving them. One might call this "sustainable risk management."

Figuring out the right hedge when markets are moving rapidly is like trying to change an insurance policy on a house while it is burning.

Reprint R0903G

To order, see the next page

or call 800-988-0886 or 617-783-7500

or go to www.hbr.org

Further Reading

The Harvard Business Review Paperback Series

Here are the landmark ideas—both contemporary and classic—that have established *Harvard Business Review* as required reading for businesspeople around the globe. Each paperback includes eight of the leading articles on a particular business topic. The series includes over thirty titles, including the following best-sellers:

Harvard Business Review on Brand Management

Product no. 1445

Harvard Business Review on Change

Product no. 8842

Harvard Business Review on Leadership

Product no. 8834

Harvard Business Review on Managing People

Product no. 9075

Harvard Business Review on Measuring Corporate Performance

Product no. 8826

For a complete list of the *Harvard Business Review* paperback series, go to www.hbr.org.

Harvard Business Review

To Order

For *Harvard Business Review* reprints and subscriptions, call 800-988-0886 or 617-783-7500. Go to www.hbr.org

For customized and quantity orders of *Harvard Business Review* article reprints, call 617-783-7626, or e-mail customizations@hbsp.harvard.edu

Harvard Business Review

www.hbr.org

U.S. and Canada
800-988-0886
617-783-7500
617-783-7555 fax