

OVERVIEW—TWO DECADES OF RISK MANAGEMENT

1990–1999

Chapter 1	A Framework for Understanding Market Crisis	7
Chapter 2	Practical Issues in Choosing and Applying Risk Management Tools	25
Chapter 3	The Three P's of Total Risk Management	39
Chapter 4	Reporting and Monitoring Risk Exposure	61

2000–Present

Chapter 5	Risk Management: A Review	73
Chapter 6	Defining Risk	113
Chapter 7	Value and Risk: Beyond Betas	125
Chapter 8	A Simple Theory of the Financial Crisis; or, Why Fischer Black Still Matters	133
Chapter 9	Managing Firm Risk	139
Chapter 10	Risk Measurement versus Risk Management	153

<http://www.pbookshop.com>

A FRAMEWORK FOR UNDERSTANDING MARKET CRISIS

Richard M. Bookstaber

The key to truly effective risk management lies in the behavior of markets during times of crisis, when investment value is most at risk. Observing markets under stress teaches important lessons about the role and dynamics of markets and the implications for risk management.

No area of economics has the wealth of data that we enjoy in the field of finance. The normal procedure we apply when using these data is to throw away the outliers and focus on the bulk of the data that we assume will have the key information and relationships that we want to analyze. That is, if we have 10 years of daily data—2,500 data points—we might throw out 10 or 20 data points that are totally out of line (e.g., the crash of 1987, the problems in mid-January 1991 during the Gulf War) and use the rest to test our hypotheses about the markets.

If the objective is to understand the typical day-to-day workings of the market, this approach may be reasonable. But if the objective is to understand the risks, we would be making a grave mistake. Although we would get some good risk management information from the 2,490 data points, unfortunately, that information would result in a risk management approach that works almost all the time but does not work when it matters most. This situation has happened many times in the past: Correlations that looked good on a daily basis suddenly went wrong at exactly the time the market was in turmoil; value at risk (VAR) numbers that tracked fairly well day by day suddenly had no relationship to what was going on in the market. In the context of effective risk management, what we really should do is throw out the 2,490 data points and focus on the remaining 10 because they hold the key to the behavior of markets when investments are most at risk.

Reprinted from *AIMR Conference Proceedings: Risk Management: Principles and Practices* (August 1999): 7–19.

This presentation considers the nature of the market that surrounds those outlier points, the points of market crisis. It covers the sources of market crisis and uses three case studies—the equity market crash of 1987, the problems with the junk bond market in the early 1990s, and the recent problems with Long-Term Capital Management (LTCM)—to illustrate the nature of crisis and the lessons for risk management. This presentation also addresses several policy issues that could influence the future of risk management.

SOURCES OF CRISIS

The sources of market crisis lie in the nature and role of the market, which can be best understood by departing from the mainstream view of the market.

Market Efficiency

The mainstream academic view of financial markets rests on the foundation of the efficient market hypothesis. This hypothesis states that market prices reflect all information. That is, the current market price is the market's "best guess" of where the price should be. The guess may be wrong, but it will be unbiased; it is as likely to be too high as too low. In the efficient market paradigm, the role of the markets is to provide estimates of asset values for the economy to use for planning and capital allocation. Market participants have information from different sources, and the market provides a mechanism that combines the information to create the full information market price. Investors observe that price and can plan efficiently by knowing, from that price, all of the information and expectations of the market.

A corollary to the efficient market hypothesis is that, because all information is already embedded in the markets, no one can systematically make money trading without nonpublic information. If new public information comes into the market, the price will instantaneously move to its new fair level before anybody can make money on that new information. At any point in time, just by luck, some traders will be ahead in the game and some will be behind, but in the long run, the best strategy is simply to buy and hold the overall market.

I must confess that I never felt comfortable with the efficient market approach. As a graduate student who was yet to be fully indoctrinated into this paradigm, I could look at the many simple features of the market that did not seem to fit.

Why do intraday prices bounce around as much as they do? The price of a futures contract in the futures market or a stock in the stock market moves around much more than one would expect from new information coming in. What information could possibly cause the price instantaneously to jump two ticks, one tick, three ticks, two ticks second by second throughout the trading day?

How do we justify the enormous overhead of having a continuous market with real-time information? Can that overhead be justified simply on the basis of providing the marketplace with price information for planning purposes? In the efficient market context, what kind of planning would people be doing in which they had to check the market and instantly make a decision on the basis of a tick up or down in price?

Liquidity and Immediacy

All someone has to do is sit with a broker/dealer trader to see that more than information is moving prices. On any given day, the trader will receive orders from the derivative desk to

hedge a swap position, from the mortgage desk to hedge out mortgage exposure, and from clients who need to sell positions to meet liabilities. None of these orders will have anything to do with information; each one will have everything to do with a need for liquidity.

And the liquidity is manifest in the trader's own activities. If inventory grows too large and the trader feels overexposed, the trader will aggressively hedge or liquidate a portion of the position, and the trader will do so in a way that respects the liquidity constraints of the market. If the trader needs to sell 2,000 bond futures to reduce exposure, the trader does not say, "The market is efficient and competitive, and my actions are not based on any information about prices, so I will just put those contracts in the market and everybody will pay the fair price for them." If the trader puts 2,000 contracts into the market all at once, that offer obviously will affect the price, even though the trader does not have any new information. Indeed, the trade would affect the market price even if the market *knew* the trader was selling without any informational edge.

The principal reason for intraday price movement is the demand for liquidity. A trader is uncomfortable with the level of exposure and is willing to pay up to get someone to take the position. The more uncomfortable the trader is, the more the trader will pay. The trader has to pay up because someone else is getting saddled with the risk of the position—someone who most likely did not want to take on that position at the existing market price because otherwise, that person would have already gone into the market to get it.

This view of the market is a liquidity view rather than an informational view. In place of the conventional academic perspective of the role of the market, in which the market is efficient and exists solely for informational purposes, this view is that the role of the market is to provide immediacy for liquidity demanders. The globalization of markets and the widespread dissemination of real-time information have made liquidity demand all the more important. With more and more market information disseminated to a wider and wider set of market participants, less opportunity exists for trading based on an informational advantage, and the growth of market participants means there are more incidents of liquidity demand.

To provide this immediacy for liquidity demanders, market participants must exist who are liquidity suppliers. These liquidity suppliers must have free cash available, a healthy risk appetite, and risk management capabilities, and they must stand ready to buy and sell assets when a participant demands that a transaction be done immediately. By accepting the notion that markets exist to satisfy liquidity demand and liquidity supply, the framework is in place for understanding what causes market crises, which are the times when liquidity and immediacy matter most.

Liquidity Demanders

Liquidity demanders are demanders of immediacy: a broker/dealer who needs to hedge a bond purchase taken on from a client, a pension fund that needs to liquidate some stock position because it has liability outflow, a mutual fund that suddenly has some inflows of cash that it has to put into the index or the target fund, or a trader who has to liquidate because of margin requirements or because of being at an imposed limit or stop-loss level in the trading strategy. In all these cases, the defining characteristic is that time is more important than price. Although these participants may be somewhat price sensitive, they need to get the trade done immediately and are willing to pay to do so. A huge bond position can lose a lot more if the bondholder haggles about getting the right price rather than if the bondholder just pays up a few ticks to put the hedge on. Traders who have hit their risk limits do not have any choice; they are going to get out, and they are not in a good position to argue whether or

not the price is right or fair. One could think of liquidity demanders as the investors and the hedgers in the market.

Liquidity Suppliers

Liquidity suppliers meet the liquidity demand. Liquidity suppliers have a view of the market and take a position in the market when the price deviates from what they think the fair price should be. To liquidity suppliers, price matters much more than time. For example, they try to take a cash position or an inventory position that they have and wait for an opportunity in which the liquidity demander's need for liquidity creates a divergence in price. Liquidity suppliers then provide the liquidity at that price.

Liquidity suppliers include hedge funds and speculators. Many people have difficulty understanding why hedge funds and speculators exist and why they make money in an efficient market. Their work seems to be nothing more than a big gambling enterprise; none of them should consistently make money if markets are efficient. If they did have an informational advantage, it should erode over time, and judging by their operations, most speculators and traders do not have an informational advantage, especially in a world awash in information.

So, why do speculators and liquidity suppliers exist? What function do they provide? Why do, or should, they make money? The answer is that they provide a valuable economic function. They invest in their business by keeping capital readily available for investment and by applying their expertise in risk management and market judgment. They want to find the cases in which a differential exists in price versus value, and they provide the liquidity. In short, they take risk, use their talents, and absorb the opportunity cost of maintaining ready capital. For this functionality, they receive an economic return.

The risk of providing liquidity takes several forms. First, a trader cannot know for sure that a price discrepancy is the result of liquidity demand. The discrepancy could be caused by information or even manipulation. But suppose somebody waves a white flag and announces that they are trading strictly because of a liquidity need; they have no special information or view of the market and are willing to discount the price an extra point to get someone to take the position off their hands. The trader who buys the position still faces a risk, because no one can guarantee that between the time the trader takes on the position and the time it can be cleared out the price will not fall further. Many other liquidity-driven sellers may be lurking behind that one, or a surprise economic announcement might affect the market.

The liquidity supplier should expect to make money on the trade, because there is an opportunity cost in holding cash free for speculative opportunities. The compensation should also be a function of the volatility in the market; the more volatile the market, the higher the probability in any time period that prices will run away from the liquidity suppliers. In addition, their compensation should be a function of the liquidity of the market; the less liquid the market, the longer they will have to hold the position and thus the longer they will be subject to the volatility of the market.

Interaction of Liquidity Supply and Demand in a Market Crisis

A market behaves qualitatively differently in a market crisis than in “normal” times. This difference is not a matter of the market being “more jumpy” or of a lot more news suddenly flooding into the market. The difference is that the market reacts in a way that it does not in normal times. The core of this difference in behavior is that market prices become countereconomic.

The normal economic consequence of a decline in market prices is that fewer people have an incentive to sell and more people have an incentive to buy. In a market crisis, everything goes the wrong way. A falling price, instead of deterring people from selling, triggers a growing flood of selling, and instead of attracting buyers, a falling price drives potential buyers from the market (or, even worse, turns potential buyers into sellers). This outcome happens for a number of related reasons: Suppliers who were in early have already committed their capital; suppliers turn into demanders because they have pierced their stop-loss levels and must liquidate their holdings; and others find the cost of business too high with widening spreads, increased volatility, and reduced liquidity making the risk-return trade-offs of market participation undesirable. It is as if the market is struck with an autoimmune disease and is attacking its own system of self-regulation.

An example of this drying up of supply can be seen during volatility spikes. Almost every year in some major market, option volatilities go up to a level that no rational person would think sustainable. During the Asian crisis in 1998, equity market volatility in the United States, Hong Kong, and Germany more than doubled. During the exchange rate crisis in September 1993, currency volatility went up manyfold. During the oil crisis that accompanied the Gulf War, oil volatilities exceeded 80 percent. Volatilities for stocks went from the mid-teens to more than 100 percent in the crash of 1987. Did option traders really think stock prices would be at 100 percent volatility levels during the three months following the crash? Probably not. But the traders who normally would have been available to take the other side of a trade were out of the market. At the very time everybody needed the insurance that options provide and was willing to pay up for it, the people who could sell that insurance were out of the market. They had already “made their move,” risking their capital at much lower levels of volatility, and now were stopped out of their positions by management or, worse still, had lost their jobs.

Even those who still had their jobs kept their capital on the sidelines. Entering the market in the face of widespread destruction was considered imprudent, and the cost of entry was (and still is) fairly high. Information did not cause the dramatic price volatility. It was caused by the crisis-induced demand for liquidity at a time that liquidity suppliers were shrinking from the market.

Market Habitat

All investors and traders have a market habitat where they feel comfortable trading and committing their capital—where they know the market, have their contacts in the market, have a feel for liquidity, know how the risks are managed, and know where to look for information. The habitat may be determined by an individual’s risk preferences, knowledge, experience, time frame and institutional constraints, and by market liquidity. Investors will roam away from their habitat only if they believe incremental returns are available to them. Someone who is used to trading in technology stocks will need more time for evaluation and a better opportunity to take a position in, say, the automotive sector, than in the more familiar technology sector.

Nowadays, the preferred market habitat for most investors and traders is expanding because of low barriers to entry and easy access to information. Anyone can easily set up an account to trade in many markets, ranging from the G-7 countries to the emerging markets. Anyone can get information—often realtime information—on a wide variety of bonds and stocks that used to be available only to professionals. The days of needing to call a broker to check up on the price of a favorite stock now seem a distant memory.

More information and fewer barriers to entry expand habitat. Higher levels of risk also tend to expand habitat. The distinction among assets blurs as risk increases. In addition, market participants become more like one another, which means that liquidity demanders all demand pretty much the same assets and grab whatever sources of liquidity are available. This situation is characterized in the market as “contagion,” but in my view, what is happening is an expansion of habitat because the risk of the market has made every risky asset look pretty much the same. If all investors are in the same markets, they will run into trouble at the same time and will start liquidating the same markets to get financing and reduce their risks.

Think of how the investor’s focus shifts as the investor moves from a normal market environment to a fairly energetic market environment, and then to a crash environment. In a normal market, investors have time to worry about the little things: the earnings of this company versus that company, P/Es, dividends, future prospects, and who is managing what. As the energy level goes up in the market, investors no longer have the luxury of considering the subtleties of this particular stock or that stock. They need to concentrate on sectors. If the technology sector is underperforming, all technology stocks look the same. If oil prices go up, an oil company’s management and earnings prospects no longer matter; all that matters is that the company is in the energy sector. Turn the heat up further to a crash environment and all that participants care about is that it is a stock and that they can sell it. All stocks look the same, and the correlations get close to 1.0 because the only characteristic that matters is that this asset is a stock or, for that matter, is risky. In fact, the situation can get even worse; junk bonds may be viewed to be similar enough to stocks that they trade like stocks. The analysis and market history of the normal market environment no longer applies. The environment is different; the habitat has changed.

An analogy from high-energy physics helps to illustrate the situation. As energy increases, the constituents of matter blur. At low energy levels—room temperature—molecules and atoms are distinct and differentiated. As energy goes up, the molecules break apart and what is left are the basic building blocks of matter, the elements. As energy goes up even more, the atoms break apart and plasma is left. Everything is a defused blob of matter.

As the energy of the market increases, the same transformation happens to the constituents of the market. In a market crisis, all the distinct elements of the market—the stocks (e.g., IBM and Intel), the market sectors (e.g., technology and transportation), the assets (e.g., corporate bonds and swap spreads)—turn into an undifferentiated plasma. Just as in high-energy physics, where all matter becomes an undifferentiated “soup,” in the high-energy state of a market crisis, all assets blur into undifferentiated risk.

One of the most troubling aspects of a market crisis is that diversification strategies fail. Assets that are uncorrelated suddenly become highly correlated, and all the positions go down together. The reason for the lack of diversification is that in a high-energy market, all assets in fact *are* the same. The factors that differentiate them in normal times are no longer relevant. What matters is no longer the economic or financial relationship between assets but the degree to which they share habitat. What matters is who holds the assets. If mortgage derivatives are held by the same traders as Japanese swaps, these two types of unrelated assets will become highly correlated because a loss in the one asset will force the traders to liquidate the other. What is most disturbing about this situation is not that the careful formulation of an optimized, risk-minimizing portfolio turns to naught but that there is no way to determine which assets will be correlated with which other assets during a market crisis. That is, not only will diversification fail to work at the very time it is most critical, but determining the way in which it will fail will be impossible.

Liquidity demanders use price to attract liquidity suppliers, which sometimes works and sometimes does not. In a high-risk or crisis market, the drop in prices actually reduces supply and increases demand. This is the critical point that participants must look for. Unfortunately, most people never know how thin the ice is until it breaks. Most people did not see any indications in the market in early October 1987 or early August 1998 that made them think they were on thin ice and that a little more weight would dislocate the market and prices would become an adverse signal. Of course, the indications seem obvious after the fact, but it should suggest something about the complexity of the market that these indications are missed until it is too late. For example, option prices, particularly put option prices, were rising before the crash of 1987. After the crash, this phenomenon was pointed to as an indicator that there was more risk inherent in the market and more demand for protection. In the month or so before Long-Term Capital Management (LTCM) had its problems, the U.S. swap spread was at its lowest volatility level in a decade. This low volatility demonstrated a lack of liquidity and commitment to the swap market. In the case of the 1987 market crash, the missed indicator was high volatility; in the case of the LTCM crisis, the missed indicator was low volatility.

CASE STUDIES

Three case studies help to demonstrate the nature of market crises: the equity market crash of 1987, the junk bond crisis, and the LTCM default.

1987 Equity Market Crash

The market crash of 1987 occurred on Monday, October 19. But it was set up by the smaller drop of Friday, October 16 and by the reaction to that drop from a new and popular strategy—portfolio insurance hedging.

Portfolio insurance is a strategy in which a manager overlays a dynamic hedge on top of the investment portfolio in order to replicate a put option. Operationally, the hedge is reduced as the portfolio increases in value and increased as the portfolio declines in value. The hedge provides a floor to the portfolio, because as the portfolio value drops beyond a prespecified level, the hedge increases to the point of offsetting future portfolio declines one for one. The selling point for portfolio insurance is that it provides this floor protection while retaining upside potential by systematically reducing the hedge as the portfolio rises above the floor.

This hedging strategy is not without a cost. Because the hedge is being reduced as the portfolio rises and increased as the portfolio drops, the strategy essentially requires buying on the way up and selling on the way down. The result is a slippage or friction cost because the buying and selling happen in reaction to the price moves; that is, they occur slightly after the fact. The cumulative cost of this slippage can be computed mathematically using the tools of option-pricing theory; the cumulative cost of the slippage should be about the same as the cost of a put option with an exercise price equal to the hedge floor.

The key requirement for a successful hedge, and especially a successful dynamic hedge, is liquidity. If the hedge cannot be put on and taken off, then obviously all bets are off. Although liquidity is not much of a concern if the portfolio is small and the manager is the only one hedging with a particular objective, it becomes a potential nightmare when

everyone in the market has the same objective, which in a nutshell is what happened on October 19.

On Monday morning October 19, everybody who was running a portfolio insurance program looked at the computer runs from Friday's market decline and saw they had to increase their hedges. They had to short out more of the exposure that they had to the market, and the hedging instrument of choice was the S&P 500 Index futures contract. Shortly after the open on October 19, the hedges hit the S&P pit.

Time mattered and price did not; once their programs were triggered, the hedge had to be increased and an order was placed at the market price. And a lot of programs were triggered. Portfolio insurance was first introduced by LOR (Leland O'Brien Rubinstein) in 1984, and portfolio insurance programs were heavily and successfully marketed to pension funds, which overlaid tens of billions of dollars of equity assets.

The traders in the S&P pit are very fast at execution. When someone wants to sell a position at the market, a trader in the pit will buy it immediately. Once the market maker takes the position, the market maker will want to take the first opportunity to get rid of it. The market makers on the floor make money on the bid–offer spread (on turnover) and not by holding speculative positions. Among the sources they rely on to unload their inventory are program traders and cash futures arbitrageurs. The program traders and arbitrageurs buy S&P contracts from the futures pit while selling the individual stocks that comprise the S&P 500 on the NYSE. If the price of the basket of stocks differs from the price of the futures by more than the transaction costs of doing this trade, then they make a profit. This trade effectively transfers the stock market activities of the futures pit to the individual stocks on the NYSE. It is here where things broke down in 1987, and they broke down for a simple reason: Although the cash futures arbitrageurs, program traders, and market makers in the pit are all very quick on the trigger, the specialists and equity investors who frequent the NYSE are not so nimble.

The problem might be called “time disintermediation.” That is, the time frame for being able to do transactions is substantially different between the futures market and the equity market. This situation is best understood with a stylized example. Suppose that you are the specialist on the NYSE floor for IBM. On Monday morning October 19, you wait for the markets to open. Suddenly, a flood of sell orders comes in from the program traders. You do not have infinite capital. Your job is simply to make the market. So, you drop the price of IBM half a point and wait. Not many people are coming, so you drop it a full point, figuring now people will come.

Meanwhile, suppose I am an investment manager in Boston who is bullish on IBM, and I am planning to add more IBM to my portfolio. I come in, glance at the screen, and see that IBM is down a half point. After coming back from getting some coffee, I check again; IBM is now down a full point. The price of IBM looks pretty good, but I have to run to my morning meeting.

Half an hour has gone by, and you and the other specialists are getting worried. A flood of sell orders is still coming in, and nowhere near enough buyers are coming in to take them off of your hands. Price is your only tool, so you drop IBM another point and then two more points to try to dredge up some buying interest.

By the time I come back to my office, I notice IBM is down four points. If IBM had been down a half point or a full point, I would have put an order in, but at four points, I start to wonder what is going on with IBM—and the market generally. I decide to wait until I can convene the investment committee that afternoon to assess the situation.

The afternoon is fine for me, but for you, more shares are piling into your inventory with every passing minute. Other specialists are faced with the same onslaught, and prices are falling all around you. You now must not only elicit buyers, but you must also compete with other stocks for the buyers' capital. You drop the offer price down 10 points from the open. The result is a disaster. The potential liquidity suppliers and investment buyers are being scared off by the higher volatility and wider spreads. And, more importantly, the drop in price is actually inducing more liquidity-based selling as the portfolio insurance programs trigger again and again to increase their selling to add to their hedges. So, because of time disintermediation and the specialist not having sufficient capital, the price of IBM is dropped too quickly, the suppliers are scared off, and the portfolio insurance hedgers demand even more liquidity than they would have otherwise.

This IBM example basically shows what happened in the crash of 1987. Demand for liquidity moved beyond ignoring price and focusing on immediacy to actually increasing as a function of the drop in price because of the built-in portfolio insurance rules. Supply dried up because of the difference in time frames between the demanders and suppliers, which led prices to move so precipitously that the suppliers took the drop as a negative signal. The key culprit was the difference in the trading time frames between the demanders and the suppliers. If the sellers could have waited longer for the liquidity they demanded, the buyers would have had time to react and the market would have cleared at a higher price.

1991 Junk Bond Debacle

Junk bonds, or more euphemistically high-yield bonds, were the mainstay of many corporate finance strategies that developed in the 1980s. The best known use of high-yield bonds was in leveraged buyouts (LBOs) and hostile takeovers. Both of these strategies followed the same course over the 1980s. They started as good ideas that were selectively applied in the most promising of situations. But over time, more and more questionable deals chased after the prospect of huge returns, and judgment was replaced with avarice. The investment banks played more the role of cheerleader than advisor, because they stood to gain no matter what the long-term outcome and they had a growing brood of investment banking mouths and egos to feed.

The size of the average LBO transaction peaked in 1987. But deal makers continued working to maintain their historical volumes even as the universe of leverageable companies declined. Volume was maintained in part by lowering the credit quality threshold of LBO candidates. The failed buyout of United Airlines in 1989 is one example of this situation, because airlines are cyclical and previously had not been considered good candidates for a highly levered capital structure. Leverage in the LBOs also increased over the course of the 1980s. Cash flow multiples increased in 1987 and 1988, from the 5× range in 1984 and 1985 to the 10× range in 1987 and 1988. This increase turned out to be fatal for many companies. An earnings shortfall that is manageable at 5 times cash flow can lead to default if the investors pay 10 times cash flow.

Although LBOs moved from larger to smaller deals, hostile takeovers went after bigger game as time went on. The RJR debt of nearly \$10 billion represented approximately 5 percent of the high-yield market's total debt outstanding. Many institutions had limitations on the total amount of exposure they could have to any one name, which became a constraint given the size of the RJR issues.

The justification for hostile takeovers was, starting in the mid-1970s, for the market value of companies to be less than their replacement cost. Thus, after a hostile takeover, the acquirer

could sell off the assets and inventories for more than the cost of buying the company holding those assets. The activity of hostile takeovers—and possibly the threat of further takeovers—woke up the market to the disparity between the market value and the replacement cost of companies' assets, and the gap closed by 1990. The arbitrage plays implicit in hostile takeovers led to an improvement of market efficiency in textbook fashion, and the *raison d'être* for the hostile takeovers disappeared. But the hope for financial killings remained and led to continued demand for the leverage of high-yield bonds as ammunition to bag the prey.

The following scenario summarizes the life cycle of LBOs and hostile takeovers. With these financial strategies still virgin territory, and with the first practitioners of the strategies the most talented and creative, the profits from the first wave of LBOs and hostile takeovers made headlines. More investors and investment bankers entered into the market, and credit quality and potential profitability were stretched in the face of the high demand for high-yield financing. Rising multiples were paid for LBOs and were accepted in hostile takeovers because of both the higher demand for financing and the increase in equity prices. The result of the stretching into lower-quality deals and the higher multiples paid for the companies led to more defaults.

The defaults hit the market even harder than did the earlier LBO and hostile takeover profits. Within a few short months, high-yield bonds were branded as an imprudent asset class. In 1991, the high-yield bond market was laid to waste. Bond spreads widened fourfold, and prices plummeted. The impact of the price drop was all the more dramatic because, even though the bonds were not investment grade, investors had some expectation of price stability. The impact on the market was the same as having the U.S. stock market drop by 70 percent. As with the 1987 stock market crash, the junk bond debacle was not the result of information but of a shift in liquidity.

In 1991, the California Insurance Commission seized Executive Life. The reaction to this seizure was many faceted, and each facet spelled disaster for the health of the market. Insurance companies that had not participated in the high-yield bond market lobbied for stricter constraints on high-yield bond holdings. It is difficult to know whether this action was done in the interest of securing the industry's reputation, avoiding liability for the losses of competitors through guaranty funds, stemming further failures (such as Executive Life), or meeting the threat of further insurance regulation. Insurance companies were anxious to stand out from their competitors in their holdings of high-yield bonds and featured their minimal holdings of junk bonds as a competitive marketing point.

A number of savings and loans (S&Ls) seized on the high-yield market as a source of credit disintermediation. Federal deposit guarantees converted their high-risk portfolios into portfolios that were essentially risk free. The S&L investors captured the spread between the bond returns and the risk-free return provided to the depositors. That this situation was a credit arbitrage at the government's expense became clear in the late 1980s. The government responded with the Financial Institutions Reform, Recovery and Enforcement Act in 1989. This act not only barred S&Ls from further purchases of high-yield bonds, but it also required them to liquidate their high-yield bond portfolios over the course of five years. The prospect of the new regulation and stiffening of capital requirements by the Federal Home Loan Bank Board led S&Ls to reduce their holdings even in early 1989 by 8 percent, compared with an increase in holdings in the previous quarter of 10 percent.

Investors reacted quickly to the weakness in the high-yield bond market. In July 1989, high-yield bond returns started to decline, hitting negative returns. For investors who did not understand the risk of high-yield bonds, the realization of negative returns must have been a rude wake-up call. Over the third quarter of 1989, the net asset value of high-yield

mutual funds declined by as much as 10 percent. The implications of erosion of principal—coupled with media reports of the defaults looming in the high-yield market—led to widespread selling.

As with any other financial market, the junk bond market had both liquidity suppliers and liquidity demanders. Some poor-quality junk bonds made it to the market, which caused some investors who normally would have been suppliers of liquidity to spurn that market because it was considered imprudent. Consequently, financing was reduced. These people then had financial problems, which demonstrated that junk bonds were imprudent and which meant more people went out of the market. So, the liquidity suppliers who were willing to take on the bonds became liquidity demanders. They wanted to get rid of their junk bonds, and the more the price dropped, the more they wanted to get rid of their junk bonds. Junk bonds were less than 5 percent of their portfolios, so owning junk bonds was not going to ruin the entire portfolio, but they could have lost their jobs. Suddenly, suppliers were disappearing and turning into demanders. The price drop created the wrong signal; it made the bonds look worse than they actually were.

The junk bond crash of 1991 was precipitated by several junk-bond-related defaults. But the extent of the catastrophe was from liquidity, not default. Institutional and regulatory pressure accentuated the need for many junk bond holders to sell, and to sell at any price. Because the usual liquidity suppliers were in the position of now needing to sell, not enough capital was in the market to absorb the flow. The resulting drop in bond prices, rather than drawing more buyers into the market, actually increased the selling pressure, because the lower prices provided confirmation that high-yield bonds were an imprudent asset class. Regulatory pressure and senior management concerns—not to mention losses on existing bond positions—vetoed what many traders saw as a unique buying opportunity.

1998 LTCM Default

Long-Term Capital Management is a relative-value trading firm. Relative-value trading looks at every security as a set of factors and finds within that set of factors some factor that is mispriced between one security and another. The manager then tries to hedge out all the other factors of exposure so that all that is left is long exposure in the factor in one security and short exposure in the factor in another security. One security is cheaper than the other, so the manager makes money. Ideally, in relative-value trading, the positions should be self-financing so that the manager can wait as long as necessary for the two prices to converge. If a spread takes, say, three years to converge, that is no problem if the position is self-financed.

The most common relative-value trading is spread trading. Spread trading is attractive because all that matters is the relative value between the two instruments. This approach has great advantages for analytically based trading because it is easier to determine if one instrument is mispriced relative to another instrument than it is to determine if an instrument is correctly priced in absolute terms. A relative-value trader can still get it right even with making an erroneous assumption, so long as that assumption affects both instruments similarly. Another advantage of relative-value trading is that a relative-value trade is immune to some of the most unpredictable features of the market. If a macroeconomic shock hits the market, it will affect similar instruments in a similar way. Although both instruments might drop in price, the relative value of the two may remain unaffected.

One of the problems of relative-value trading, and of working with spread trades in particular, occurs because the spreads between instruments are typically very small. These small spreads are a direct result of trading between two very similar instruments, where the variations

between the prices are very small. Although in the end the dollar risk may be the same as an outright trade to put on this risk—and thereby get double-digit expected returns—the relative-value trader is usually highly leveraged.

Relative-value trading has other problems as well. First, these very big positions are hard to liquidate, and the newer, less-liquid markets are usually the very markets that exhibit the spread discrepancies. Yet these are the very markets where experience is limited and observers have not seen the risks played out over and over. Second, in a relative-value trade, the manager requires price convergence between the two assets in a spread position. Sooner or later that convergence should take place, but the manager does not know when and thus may have a long holding period. Third, because of the myriad risks and small spreads, the modeling in relative-value trading has to be very precise; if a manager has \$10 billion long in one instrument and \$10 billion short in another instrument and if the manager is off by 1 percent, then the manager stands to lose a lot of money.

In terms of relative-value trading at LTCM, the traders were doing such things as buying LIBOR against Treasuries, so they were short credit risk. They were buying emerging market bonds versus Brady bonds and mortgages versus Treasuries. While they had the trades on, they decided to reduce their capital. In the early part of 1998, LTCM returned nearly \$3 billion of capital to its investors, reducing its capital base from about \$7 billion to a little more than \$3 billion.

Normally, LIBOR, Treasuries, and mortgages—the markets that LTCM invested in—are very liquid. The liquidity that the traders at LTCM had, however, was lower than what they expected for several reasons, some completely unanticipated. Even in a normal market environment, if a trader is dealing with really large size, the market is not very liquid; if the trader starts to sell, nobody wants to buy because they know there is a lot more supply where that came from. LTCM's real problems, however, started on July 7, 1998. On that day, the *New York Times* ran a story that Salomon Smith Barney was closing its U.S. fixed-income proprietary trading unit. Even though I was the head of risk management at Salomon, I did not know this decision had been made. I certainly questioned the move after the fact on several grounds; the proprietary trading area at Salomon was responsible for virtually all the retained earnings of Salomon during the previous five years. Furthermore, this was an announcement that no trader would ever want made public. Closing the trading unit meant that Salomon's inventory would probably be thrown into the market. If Salomon was closing its proprietary trading area in the United States, it probably would do so in London as well. So, the logical assumption was that Salomon's London inventory would be coming into the market as well. The result was that nobody would take the other side of that market; who wants to buy the first \$100 million of \$10 billion of inventory knowing another \$9.9 billion will follow? Salomon should have quietly reduced its risk and exposure. Once the risk and exposure were down and inventory was low, then Salomon could have announced whatever it wanted. As it was, the nature of the announcement worked to dampen demand in the market, which did not bode well for LTCM.

Another event that was not favorable for LTCM occurred in August 1998; Russia started to have problems. LTCM, like everybody else, had exposure to Russia. The result was that LTCM had to liquidate assets because its cash reserve was gone. Liquidating assets is only a big deal when nobody wants the assets. Not only did nobody want the assets because of the glut of inventory resulting from the closing of Salomon's proprietary trading units; they now did not want the assets because they knew LTCM was selling because it had financial problems and because they did not know how deep LTCM's inventory was. At the time LTCM was demanding immediacy, liquidity suppliers did not exist in the market.

To make matters worse, LTCM was itself a major liquidity supplier in the market. LTCM was providing the other side of the market for people who wanted to hedge out their credit exposure in various instruments. The reason LTCM was making money was that it was supplying liquidity. It was providing a side of the market that people needed. Once LTCM was gone, not many other people were left. And those who were left were not going to stay in the face of this huge overhang of supply. So, when LTCM had to sell, a market did not exist for its positions, because LTCM was the market. LTCM's selling drove the price down enough so that, just as in the case of portfolio insurance, LTCM had to sell even more. LTCM did manage to sell some of its positions but at such low prices that when it marked to market its remaining holdings, they dropped so much as to require even more margin and to require even more selling. So, a cycle developed, and as the spreads widened, anybody who would have provided liquidity on the other side was not willing to.

If people had had more time, the downward cycle would have been halted; someone would have taken the assets off LTCM's hands because the assets were unbelievably mis-priced, not only in terms of price levels but also in totally different directions. How could fixed-income instruments in Germany have almost historically low volatility while LIBOR instruments in the United Kingdom have historically wide spreads? The issue was strictly one of liquidity and immediacy; buyers simply were not there quickly enough.

Many things have been written about LTCM, some of which are not very favorable to the principals of the firm. But the fact is that the principals are among the brightest people in finance. They have done relative-value trading longer than anybody else on Wall Street. The failure of LTCM says more about the inherent risk and complexity of the market than it does about LTCM; the market is sufficiently complex that even the smartest and most experienced can fail. Who would have anticipated a closing of U.S. fixed-income proprietary trading at Salomon? Who would have anticipated that this closing would be revealed in a public announcement? Who would have anticipated the speed and severity of the Russian debacle hard on the heels of the Salomon announcement? It is that very complexity that the risk analysis models failed to capture.

LESSONS LEARNED

These market crises share some common elements that can teach all of us important lessons about risk management.

First, it is not just capital that matters. What matters is the willingness to put that capital into the market, to commit capital at times of crisis and high risk. During the LTCM crisis, if somebody had been willing to commit capital at a time when the spreads were at unbelievably wide levels, the crisis would have been averted. I was in charge of risk management at Salomon Smith Barney at the time of this crisis and encouraged—unsuccessfully, it turned out—a more aggressive position in the market. Salomon Smith Barney was in a position to stay in these spread trades, because the firm had sizeable capital and, through its proprietary trading group, more expertise on staff than anybody else in the world. (Remember that LTCM was dubbed “Salomon North” because the bulk of its talent came from Salomon, but Salomon retained an exceptional talent for relative-value trading even after John Meriwether and others left the firm.) Nevertheless, in spite of its far stronger capital position and its trading expertise, Salomon Smith Barney was just as quick to get out of the market as LTCM. So, what matters is not just capital or expertise. What matters is capital and expertise and the willingness to use that capital at the time the market really needs liquidity.

Second, speculative capital is needed to meet liquidity demand. Either the markets must slow down to allow people more time to respond to the demand for immediacy, or more participants must enter the markets who can act quickly and meet that immediacy. In the crash of 1987, circuit breakers would have slowed things down so that the portfolio insurance programs could have triggered at a pace that the traders in New York and elsewhere could have matched. Or on the futures side, more speculators with capital could have made the market and held onto those positions. Or on the stock exchange side, specialists with more capital and staying power could have held onto the inventory until the stock investors had gotten settled for the day.

Third, the markets must have differentiated participation. As the financial markets become more integrated, there is increasing focus on systemic risk—the risk inherent in the institutions that comprise the financial system. A nondifferentiated ecosystem has a lot of systemic risk. One little thing goes wrong, and everything dies. Complexity and differentiation are valuable because if one little thing goes wrong, other things can make up for it. Systemic risk has its roots in the lack of differentiation among market participants. Modern portfolio theory focuses on the concept of diversification within a portfolio, which is fine in a low-energy market. As a market moves to a high-energy state and habitats expand, what matters is not so much diversification among asset classes but diversification among market participants.

If everything I hold is also held by other market participants, all of whom have the same sort of portfolio and risk preferences that I have, I am not diversified. In a low-energy state, this lack of diversification will not be apparent, because prices will be dictated by macroeconomics and firm performance. As the market moves to a high-energy state, things change. What matters then is which assets look like which other assets based on the liquidity demanders and suppliers who will be dumping assets into the market. So, in a low-energy state, I am well diversified, but in a high-energy state, everything goes against me because what matters now is not what the assets are but the fact that they are pure risk and that they are all held by the same sort of people.

Finally, Wall Street has experienced a lot of consolidation—Citigroup and Morgan Stanley Dean Witter, for example. Big firms are sensitive to institutional and political pressure; they have to go through many checks and sign-offs and thus are slow to react. The habitat is becoming less diverse, and more systemic failures are occurring because everybody looks the same and is holding the same assets. Big firms never seem to be as risk taking as their smaller counterparts. When two firms merge, the trading floor does not become twice as large. The trading floor stays about the same size as it was before the two firms merged. The total risk-taking capability, however, is about half of what it was before. In fact, the situation gets even worse because two firms do not merge into one big firm in order to become a hedge fund. Firms merge in order to conduct retail, high-franchise business. Risk taking becomes less important, even somewhat of an annoyance. Although with consolidation the firm has more capital and more capability to take risk, it is less willing to take risk.

POLICY ISSUES

The markets are changing, and thus, risk management must change along with them. But often, changes resulting from reactions to market crises create more problems than they solve. Policy issues surrounding transparency, regulation, and consolidation could dramatically affect the future of risk management.

Transparency

The members of the LTCM bank consortium (the creditors of LTCM that took over the firm in September 1998) complained that they were caught unaware by the huge leverage of the hedge fund. Reacting to the losses and embarrassment they faced from the collapse, some of the consortium members entered the vanguard for increased transparency in the market. They argued that the only way to know if another LTCM is lurking is by knowing their trading clients' positions.

The issue of hedge fund transparency may deserve a fuller hearing, but opaqueness was not the culprit for LTCM. A simple back-of-the-envelope calculation would have been sufficient to demonstrate to the creditors that they were dealing with a very highly leveraged hedge fund. The banks—and everyone else in the professional investment community—knew that LTCM's bread and butter trading was swap spreads and mortgage spreads. Everyone also knew that on a typical day, these spreads move by just a few basis points—a few one-hundredths of a percent. Yet historically, LTCM generated returns for its investors on these trades of 30 percent or more. The only way to get from 5 or 10 basis points to 30 or 40 percent is to lever more than 100 to 1.

If the banks were unable to do this simple calculation, it is hard to see how handing over reams of trading data would have brought them to the same conclusion. Often in trading and risk management, it is not lack of information that matters; it is lack of perceiving and acting on that information. Indeed, looking back at the major crises at financial institutions—whether at Barings Securities, Kidder, Peabody & Co., LTCM, or UBS—finding even one case in which transparency would have made a difference is hard. The information was there for those who were responsible to monitor it. The problem was that they either failed to look at the information, failed to ask the right questions, or ignored the answers.

Indeed, if anything, the LTCM crisis teaches us that trading firms have good reasons for being opaque. Obviously, broadcasting positions dissipates potential profit because others try to mirror the positions of successful firms, but it also reduces market liquidity. If others learn about the positions and take them on, fewer participants will be in the market ready to take the opposite position. Also, if others know the size of a position and observe the start of liquidation, they will all stand on the sidelines; no one will want to take on the position when they think a flood of further liquidation is about to take place. Transparency will come at the cost of less liquidity, and it is low liquidity that is at the root of market crisis.

Regulation

Regulation is reactive. It addresses problems that have been laid bare but does not consider the structure that makes sense for the risks that have yet to occur. And indeed, by creating further rules and reporting requirements to react to the ever-increasing set of risks that do become manifest, regulation may actually become counterproductive by obscuring the field of view for financial institutions to the areas of risk that have yet to be identified. At some point, the very complexity of the risk management system gets in its own way and actually causes more problems than it prevents. We are not at that point yet in the financial markets, but some precedence exists for this phenomenon in other highly regulated industries, such as airlines and nuclear energy.

The thing to remember is that every new risk management measure and report required by regulation is not only one more report that takes limited resources away from other, less well-defined risk management issues; it is also one more report that makes risk managers more complacent in thinking they are covering all the bases.

Consolidation

I have already discussed the implications of consolidation on risk taking. With every financial consolidation, the capacity of the market to take risk is reduced. Large financial supermarkets and conglomerates are created to build franchise, not to enhance risk taking.

Consolidation also increases the risk of the market, especially the risk of market crisis. The increase in risk occurs because the market becomes less differentiated. A greater likelihood exists that everyone will be in the same markets at the same time and will share the same portfolios. The investment habitat becomes less diverse.

The drop in habitat diversity from financial consolidation looks a lot like the drop in retail diversity that has occurred as interstate highways and mass media have put a mall in every town and the same stores in every mall. Whether in food, clothing, or home furnishings, regional distinctions are disappearing. “The malling of America” is creating a single, uniform retail habitat.

Coming soon will be “the malling of Wall Street.” Broker/dealers are consolidating into a small set of investment “super stores.” On the investor side, more and more investors are taking advantage of ready access to information and markets, but along with this information advantage comes a convergence of views among investors—particularly the retail or individual investors—because the information sources are all the same.

When the Glass–Steagall Act was passed, in all likelihood Congress did not have in mind diversifying the ecosystem of the financial markets. Glass–Steagall created a separation between different types of financial institutions in order to protect investors. The separation and resistance to certain types of consolidation is still needed but now for another reason—to maintain a diverse habitat. The goal of any Glass–Steagall-type reform should be to maintain different types of risk takers. It should encourage differentiation among financial market participants so that if one liquidity supplier is not supplying liquidity in a particular adverse circumstance, another one is, thus helping to prevent or minimize a full-blown crisis.

Some people think of speculative traders as gamblers; they earn too much money and provide no economic value. But to avoid crises, markets must have liquidity suppliers who react quickly, who take contrarian positions when doing so seems imprudent, who search out unoccupied habitats and populate those habitats to provide the diversity that is necessary, and who focus on risk taking and risk management. By having and fostering this differentiated role of risk taking, market participants will find that crises will be less frequent and less severe, with less onerous consequences for risk management systems. The hedge funds, speculative traders, and market makers provide this role.

QUESTION AND ANSWER SESSION

Question: Could you discuss the U.S. Federal Reserve’s role in the LTCM crisis?

Bookstaber: Other solutions could probably have been found if more time had been available.

The Fed could have waited until things worked out, but the Fed took another course because it perceived a time of real financial crisis. These were the major financial markets of the world, and if something had not been done, the situation could have been much worse. It was already much worse from a systemic standpoint than the crash of 1987, but from the perspective of most individual investors, the crisis was behind the scenes because it dealt with esoteric instruments. For the financial marketplace, however, these were the primary financial instruments.

The Fed has taken a lot of heat for its activist role, but in that position, you have to step up and do what you think is right even if you have to explain afterwards. It is a mark of courage and perspicacity on the part of the Fed that it would take the step that was necessary, even if the action was unorthodox and opened the Fed up to criticism. The alternative would have been far worse. At least we have the luxury of debating the propriety of the Fed's actions and whether there was some conflict of interest. I would rather be debating than dealing with the aftermath if nobody had protected these markets.

Question: How do investors protect themselves from the malling of Wall Street and lack of diversification among participants?

Bookstaber: If you are an individual investor, the malling of Wall Street probably does not matter quite so much because your positions are small and you can get out quickly. If you are an institutional investor, you have to start looking at diversification in a different dimension. Low-energy diversification is the Markowitz diversification. High-energy diversification is looking at diversifying among net asset classes, among market participants, and among habitats so that if something happens in one area, it is less likely to affect your holdings in other areas. The more that globalization and the malling of Wall Street occurs, the harder it is to do that high-energy diversification, because Wall Street goes beyond the boundaries of Wall Street or the United States. Capital can flow from anyplace to anyplace else.

Question: If these crises are the result of a time disintermediation between liquidity suppliers and demanders, why don't the markets recover much faster?

Bookstaber: If you think it took a long time for recovery—whether it was the crash of 1987, LTCM, or the junk bond crisis, which was a multiyear ordeal—that is, unfortunately, the nature of systemic risk. Recovery could have been much slower and more painful than it was. In a normal market, liquidity demanders are serviced by liquidity suppliers who are in the market, and participation in the market is a function of price. When a cycle is created in which prices do the opposite of what they are supposed to do and suppliers disappear or become demanders themselves, that is a wrenching experience for all concerned, especially those who have not had such a previous experience. As is the case with any experience that shatters our illusions and causes us to rethink long-held assumptions, recovery comes slowly.

If the suppliers had been there at the same time as the demanders, October 19, 1987, would have just been another day and prices would not have dropped 20 percent. If the suppliers had been there for LTCM so that when LTCM had that first margin call it could have sold at a reasonable price and met the margin, then life would have gone on. Neither scenario happened, and recovery was difficult.

Question: How would you describe your view of risk management?

Bookstaber: I think about the markets as a scientific enterprise rather than an accounting enterprise. Many facets of the markets are accounting oriented, or the mathematical equivalent of accounting; examples include modern portfolio theory and the capital asset pricing model. These accounting-type models are important, but we have to look beyond the simple relationships and resulting output.

During the oil crisis in the mid-1970s, the speed limit was dropped to 55 miles per hour. One firm ran this information through its models and discovered that auto insurers would profit from the reduction in the speed limit. We have to learn to make this type of connection between an oil crisis, lower speed limit, and the decision to buy stock in auto insurance companies. When Chernobyl blew up, a lot of people saw it

only as a terrible event, but somebody saw it as an opportunity to buy wheat futures. Making that kind of connection is easy to do after the fact and does not require deep analytical tools, but it does require a scientific or analytical view of how the world is tied together.

Looking at risk management from a scientific perspective is important because the risk that finally hurts most is the risk that you do not know about. Refining our bread and butter measures of risk—VAR, stress tests, and similar tools—will not bring us much closer to uncovering the most critical risks. Granted, they are valuable tools for measuring well-known risks, and they are capable of assessing the likelihood of somebody losing money because a known market factor, such as interest rates or equity prices, moves precipitously. But what matters most are the risks we do not recognize until they occur; after the fact, it is always easy to say, “I should have known that.” The challenge is to try to see the risk ahead of time, to imagine the unimaginable.

<http://www.pbookshop.com>