



MARKETS, HEDGE FUNDS, AND THE
PERILS OF FINANCIAL INNOVATION

A DEMON OF
OUR
OWN DESIGN

RICHARD BOOKSTABER

A DEMON
OF OUR
OWN DESIGN

A DEMON OF OUR OWN DESIGN



*Markets, Hedge Funds,
and the Perils of
Financial Innovation*

RICHARD BOOKSTABER



John Wiley & Sons, Inc.

Copyright © 2007 by Richard Bookstaber. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

Published simultaneously in Canada.

Wiley Bicentennial Logo: Richard J. Pacifico

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600, or on the Web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books. For more information about Wiley products, visit our Web site at www.wiley.com.

Library of Congress Cataloging-in-Publication Data:

Bookstaber, Richard M., 1950–

A demon of our own design : markets, hedge funds, and the perils of financial innovation / Richard Bookstaber.

p. cm.

Includes bibliographical references and index.

ISBN: 978-0-471-22727-4 (cloth)

1. Hedge funds. 2. Risk management. I. Title.

HG4530.B66 2007

332.64'524—dc22

2006034368

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

In memory of my son, Joseph Israel Bookstaber



CONTENT

<i>Acknowledgments</i>	ix
<i>About the Author</i>	xi
CHAPTER 1 ~ <i>Introduction: The Paradox of Market Risk</i>	1
CHAPTER 2 ~ <i>The Demons of '87</i>	7
CHAPTER 3 ~ <i>A New Sheriff in Town</i>	33
CHAPTER 4 ~ <i>How Salomon Rolled the Dice and Lost</i>	51
CHAPTER 5 ~ <i>They Bought Salomon, Then They Killed It</i>	77
CHAPTER 6 ~ <i>Long-Term Capital Management Rides the Leverage Cycle to Hell</i>	97
CHAPTER 7 ~ <i>Colossus</i>	125
CHAPTER 8 ~ <i>Complexity, Tight Coupling, and Normal Accidents</i>	143
CHAPTER 9 ~ <i>The Brave New World of Hedge Funds</i>	165
CHAPTER 10 ~ <i>Cockroaches and Hedge Funds</i>	207
CHAPTER 11 ~ <i>Hedge Fund Existential</i>	243
<i>Conclusion: Built to Crash?</i>	255
<i>Notes</i>	261
<i>Index</i>	271

ACKNOWLEDGMENTS

Nearly a decade ago I made a presentation to the Institute for Quantitative Research in Finance on the origins of market crisis. One of those in attendance was the economist and author Peter Bernstein, who encouraged me to pursue the topic as the basis for a book and provided me with initial guidance for the process, including an introduction to the editors at John Wiley & Sons.

The road from that point to completion was far longer than I, Wiley, or just about anyone I knew would have anticipated. The editorial staff at Wiley has been patient in waiting out the manuscript as it moved forward in fits and starts. The book only reached its final form with the editorial guidance of Bill Saporito, business editor of *Time* magazine, who corralled the mesh of my academic prose, historical vignettes, and biographical events into a cohesive and readable result. My wife, Janice Horowitz, formerly a journalist with *Time*, supported me throughout the writing of the book and contributed her expertise in editing the final product.

The book has benefited from, and indeed to a large extent has as its topic, those who have enriched my professional life. Many are mentioned in the book so I will not list them here. But I wish to close with a nod of appreciation to those who introduced me to many facets of the exciting and challenging field of finance and who have worked as my colleagues with intensity and integrity through periods of exuberance and crisis.

ABOUT THE AUTHOR

Richard Bookstaber runs a market neutral equity hedge fund at FrontPoint Partners.

He came to FrontPoint from Ziff Brothers Investments, where he was the managing director responsible for risk management and the Quantitative Strategy Group. In the latter capacity he developed and managed the firm's quantitative long/short equity portfolio. Before joining Ziff in 2002 he was responsible for risk management at Moore Capital Management. Prior to joining Moore, Dr. Bookstaber was the managing director in charge of firmwide risk management at Salomon Brothers, and served on that firm's Risk Management Committee. He remained in these positions at Salomon Smith Barney after the firm's purchase by Travelers in 1997 and the merger that formed Citigroup.

Before joining Salomon in 1994, Dr. Bookstaber spent 10 years at Morgan Stanley in quantitative research and as a proprietary trader. He also managed portfolio hedging programs as a fiduciary at Morgan Stanley Asset Management. With the creation of Morgan Stanley's risk management division, he was appointed as the firm's first director of market risk management.

Richard Bookstaber also is the principal of Scribe Reports, a firm that provides analytics for skill assessment of long/short equity portfolio managers. He is the author of a number of books and articles on finance topics ranging from option theory to risk management, and has

ABOUT THE AUTHOR

received various awards for his research, including the Graham and Dodd Scroll from the Financial Analysts Federation and the Roger F. Murray Award from the Institute of Quantitative Research in Finance.

He received a Ph.D. in economics from MIT.

CHAPTER 1

INTRODUCTION: THE PARADOX OF MARKET RISK

While it is not strictly true that I caused the two great financial crises of the late twentieth century—the 1987 stock market crash and the Long-Term Capital Management (LTCM) hedge fund debacle 11 years later—let’s just say I was in the vicinity. If Wall Street is the economy’s powerhouse, I was definitely one of the guys fiddling with the controls. My actions seemed insignificant at the time, and certainly the consequences were unintended. You don’t deliberately obliterate hundreds of billions of dollars of investor money. And that is at the heart of this book—it is going to happen again. The financial markets that we have constructed are now so complex, and the speed of transactions so fast, that apparently isolated actions and even minor events can have catastrophic consequences.

A DEMON OF OUR OWN DESIGN

My path to these disasters was more or less happenstance. Shortly after I completed my doctorate in economics at the Massachusetts Institute of Technology and quietly nestled into the academic world, my area of interest—option theory—became the center of a Wall Street revolution. The Street became enamored of quants, people who can build financial products and trading models by combining brainiac-level mathematics with massive computing power. In 1984 I was persuaded to join what would turn out to be an unending stream of academics who headed to New York City to quench the thirst for quantitative talent. On Wall Street, too, my initial focus was research, but with the emergence of derivatives, a financial construct of infinite variations, I got my nose out of the data and started developing and trading these new products, which are designed to offset risk. Later, I managed firmwide risk at Morgan Stanley and then at Salomon Brothers. It was at Morgan that I participated in knocking the legs out from under the market in October 1987 and at Solly that I helped to start things rolling in the LTCM crisis in 1998.

The first of these crises, the 1987 crash, drove the Dow Jones Industrial Average down more than 20 percent, destroying more market wealth in one day than was generated by the world economies in the previous two years. The repercussions of the LTCM hedge fund default sent the swap and credit markets, the backbone of the world's financial system, reeling. In the process it nearly laid waste to some of the world's largest financial institutions. Stunning as such crises are, we tend to see them as inevitable. The markets are risky, after all, and we enter at our own peril. We take comfort in ascribing the potential for fantastic losses to the forces of nature and unavoidable economic uncertainty.

But that is not the case. More often than not, crises aren't the result of sudden economic downturns or natural disasters. Virtually all mishaps over the past decades had their roots in the complex structure of the financial markets themselves.

Just look at the environment that has precipitated these major meltdowns. For the crash of 1987, it was hard to see anything out of the ordinary. There were a few negative statements coming out of Washington and some difficulties with merger arbitrage transactions—traders who play the market by guessing about future corporate takeovers. What else is new? The trigger for the LTCM crisis was something as remote as a Russian default, a default we all saw coming at that. Compare these with the market

INTRODUCTION: THE PARADOX OF MARKET RISK

reaction to events that shook the nation. After 9/11, the stock market closed for a week and reopened to a drop of about 10 percent. This was a sizable decline, but three weeks later the Dow had retraced its steps to the pre-9/11 level. Or go back to the assassination of President John F. Kennedy in 1963 or the bombing of Pearl Harbor in 1941. Given the scope of the tumult, the market reactions to each event amounted to little more than a hiccup.

There is another troublesome facet to our modern market crises: They keep getting worse. Two of the great market bubbles of the past century occurred in the last two decades. First, the Japanese stock market bubble, in which the Nikkei index tripled in value from 1986 through early 1990 and then nearly halved in value during the next nine months. The second was our own Internet bubble that witnessed the NASDAQ rise fourfold in a little more than a year and then decline by a similar amount the following year, ultimately cascading some 75 percent.

This same period was peppered with three major currency disasters: the European Monetary System currency crisis in 1992; the Mexican peso crisis that engulfed Latin America in 1994; and the Asia crisis, which spread from Thailand and Indonesia to Korea in 1997, and then broke out of the region to strike Russia and Brazil. The Asia crisis triggered losses that wiped out the majority of the market value that the Asian “Tiger” economies had amassed in the prior decade of booming growth. LTCM seemed just as cataclysmic at the time, but it centered on a single \$3 billion hedge fund in 1998, albeit one that had more than \$100 billion at risk. As a debacle, it was later overshadowed by the spectacular failures of Enron, WorldCom, and Tyco after the dot-com collapse. Yet, did anyone even notice the convertible bond collapse that erupted for no apparent reason in 2005 or the \$6 billion of losses by Amaranth in September 2006? It’s only money.

One of the curious aspects of worsening market crises and financial instability is that these events do not mirror the underlying real economy.¹ In fact, while risk has increased for the capital markets, the real economy, the one we live in, has experienced the opposite. In recent decades the world has progressively become a less risky place, at least when it comes to economics. In the United States, the variability in gross domestic product (GDP) has dropped steadily. Year by year, GDP varies half as much as it did 50 years ago. The same holds for disposable personal income. With greater stability in economic productivity and earnings, and with greater

and broader access to borrowing—think of your home equity line of credit—the variability of consumption year by year is less than a third of what it was in the middle of the twentieth century. And while recessions still occur, they have become shallower. This same pattern is true in Europe, where both GDP and consumption have become more stable over the course of the past 50 years.

There is ample reason for the increased economic stability. In the United States, the federal government provides unemployment insurance and Social Security, most corporations support 401(k) accounts, and many provide pensions. Governments worldwide stabilize commodity and farm prices with massive subsidies. Monetary and fiscal policy has improved with experience and study, and it benefits from improving coordination and real-time access to data.

The workforce is more diversified, with a much greater proportion employed in noncyclical sectors such as technology and services than in the past. The economic sectors themselves are also far more diversified. In the early twentieth century, there were no technology, telecommunications, media, or health care sectors. The industrial economy revolved around a few highly integrated, large-scale industries. A coal miners' or steelworkers' strike would cripple the country, shutting factory floors and shipping yards. Even as late as the 1970s, the industrialized nations were so energy dependent that an oil shock precipitated a global recession. Today, high gasoline prices cause lots of grumbling, but little real pain.

Similarly, as progress and refinement reduce risk, so should they also level the playing field for market participants. There should be less of a gap between your investment returns and those of Wall Street pros. Do you think that's happening? Sure, the trappings are there: Information is released more quickly and to a broader constituency of investors, and limitations are imposed on insider trading and nonpublic disclosure. Trading costs are a tenth of what they were 30 years ago. Ample liquidity and innovative financial products—all manner of swaps and options, weather futures, exchange-traded funds, Bowie bonds—accommodate trading in more areas. With all these improvements we are moving ever closer to the notion of perfect markets—and perfect markets should not offer unusual profit opportunities for a subgroup of investors and traders.

INTRODUCTION: THE PARADOX OF MARKET RISK

That does not seem to be happening. The market remains volatile and the returns widely uneven. In spite of 40 years of progress and a drop in real economic risk by 50 percent or more, the average annual standard deviation in the S&P 500 index was higher during the past 20 years than it was 50 years earlier. The fact that the total risk of the financial markets has grown in spite of a marked decline in exogenous economic risk to the country is a key symptom of the design flaws within the system. Risk should be diminishing, but it isn't.

Meanwhile, there is a proliferation of hedge funds that continue to capture differentially higher returns. Over the past five years, the assets under management by hedge funds have grown over sixfold from \$300 billion to more than \$2 trillion. And this does not include the operation of the quasi-hedge fund proprietary trading desks at firms like Goldman Sachs or Deutsche Bank. It's a zero-sum game, though, so if hedge funds are able to extract differentially higher returns, someone else is paying for them with comparably subpar returns. Maybe it's you.

This is not the way it is supposed to work. Consider the progress of other products and services over the past century. From the structural design of buildings and bridges, to the operation of oil refineries or power plants, to the safety of automobiles and airplanes, we learned our lessons. In contrast, financial markets have seen a tremendous amount of engineering in the past 30 years but the result has been more frequent and severe breakdowns.

These breakdowns come about not in spite of our efforts at improving market design, but because of them. The structural risk in the financial markets is a direct result of our attempts to improve the state of the financial markets; its origins are in what we would generally chalk up as progress. The steps that we have taken to make the markets more attuned to our investment desires—the ability to trade quickly, the integration of the financial markets into a global whole, ubiquitous and timely market information, the array of options and other derivative instruments—have exaggerated the pace of activity and the complexity of financial instruments that makes crises inevitable. Complexity cloaks catastrophe.

My purpose here is to explain why we seem to be doing the right things but the results go in the other direction. The markets continue to develop new products to meet investors' needs. Regulation and oversight

A DEMON OF OUR OWN DESIGN

seek to ensure that these advances land on a level playing field, with broad and simultaneous dissemination of information and price transparency. But the innovations are somehow making our investments more risky. And more regulation, ironically, may be compounding that risk. It would seem there is a demon unleashed, haunting the market and casting our efforts awry: a demon of our own design.

CHAPTER 2

THE DEMONS OF '87

When I got there in the summer of 1984, Morgan Stanley was still the exclusive partnership it had been since its inception in 1935. The firm's investment bankers lorded over that sexy part of the business, but I was headed for the fixed income division—bonds. You couldn't get less glamorous than fixed income—unless, of course, you worked in fixed income research (FIR), which is exactly where I would spend my first years on Wall Street.

Bob Platt wanted to change all that. A former midlevel insurance executive, he had been snatched from obscurity to run Morgan Stanley's fixed income research division. Obscurity in this case was the giant institutional machine called the Equitable Life Insurance Company, headquartered at 52nd Street and Seventh Avenue, not far from Morgan Stanley's offices at 50th Street and Sixth. When Bob arrived at Morgan Stanley in

1982, the fixed income domain was still just a step or two removed from the backwater of green-eyeshaded bookkeepers, ledgers stacked on their Steelcase desks, tracking bond coupon payments.

This world of fixed income, however, would soon become the vanguard of a revolution made possible by mathematics. At regular intervals the U.S. government issues bonds of varying maturities, from 30 days up to 30 years, with the longer-term bonds typically yielding higher interest rates than the shorter-term ones, to reflect the risk of having your money tied up longer. Plotting these payouts forms the well-known yield curve. Although the prices of the bonds and their respective yields vary as circumstances change—inflation, recession, war—each interest rate along the yield curve is flexibly but securely tethered to its neighboring rates in a way that can be described mathematically.

CRUNCH TIME AT MORGAN STANLEY

The job of the quants descending on Wall Street was to exploit the relationships along the yield curve, to develop mathematical models that would tease a higher return out of a bond portfolio or a bond trading operation than the green-eyeshade gang could. By the early 1980s, a number of other firms were already riding the number crunching wave. Marty Leibowitz at Salomon had built a strong team that was at the top of the heap for fixed income portfolio strategy and yield-curve trading. This group would provide the raw material for Salomon's gold rush into proprietary fixed income trading a few years later.

At Morgan Stanley, Platt wanted to use fixed income research to scale another mountain. An opera aficionado who fancied himself a Brahmin intellectual, he was uncomfortable in the ranks of the meat-and-potatoes bonds crowd. His vision for fixed income research was to slide it away from backwater trader support and propel it into an investment banking role, where the prestige was. The idea, as Bob laid it out to me several times, was to create an investment bank within an investment bank. He would develop products that clients would be willing to pay for. They would then do trades with Morgan Stanley as a way of compensating for the services, with the profits clearly attributed to fixed income research, so that FIR would be a stand-alone profit center. He hired Jim Tilley, a Harvard-trained physicist turned actuary, to run the insurance effort; Jeremy Gold, a pen-

sion expert, to deal with pension funds; and Alden Toevs to tie in with banks and thrifts. Each headed his own small group devoted to client consulting and new business. Bob initially brought me into his department in 1984 to provide the analytical backbone for his effort, but when an opportunity showed itself, he was quick to move me into a profit center role.

THE FORMULA

That opportunity came in the form of a new investment strategy called portfolio insurance, two words that would shortly make an unexpected mark on the investing world. Portfolio insurance is a strategy designed to protect a stock portfolio from dropping below a prespecified floor value. The strategy works by using a hedge—selling S&P 500 futures, for example. If the portfolio increases in value and moves above the desired minimum floor value, the hedge is reduced, allowing the portfolio to enjoy a greater fraction of the market gain. If the portfolio declines in value, the hedge is increased, so that finally, if the portfolio value falls well below the floor price, the portfolio is completely hedged. Thus the portfolio is hedged when it needs it and is free to take market exposure when there is a buffer between its value and the floor value. Because the hedge increases and decreases over time, it is called a dynamic hedge.

The hedging method of portfolio insurance is based on the theoretical work of Fischer Black, Robert Merton, and Myron Scholes. Their work is encapsulated in the Black-Scholes formula, which makes it possible to set a price on an option. No other formula in economics has had as much impact on the world of finance. Merton and Scholes both received the Nobel Prize for it. (Fischer Black had died a few years before the award was made.) The theory and mathematics behind it were readily embraced by the academic community. Adopted from the mathematics of the heat transfer differential equation of physics and employing the new tools of stochastic calculus, it appealed to an academic core that seemed to derive a twisted pleasure from the mathematically arcane.

Despite its esoteric derivation, the formula was timely and—a rarity for work on the mathematical edge of economics—was immediately applicable. First, there was a ready market that required such a pricing tool: the Chicago Board Options Exchange (CBOE) opened for business in 1973, the same year both the paper presenting the Black-Scholes formula and a

more complete exposition on option pricing by Merton were published.¹ Second, although the formula required advanced mathematics and computing power, it really worked, and it worked in a mechanistic way.

The formula gave rise to portfolio insurance through the work of two University of California at Berkeley finance professors, Hayne Leland and Mark Rubinstein.² With John O'Brien, their marketing partner, they founded a management company, Leland O'Brien Rubinstein Associates (LOR), in 1981 to sell their technique. Within a few years it was programmed for action in the computers of some of the largest investment firms in the world. At the start of each day LOR sent its portfolio manager clients hedging instructions based on its runs of the Black-Scholes model. The managers did the hedging themselves. For the service, LOR received an annual fee of 15 basis points of the assets being hedged. For a pension fund of \$1 billion, this was \$1.5 million a year.

Executed successfully, LOR's strategy really did look like buying insurance on the portfolio, because it provided protection against loss for a known cost. It even adopted terminology from the insurance world. The difference between the floor and the current portfolio value was called the deductible; the cost of the trading slippage from the repeated buying and selling for the stock was called the premium.

With portfolio insurance, a money manager could lock in gains without giving up continued appreciation. Many portfolio managers and pension funds found this alluring, especially given what the spectacular rise in the equity market had done for their portfolios. LOR quickly racked up client portfolios worth billions of dollars, with a number of competitors—including me—nipping at their heels. By 1987 portfolio insurance was the driving force in trading decisions, with more than \$60 billion of assets being hedged with the technique.

MY LIFE AS AN INSURANCE SALESMAN

The basic option technology for portfolio insurance was well known, so there were no barriers to other firms following LOR into the market to provide this hedging advice. I spearheaded the effort at Morgan Stanley. Getting the portfolio insurance effort off the ground involved a few months of computer programming and testing. Although the broader investment community considered it complicated, anyone who had been

well trained in option theory, as I had been, found it to be a relatively simple application. As the programs were being put in place, Platt and I made the rounds internally to market the idea. Portfolio insurance required us to take on hedge positions on behalf of clients, which meant acting as a fiduciary. The place in the firm where this could be accomplished was Morgan Stanley Asset Management (MSAM). So key to the effort was a buy-in by Barton Biggs, who was running the asset management arm. Barton was only in his early fifties at the time, but with all the new blood pumping into the firm he already had the role of elder statesman. He had client breakfasts and client lunches nearly every day in the partners' dining room—and made up for it by working out religiously in the afternoon in the company gym, riding the stationary bike while poring over research. For a while I timed my workouts to coincide with his and plied the benefits of portfolio insurance in expanding MSAM's product profile. He finally bit.

After Barton was on board came months of mind-numbing presentations to potential clients. When I taught, I used to think it was boring to give the same lecture once each semester; yet now I was going out on the road and repeating essentially the same spiel dozens of times a month. There were times when I felt as if I was simply exhaling and the words were being produced by muscle memory. But the results of my efforts were surprisingly quick in coming. In less than a year I had more than \$3 billion of assets to hedge. At the time, that was considered a lot of money.

This activity engulfed segments of the firm that rarely related to one another. I was in the fixed income division marketing an equity product to investment banking clients and then managing the resulting programs as a fiduciary in Morgan Stanley Asset Management. I ran portfolio insurance programs for some of the firm's blue-chip clients, including Chrysler, Ford, and Gillette.

This strategy was considered to be at the leading edge of market innovation; rather than buying an existing security, portfolio insurance was ushering in the brave new world of synthetic instruments, financial instruments created on the fly using dynamic trading strategies. In fact, the Japanese, who at the time were in the mode of copying and then improving on every Western innovation, subscribed to the strategy mostly in an effort to reverse-engineer it. (On one visit a Japanese client brought along a stack of spreadsheets that methodically stepped through our daily trade

adjustments, working backward, to infer the functional form of our hedging model.) I made a number of trips to our burgeoning Japanese office and ended up with a dozen or so clients there. In almost every case, they applied the program to a \$100 million portfolio—a drop in the bucket for them and the minimum they felt was needed to pick my brain.

I used the dynamic hedging strategies to synthesize some innovative financial instruments. For example, I teamed up with Vikram Pandit (who later rose to become the head of institutional sales and trading before leaving the firm during the death throes of Phil Purcell's tenure in 2005) to issue the first long-dated options for both the U.S. and Japanese equity markets, using the Amex Major Market Index (XMI) futures (which mimicked the DJI 30) in the United States and the Nikkei futures in Japan. We issued these as principal, which meant that Morgan Stanley took the other side of the hundreds of millions of dollars' worth of option positions, using the dynamic hedging methods of portfolio insurance to maintain the book.

RUNNING FOR THE LIFEBOAT

The equity market was ripe for the promise of portfolio insurance, because there was a lot to insure. From 1982 to its precrash peak in August 1987, the Dow Jones Industrial Average went on a bull run that nearly tripled the index. The U.S. economy cooperated, providing five years of uninterrupted economic expansion. By 1987 the market was moving forward at an exponential rate; from the start of the year to late August the Dow rose more than 40 percent.

By mid-October, though, as the first fissures in the market started to show, the protection provided by portfolio insurance began to look like a very good idea. After the close on October 16, the market seemed to be in a totally different world. The Dow had already fallen nearly 500 points since August, washing away nearly half the year's gains. And then decline became free fall. In the week preceding the October 19 crash the market dropped 4 percent on two separate days. On Wednesday, October 14, the Dow dropped by a one-day-record 95 points, and on Friday, October 16, a new record was set: It dropped more than 100 points.

As the week progressed, equity trading turned into a spectator sport. At Morgan Stanley, fixed income traders and salespeople filtered down the

stairs from the 32nd floor to the equity trading floor to watch the frenetic scene. The equity markets benefited from the built-in structure of having listed exchanges for stocks, futures, and options, and usually were much calmer than the fixed income markets, where every bond the desk purchased had to be taken out to many clients to find a new home. The crowds watching this train wreck amplified the crisis mentality. This finally led Anson Beard, the head of the equity division for Morgan Stanley, to post signs declaring that "Unauthorized Personnel Loitering in the Trading Area Are Subject to Immediate Dismissal." That took care of the riffraff, but the firm's managing directors still found their way to the floor.

And if they were on the floor on Monday, October 19, they got an eye-ful. The open of the futures market at 9 A.M. started a cascade of selling. A half hour later, the New York Stock Exchange opened to an apparently insatiable demand to sell stocks to keep pace with the selling of S&P futures in Chicago. The imbalance of buying and selling demand was so severe that many stocks did not even open, and the rapid decline in the prices of the stocks that did, coming on the heels of the previous week, left most investors frozen in their tracks. The few brave souls who decided to "buy on the dip"—one of whom entered orders to our futures desk with every 100-point decline in the Dow—could not stem the tide of selling. By the end of the day the market had suffered its worst one-day percentage drop in history, down more than 22 percent. The S&P futures fared even worse. The program trading that normally linked the futures' intraday prices to the S&P cash market could not keep up with the selling demand in the futures pit, so the futures dropped even further—nearly 29 percent. Overnight the panic spread around the globe to other equity markets. In the 18-hour period after the New York market opened, wealth equal to several years' worth of gross domestic product (GDP) was wiped away.

The postmortem of the 1987 crash is filled with reams of reports that have tried to fit it into the efficient market, information-based mold of contemporary economic dogma. Events and announcements for the week were scoured to see what could have triggered the decline. There were some negatives in the news; there always are. In the week before October 19 the market was greeted with a surprisingly high merchandise trade deficit; the 30-year bond rose above the 10 percent level for the first time in more than two years; legislation was making its way through the House Ways and Means Committee that would tax so-called greenmail. Over the

weekend Secretary of the Treasury James Baker warned that the United States might stop providing support for the dollar, which led to aggressive selling of dollars in Japan and Europe.

But none of the analysis stood out, either individually or in aggregate, as the source for any sort of major rethinking of the market. The 1987 crash simply was not the result of a rational reaction to new information. What sort of information could have led the market to drop more than 20 percent on October 19 and jump 12 percent early the next morning, only to fall 10 percent in the following few hours? Nor was it a matter of herd psychology. The moon and stars did not align to lead broad segments of the market to wake up Monday morning and decide to dump their shares. In fact, a select and concentrated set of firms generated the selling demand on the 16th and the 19th.

Perhaps because portfolio insurance had sprung forth from academic circles, the academic world was unable to see the root of the problem and looked everywhere but home for the cause of the crash. It didn't take a genius to see that the source of the crash was market illiquidity, an unintended byproduct of the new and wildly successful portfolio insurance strategy—and one that I had helped to popularize and implement.

READ THE FINE PRINT

Unfortunately, the term *portfolio insurance*, coined by LOR, while a great marketing hook, was not an entirely accurate representation of the scheme. Portfolio insurance was, after all, not really an insurance policy, but a dynamic hedging strategy that attempted to create insurance-like protection. The strategy had all the potential pitfalls of any hedging strategy. For the option pricing theory to lead to a predetermined hedging cost, the parameters of the hedge, especially the volatility of the underlying security, have to be correctly calculated, and the statistical assumptions about stock price movements—in particular, that the price will move without sudden gaps—have to be correct. More to the point, the market must be liquid. If a hedge cannot be readily adjusted, then obviously all bets are off. This was especially true in October 1987.

If one small portfolio uses this sort of strategy, liquidity will not be an issue. If everyone in the market is trying to do it, it can become a nightmare, a little like everyone on a cruise ship trying to pile into a single

lifeboat: it won't float. Neither did the market. On Monday morning, October 19, 1987, everybody who was running a portfolio insurance program pulled out the computer runs from Friday's record drop. Some firms had been caught by surprise, so they poured sell orders into the S&P futures pit by the truckload. But other such trucks had started down the road a week earlier, barreling into the option and futures markets from a distant part of the equity trading floor that was run more by rumors and personal contacts than by computers and mathematical models.

RISK ARBITRAGE: THE SNOWBALL AT THE MOUNTAINTOP

No one knew it that day, but the risk arbitrageurs had started the snowball down the mountain. It's not that they actually cared about anyone else, but even the arbs would have tried to pull back had they known otherwise. During the takeover boom of the 1980s, they were voracious, sharks that preyed on other sharks. They bet on who would take over whom. The most prominent names in this business included Ivan Boesky, who would later be arrested, and Robert Rubin, who would later be Secretary of the Treasury.

Morgan Stanley's risk arbitrage desk had neither the visibility nor the profitability of the desks at other firms. Since we were such a force in investment banking, we ended up on one side or the other of almost every deal and so were restricted from taking on positions as a principal on the trading desk. The risk arb department was set up more for information gathering than for moneymaking. It was run by Barry Allardice, an investment banker with a Harvard pedigree who was devoid of pretense, often wore shirts torn at the elbow, and slept in an apartment that looked like a college dorm room. Allardice could take positions in some deals, but the main motivation for setting up his operation was networking. After talking with other deal makers, the risk arbs would alert our investment bank to where the next opportunity or threat would be.

In the fall of 1987, the list of potential takeovers seemed endless: USG Corporation, Tenneco, Gillette, Newmont Mining, and Santa Fe Southern Pacific, to name a few. The weapons for takeovers were high-yield bonds, which provided the capital for a raider to make a run on a company. The payoff as often as not was greenmail, a payment made in essence to buy off

the attacker. The tactic was more akin to institutionalized extortion than anything else, but in the perspective of the times it seemed mostly mildly entertaining. On Wednesday, October 14, a bill wending its way through the House Ways and Means Committee that threatened to tax greenmail-oriented transactions somehow came to a focus on Wall Street. This bill would dim the prospects for many takeovers in the works. As a result, on Wednesday risk arb firms began a broad liquidation.

The run had begun. The drop in the market tied to these liquidations was echoed by portfolio insurance hedgers, who had to respond to the drop with their own programmed selling. The portfolio insurance selling clipped the market by another 50 points, and this was capped off at the end of the day by a third wave of selling. And so the Dow experienced its largest one-day point drop in history, down nearly 100 points.

Thursday was worse. Either caught unaware or waiting for the next day by design, the reaction of portfolio insurance trades brought the Dow down another 50 points. The results of these two days should have sounded the alarm bells for the potential impact of portfolio insurance programs. In retrospect, it was clear that we were getting dangerously close to critical mass in terms of the effect of portfolio insurance on the market, the point where one drop in the market would trigger another, each tied to portfolio insurance hedging and each magnifying the previous one.

IGNORING THE CASSANDRAS

There were warning signs before any of this happened, of course. By the beginning of October there had been rising concerns that things had gone too far, even though the market had retrenched from its August highs. One indication on the technical side was that the premiums for put options were increasing because of a rise in the implied volatility of these options. On the intellectual side was John Kenneth Galbraith, who wrote an article for the January 1987 *Atlantic Monthly* entitled "The 1929 Parallel" that stated bluntly: "The market at this stage is inherently unstable." Galbraith, who had lived through the first Crash, cited the market's spectacular and constant rise, in part because of "the present commitment to seemingly imaginative, and eventually disastrous, innovation in financial structure." By October, though, most of us were too involved in keeping

our heads above water in the face of waves of unprecedented price moves and volatility to think broadly.

Most of us, but not all of us. One salesman on the options desk saw what was coming and retired a few weeks later with a small fortune. Thursday morning, October 15, a young options salesman named Peter Palmedo took me aside and showed me an overlay of the path of the stock market in 1987 versus the path the market took in 1929. The chart had been making the rounds on Wall Street for a few weeks and eventually ended up in the *Wall Street Journal* on October 19. Some have argued that its appearance in the paper was one of the triggers for that Monday's crash. Both markets had exponential growth, and we were just a month or so short of matching the October 1929 decline. The relationship of the patterns at first seemed remarkable, but was actually a natural product of constant positive returns. A steady and stable rate of return will show up as an exponentially growing curve when prices, rather than returns, are plotted. For example, a stock that starts out at \$100 and grows at 20 percent a year will be worth \$120 the second year, \$144 the third and over \$170 by the fourth. The leaps keep getting bigger.

But Palmedo's argument extended beyond a simple comparison of the market of the day with that leading to Black Monday in 1929. He knew about portfolio insurance and had run through a simple thought experiment with me. "If the S&P drops, what are you going to do? And what will LOR, J.P. Morgan, and every other portfolio insurer do?"

Because all of the firms running these programs had the same type of hedging strategies in place—albeit with different clients—in the face of a market drop we would all have to sell to increase our hedges. Palmedo's conclusion: "You have \$3 billion of positions you need to hedge. Between you and LOR and everyone else who is doing this—not counting some places that have started doing this sort of hedging in-house—there must be 20 times that amount. What do you think the effect will be on the market when you all go into the futures market to hedge?"

Palmedo had figured it out. In his view, actions would drive the market down even further, and this in turn would require us to increase our hedges further, which would lead to a downward spiral in prices. It was an easy conclusion to make; anybody could have done it. But everyone seemed to be having too much fun marketing this latest innovation and making money to think seriously about this sort of unpleasant scenario.

Palmedo walked the talk; he bought \$60,000 of out-of-the-money put options, options that would grow in value if the market collapsed.³

A LONG WEEKEND

Wall Street had all weekend to think through the implications of the Friday decline. We had adjusted our hedge gradually over the course of that Friday as we knew that a number of the largest portfolio insurance providers would be waiting until Monday to do so. This was in part for a technical reason—the models were only run once a day based on the market close—but could also be justified on efficiency grounds. Intraday adjustments can lead to unnecessary whipsawing and increase transaction costs as the market moves up and down. Even though Friday was not a normal day, many firms were locked into the next-day adjustment process and could not have made intraday adjustments, even though the market had declined by more than 70 points by midafternoon.

We estimated that the overhang from LOR clients alone would be more than \$5 billion, and while LOR was the market leader, the total overhang across all of the portfolio insurance purveyors could be double this amount. On a normal day, this overhang would be evident ahead of time. LOR advocated what is called “sunshine trading,” where it would preannounce its intended trades. Hayne Leland argued that announcing the trades ahead of time provided better liquidity in two ways. First, market makers would know what to expect and could position accordingly. And second, when the trades came into the market, everyone would know that they were purely liquidity based and that there was no other information behind the trades that could catch the market maker by surprise. That made sense, but it accounted for only one segment of the market. The program trading desk was getting a wave of sell orders from mutual fund clients. Panic had set in among the civilians on Friday, so redemptions made that day and over the weekend meant that a commensurate amount of their equity positions had to be sold.

By Monday morning everyone was lined up at the gate to be the first to get orders filled. Portfolio insurance firms sold nearly half a billion dollars of S&P futures, amounting to about 30 percent of the public volume. The futures prices dropped precipitously, and the stock market had not even opened. About 15 minutes into the futures market decline, we

started to see activity from an unexpected quarter, cash-futures arbitrageurs. Their attempts to capitalize on the apparent chasm between the cash and futures prices would be the red flag that triggered the stampede in the NYSE.

The futures desk on the equity floor had one junior member whose job throughout the entire trading day was to sit with a phone cradled on his shoulder. On the other end of the line was someone whose job during trading hours was to watch the pit and relay the latest trade or bid and offer to the guy on the desk. When a new level was reached, or unusual activity occurred, he would announce it in a monotonic cadence that blended into the background noise of the trading floor. But on the morning of the 19th we were all ears.

The cash-futures spread was reaching levels that were many times what was generally needed to make an arbitrage worthwhile, so a host of cash-futures traders began to bring their orders to our program desk. The only problem—and in our minds it was a big problem—was that the stock market was not even open yet. The discrepancy these traders were observing was based on the current futures price versus the price of the stock market on Friday's close. The traders were basing their actions on stale data; there could be no telling where the stock market would actually open. They put in orders to sell at the market price at the open, under the assumption that the open would be close enough to the Friday close to still make the discount in the futures contracts a profitable trade.

That was a big bet and a far cry from the relatively low-risk enterprise of the usual cash-futures trade. And in this environment, it was even more risky because when the stock market did open, it was almost certain to open down. The execution of the program trade would then be complicated by the downtick rule, which proscribes short-selling a falling stock. The arbitrageurs wanted to buy the futures and sell the stocks short against them. If the market is in free fall, upticks are few and far between, and there are many short sellers trying to squeeze in their execution. It can take a long time to get a trade off. In the meantime, the long futures position is being held unhedged. If the market drops, the trader loses.

The portfolio insurance hedgers found the other side of the market for their trades in the cash-futures traders and market makers. And these traders in turn were depending on the stock market for the other side of their bets. In effect the cash-futures traders were taking the market impact

from the futures pit and transmitting it back to the individual stocks on the NYSE. When they bought the futures and then sold the individual stocks, the individual stocks would finally register the intense selling pressure coming from the futures pit in Chicago. In theory this arbitrage is a natural market mechanism for tying the cash and futures stock markets together, and up to Monday, the 19th, it worked smoothly. But on the 19th the speed and magnitude of the normally smooth waves of selling metamorphosed into a tidal wave that rushed in from Chicago. It was more than the stock market could absorb.

THE AVALANCHE BURIES THE BUYERS

Program traders and arbitrageurs take positions on the S&P contract trading in the futures pit while simultaneously taking opposite positions on the individual stocks that make up the S&P on the NYSE. When the S&P futures contract sells for less than the price of the basket of the individual stocks in the S&P, then the cash-futures arbitrageur buys the S&P and sends in orders to sell the individual stocks. If the price difference is greater than the transaction costs of doing this trade, then they make an almost certain profit. This trade effectively transfers the stock market activities of the futures pit to the individual stocks on the NYSE. That's where things broke down in 1987, and broke down for a simple reason: Stocks are not as liquid as futures.

The problem was that the traders in the S&P pit are mostly market makers, jammed together gesticulating and shouting out orders in hopes of scalping a few ticks. They thrive because of their quick reactions to the market and their speed of execution. By contrast, equity investors of the day who frequented the NYSE were not particularly focused on speed of execution, nor were they concerned with the minute-by-minute movement of the market. As the futures traders reacted to the market and the cash-futures arbitrage traders transmitted that activity to the NYSE floor, the flow hit a wall. Equity investors were not glued to their screens, ready to react en masse. The futures market was operating in broadband and the NYSE on dial-up.

The specialists at the NYSE tried to elicit more buyers by dropping the price, but there was a limit to how much more buying interest they could attract. No matter how quickly the price was dropped, the decision making

by the equity investors took time; unlike the twitch-quick futures pit traders, they made portfolio adjustments only after reasoned consideration. With their limited capital, the specialists were not willing to wait for the process to unfold, and their increasingly aggressive offers ended up backfiring. Prices dropped so violently that many potential buyers started to wonder what was happening and backed off completely.

The root dynamic was what I call time disintermediation—the time frame for being able to do transactions in the futures market was substantially different from the time frame in the equity market, yet these two markets had been linked together through market arbitrageurs. We could see the effect on our trading floor. One of our institutional clients in Boston was bullish on IBM and had discussed strategies for adding more of the stock to his substantial portfolio. His salesman tried to grab him as IBM started to tank, but he was in a meeting. A second call could not locate him. We could imagine him heading off to grab some coffee and leaf through his morning faxes, unaware that the markets had begun to slide.

Back at the NYSE, a day's worth of activity had passed in what seemed to be 10 minutes. The IBM specialist was starting to panic. A flood of sell orders was coming in and there were nowhere near enough buyers to take them off of his hands. With price as his only tool, he dropped IBM another point, and then two more points, to try to dredge up some buying interest.

The institutional portfolio manager in Boston finally got back to his desk, saw the beating IBM was taking, and gave us a call. If IBM had been down a half point or a full point, he would have put in an order. But with stocks in free fall he hesitated to buy, waiting instead to get a read on what was going on with IBM and the market generally. As he spoke with us, he was interrupted and returned to announce that the fund's director of equity investments had asked that no trades be executed until all the portfolio managers could meet to assess the situation.

Because this client worked for an asset management firm with a long-term investment horizon, he could put on a position just as easily tomorrow as today. He watched the market's downward cascade with something approaching detached curiosity. As the markets fell further, from 5 to 10 to 20 percent, he felt some panic. But for the moment he and many of his counterparts were on the sidelines.

Meanwhile, for the specialist, more shares piled up in inventory with each passing minute. Other specialists were faced with the same onslaught and prices fell all around, so now the IBM specialist found that eliciting buyers was even more difficult, because he had to compete with the other falling stocks for attention. It was not long before the offer price for IBM, Big Blue, the bluest of chips, was down 10 points from the open.

The result was a disaster. The potential liquidity suppliers and investment buyers were being scared off by the higher volatility and wider spreads. And, more importantly, the drop in price was actually inducing more liquidity-based selling. With each point drop, the portfolio insurance programs triggered more selling, and the portfolio insurance managers threw more sell orders into the futures market. Because of the dislocation between the hair-trigger execution of the futures and the ponderous decision making on the cash equity side, compounded by the insufficient capital of the specialist to bridge the gap between the incoming supply and the time frame of the potential buyers, the specialist was forced to drop the price of IBM too quickly. The potential suppliers who could have taken on the selling demand—and who would have been willing to do so with modest price concessions had the move been more gradual—got spooked, and the portfolio insurance hedgers demanded even more liquidity than they would have otherwise.

A price drop is normally the dinner bell for buyers. So a precipitous drop should have had traders licking their chops. But it doesn't work that way; if prices drop too far and too fast, it sends the wrong signal. Rather than taking the drop as an indication of liquidity demand, they wondered if it could be the result of some new market information. They viewed themselves at a disadvantage and elected to stay out of the market.

Replay this portfolio manager's reaction over many times, and you basically have the fault line of the crash of 1987. Selling demand increased as prices dropped because of the prewired hedging rules of the portfolio insurance programs. Supply dried up because of the difference in time frames between the demanders and suppliers. By the time equity investors could have reacted to the prices and done some bargain hunting the specialists had moved prices so precipitously that these potential liquidity suppliers were scared away. The key culprit was the difference in liquidity because of the different trading time frames between the demanders and the suppliers in the two markets. 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.

COLLATERAL DAMAGE

The enormous selling pressure from the shorts caused delayed openings, and when the specialists finally did open a stock for trading, it was at a breathtaking discount to the previous Friday close. The cascade in stock prices meant losses for the cash-futures arbitrage traders. The more prudent traders had put in orders to sell the stocks, but waited to buy the futures contracts until those shorts were executed. They succeeded in filling some or all of their baskets, but at prices that washed away the arbitrage. The open for the stock market was at such a discount that it actually led the futures, which for the first hour or so of trading was showing a discount to the cash market of 20 points or more—a normal discount would have been a point or less—to trade back at fair value. With the opportunity gone and the market both volatile and illiquid, these traders quickly pulled out of their positions, buying back the stocks they had shorted and leaving the trade without having executed the futures leg of the strategy. The less prudent traders had put in simultaneous orders to buy the futures and to sell at the market at the open for the stocks. By the time the stocks finally did open, the futures legs of their positions had dropped precipitously, and their cash stock position was only partly filled if it was filled at all. Many of them had lost more through the drop in the futures contract than they would usually make in a year of cash-futures arbitrage. They also pulled out, selling their futures positions, thereby contributing to the selling pressure on the futures exchange, while closing out their shorts in the stocks.

The result was a short respite in the decline in the market and in the disconnect between the cash and the futures. But now with cash-futures traders limping to the sidelines, many with substantial losses, the market had lost the mainstays of liquidity. Others who might have hung in could not be assured they could get the hair-trigger execution they required. With the high volatility, a few seconds' lag between one side of a trade and the other could wipe out its potential profit. The available technology was of no help. Trades were transmitted to the exchange over a system grandly called the Super Designated Order Turnaround (SuperDOT) system, but

there was nothing super about it. The markets were in the infancy of electronic order systems; orders were being sent using primitive 386 PCs communicating via a Hayes Micromodem. Between the short-sale restrictions and bottlenecks from the excessive volume, there was no guarantee that orders would get executed at all.

Buyers were scarce because they did not have time to digest the implications of the morning's turmoil. Meanwhile, the portfolio insurance programs continued their selling. From noon until two o'clock, more than \$1 billion more of portfolio insurance sell orders, nearly half of the public volume in the futures, flooded into the futures pit. Another billion dollars' worth was traded directly in the stock market through program trades.

BAD GAMMA

The problems created by portfolio insurance were compounded by a feature of the option strategy being implemented. When the portfolio is far from the floor price, the hedge is small, and changes in the value of the portfolio require only a small adjustment in the hedge. As the portfolio value drops and nears the floor, the size of the hedge increases and the amount of adjustment that is made for any change in the market increases as well. Another way to think about this is that the change in the value of the option being created (because that is what the hedge is doing) with a change in the underlying security (in this case the market) varies depending on how far the option is from the exercise price.

In the mathematics of option theory this change in the amount of the hedge in relation to a change in the price of the underlying security is called the option's gamma. As the market declined and brought the portfolio insurance programs closer and closer to the floor, the gamma of the hedge increased, leading to more aggressive hedging. This meant the impact on the portfolio hedges, bad as it was early in the day, got worse and worse.

Nothing could change the momentum. The lack of program trading caused the discount in the price of the futures relative to the cash equity market to reappear. Even though the discount extended to 20 S&P points and nearly 150 points in the Dow, many investors who might have dipped in at these levels still thought stocks overpriced. They expected stocks to drive down to the futures level, and so waited to get in. In the middle of all of this, rumors that the exchange might close circulated, and not without

foundation: The chairman of the Securities and Exchange Commission (SEC) was quoted as saying, "I would be interested in talking to the NYSE about a temporary . . . halt in trading."

In the midst of the market disconnect, the inability to get timely execution, the vanquishing of many of the liquidity providers, and the fears of a total break in liquidity through an early market close, the portfolio insurance programs continued robotically to spit out sell orders, oblivious to anything but the current market level and the mathematical requirements of the hedge. In the last 75 minutes of the trading day, the Dow dropped by 300 points, three times as much in a little over an hour as it had in any other full trading day in history.

THE PHYSICS OF THE MELTDOWN

I spent the week chained to my desk, my eyes frozen on the Quotron screen as I struggled to maintain the hedges demanded of the portfolio insurance programs I ran. Prices were moving all over the place, swinging more violently minute by minute than they usually did in an entire day, and the spread required to buy or sell the S&P futures—still the most liquid instrument in the equity market—was a dollar or more, 20 times normal. I had to weigh the implications of holding off on a hedge adjustment on the one hand with the incredible transaction costs in executing in the market on the other.

The huge volatility of the market broke down all but the most fundamental relationships between the market securities. The usual day-to-day world where investors cared about subtleties like corporate earnings or analyst forecasts dissolved as the energy of the market was turned up. All stocks moved together; if it was a stock, it was sold. The market hardly differentiated between domestic and foreign, small cap or large. It was like plasma physics: As matter becomes hotter it becomes less differentiated. The forces that bond atoms together in the form of molecules are overwhelmed, so that rather than having a myriad of different substances, we have the elemental building blocks of the atoms. Turn up the heat even more and the atoms themselves are melded into plasma, positively charged ions and negatively charged free electrons: matter in its most uniform and nondifferentiated state, no longer hydrogen atoms and oxygen atoms, just a seething white-hot blur of matter.

Just as high-energy physics creates a state that is no more differentiable than to say that it is matter, so the high energy in the financial markets created a world where securities were no more differentiable than that they contained risk.⁴

This melding even extended beyond stocks. High-yield bonds, which usually tracked fairly closely to Treasury bonds, suddenly became simply high-risk bonds and traded just like stocks. Meanwhile, Treasury bonds, the antimatter of the world of risk, were grabbed in the flight to quality and traded up in price. This behavior demonstrates a characteristic I have observed and expounded repeatedly: As the market moves into crisis, the absolute value of the correlation of assets approaches one. The problem is you cannot always predict ahead of time if the correlation will be positive or negative. A particular asset might end up hedging another, or it might end up doubling your exposure.

Because the market plasma temporarily strips away many of the economic relationships between securities, a crisis can offer spectacular profit opportunities for those who can stand the heat. While we sat in a conference room late that Monday night, eating pizza and generating risk reports to understand our losses and where our exposure now stood, with Platt pondering his prospects at Morgan Stanley (which turned out not to be good—he spent two months after the crash sitting alone in his office with the door closed, and then got the call that ended his tenure at Morgan Stanley and, as it turned out, his career in finance), downtown at Salomon some of my MIT classmates were hatching plans to seize opportunities that the day's panic had revealed. They had been hit far harder than we were, but they kept their trading hats on, and their analysis quickly showed that some of the most commonsense security relationships had failed.

Although the center of the storm was in equities, it extended into the fixed income markets as well. There was a demand for liquidity and a flight to quality, which meant that demand moved away from corporate bonds into Treasury bonds, and in Treasuries moved toward the most liquid bonds. They sold short the newly issued, on-the-run 30-year U.S. Treasury bond that everyone else on the street was clamoring to buy and bought an equal exposure in the off-the-run 30-year bond the Treasury had issued three months earlier, a bond that had been supplanted for trading by the new on-the-run Treasury and was less liquid.

The two bonds were basically identical—they differed only in that the new bond would have one more coupon payment, 30 years hence—so they should have traded at very close to the same price. But in the frenzy of the crisis, the liquidity premium of the on-the-run bond had gotten totally out of hand, and no one cared or took the time to notice that such a subtle relationship had been broken. The Salomon gang bet that once the market cooled the normal physics would reemerge, pick up on this mispricing, and reestablish the relationship, which indeed occurred. Between this and several similar trades they picked up more than \$100 million. (Ironically, 11 years later this same team, having left Salomon and riding high at Long-Term Capital Management, would trigger a market crisis that would lead to a far more egregious mispricing between various 30-year bonds. By then I would be sitting at Salomon arguing at a risk management committee meeting for the firm to take on a similar position. But with its trading spirit dulled by the Citigroup merger, Salomon would pass on the opportunity.)

The damage to the stock market was accentuated by the fact that many large institutions had increased their stock holdings on the premise that portfolio insurance would protect them from losses. And the damage extended beyond the United States, even though the proximate cause of the crash was a liquidity crisis that was strictly a U.S. phenomenon. Investors in foreign markets could not help inferring that the sizable drop was due to some major new information that was somehow obscured from their view in the near-panic conditions of the day. And while the focus was on the losses in equities, the damage extended to other financial instruments as well.

The frenzy also led to near disasters in the usually routine back-office backwater of stock clearance and finance. Morgan Stanley's head of finance and operations, Louis Bernard, came to work the next day at 7 A.M. to find urgent messages piled up from three money-center banks and the Federal Reserve. Morgan Stanley owed \$700 million in futures margin that had not arrived at the exchange. There was concern among the firm's major creditors and the Fed that Morgan Stanley might be in default. A call confirmed that the funds had been sent, but they had somehow been lost en route. A few hours later it was determined that a harried Telex operator deep in the cavern of Continental Bank in Chicago, overwhelmed by the volume of paper he needed to process, had pushed a stack off to the

side to be dealt with the next day. Included in the stack was a transfer for \$700 million on behalf of Morgan Stanley.

Heroes emerge from panicky times, and on this occasion, John Mack, the head of Morgan Stanley's fixed income division, was one of them. Mack was a clubby manager with a close set of loyal and trusted friends surrounding him, most of whom also shared his passion for golf. Sometimes I felt like I was part of the kitchen staff in a country club. The head of human resources once played this down, saying, "You don't have to be a friend of John's to be successful here. You can just be a friend of one of his friends." But Mack's instinct was always for leadership over political expediency. He took the helm to deal with the portfolio insurance crisis while the other business heads who should have grabbed hold of the reins scurried for cover. This led to the strange result that the greatest equity crisis to hit the firm was being managed by the head of the fixed income division and a bunch of bond option traders.

Once the mechanism behind the crash was understood, one thing was clear: Through my role in implementing portfolio insurance, I had helped precipitate a financial crisis of monumental proportions. Strangely, on the home front at least things remained more or less intact. My clients had lost money just as had everyone who employed the strategy, but they knew what they were getting into. I had gone to great pains to explain that this was a hedging program, and that if the markets were not there to do the hedge, all bets were off. They were not happy about it, but they knew that we had weathered a market where hedging was near impossible, and we received no complaints about the results of the program.

That being said, it was not like anyone rushed to reenroll. Our internal hedges for the long-dated option business had fared no better than had our client portfolios. The inability to hedge and the huge rise in the market volatility meant those options now cost us far more to hedge than we had expected when we sold them. But we ended up recouping a good share of the losses through the increased trading activity that the huge run-up in market volatility precipitated.

On the brighter side, Peter Palmedo saw his \$60,000 investment in out-of-the-money put options explode in value. Those desperate to hedge out positions ended up bidding for the put options he had bought in the weeks before. The option market was nearly frozen from the market gaps, and rather than having all the options trade, the CBOE went into rotation.

One type of option at a time would open up for trading. With each announcement of a change in rotation, the market makers moved from booth to booth like a herd roaming between watering holes, and with each rotation Palmedo unloaded more of his options.

Our own trading desk was one of the buyers. David Booth, who John Mack had assigned to the task of interim risk management, had us pull in options to hedge our remaining risk from wherever they could be found. Our demand was replayed in many other firms, and options prices reflected this. On a so-called fair value basis—that is, priced based on a Black-Scholes methodology with reasonable assumptions for the postcrash volatility—we ended up overpaying for these options by \$3 million. The problem was, as it almost always is, that during the crisis there was not enough time to worry about the costs, and those who were the most knowledgeable at assessing the costs, like me, were also the ones in the middle of the losses, our judgments discredited.

In any case, the options Palmedo had bought a few weeks earlier for pennies had now grown in value to \$20 to \$30 each. When the dust settled at the end of the week, Palmedo, at age 27, made his farewells and retired to Sun Valley, Idaho, with his family and \$7 million in profits.

The surprise for Morgan Stanley was that our biggest losses came not from equities, but from high-yield bonds. The flight to quality moved investments away from equities into Treasury bonds, with the result that as equity prices declined, bond prices shot upward. In the aftermath of October 19, the interest rate on 90-day Treasury bills dropped almost two percentage points to just over 5 percent, and the benchmark 30-year Treasury bond shot up by more than 11 points. Corporate bonds did not share in the shift to fixed income; they went south instead.

Most of the time, the specter of corporate defaults is remote. But after the crash, the possibility of long-term economic damage could not be dismissed. As a result, the corporate bond market started to tank. Not surprisingly, we were hit hardest on the lowest-quality bonds, the high-yield or junk bonds. These bonds fueled the merger mania of the mid-1980s, a market that was absolutely dominated by Drexel Burnham Lambert, but we had put serious resources into establishing a presence. To play catch-up we had to trade aggressively and take on risks, make a market for clients when others wouldn't, and trade in size when clients demanded it. Since we did not have established distribution channels, we held more bonds in

inventory than would make sense in an established business. Large inventory meant large risk, but the traders dealt with this by hedging the positions with a strategy I had developed—and published to some fanfare—called a composite hedge.⁵

WE CAN SEE THE FUTURE OF MARKETS, AND IT'S UGLY

The composite hedge took advantage of academic work that demonstrated how a corporate bond payoff can be replicated by holding a weighted position in the company's equity and in Treasury bonds. The weighting varied over time in much the same way that the portfolio insurance hedges did. The bulk of the hedge was in Treasuries, because the effect of interest rates (except for the lowest-rated bonds, of course) is larger than the effect of idiosyncratic movements in the stock. But the relative weights of the equity and Treasury positions in the hedge moved dynamically as the fortunes of the firm and the quality of the bonds varied. So here was yet another derivatives-based trading strategy that depended on the ability to make smooth adjustments in the hedge over time, and it did not fare any better in the gap moves in prices during the crash.

The bond positions were hedged with short positions in Treasuries, and added to that was a short position in the equity of the firm. Usually if things do not go well for the firm, both the bonds and the equity drop in price, so a short position in the equity counteracts that effect. And if interest rates rise, the bonds will drop in value, so a short position in Treasuries helps to offset that.

The composite hedge was great in theory again but the hedge did not work well because bonds, especially high-yield bonds, started to trade like equities, going down points at a time. The short stock positions tempered this, but only slightly, because they were established for a normal market environment when the equity component of the bonds is small. In theory, the equity exposure would have been ratcheted up, but of course there was no time to do this. Meanwhile, the Treasury component of the bond hedges—put on under the usually reasonable assumption that the corporates would more or less trade like interest rate instruments—went into reverse. The desk lost spectacularly on the corporate bonds it held in inventory, and then lost even more on the short Treasury positions that were suppose to hedge them. The same problems

THE DEMONS OF '87

were faced by several buy-side institutions that had embraced the composite hedge concept and used it to add to their high-yield risk and, they had assumed, to their returns.

So not only had I managed to be at the center of client and desk losses in the equity division, but I had also been able to do the same for the fixed income desk and some of its clients as well. That's versatility for you.

October 19 ushered in, spectacularly so, the context for a type of risk that would embrace the markets in the future. At the root of the crash were two things: One was computer-assisted liquidity, which took the form of rapid, programmed cash-futures execution. The other was an innovative product, portfolio insurance, that demanded that liquidity for successful implementation. The market was developing to allow lightning-fast reactions while at the same time driving complex innovations in market instruments and strategies that required ever more time for investors to analyze and absorb. This combination of speed and complexity would be the source of many future crises.

CHAPTER 3

A NEW SHERIFF IN TOWN

Until the early 1970s, investment banking was more an elite fraternity than a competitive business. It lacked risk, it lacked capital, and it lacked imagination. And it was all about to change. Morgan Stanley, though one of the loftiest of the firms, had only 30 partners and less than \$20 million in capital. The primary business of the investment banks was mergers and acquisitions (M&A) advice and underwriting. The order of things was that the top-tier investment banks advised their corporate clients and managed their underwriting, and then the bonds or equity shares were farmed out to less prestigious trading houses for distribution. The M&A business was done among friends; old school ties and country club memberships counted for as much as innate talent. On the one hand, the underwriting was clear-cut and highly profitable. It could be run with little capital and overhead.

On the other hand, the distribution was closer to a boiler room, with calls to investors, teams of traders and salespeople to manage, and capital at risk if an issue did not sell.

The first challenge to the old order came in the mid-1970s when one of the key distributors, Salomon, decided it wanted in on the underwriting. Salomon figured that if it was going to do the dirty work of distribution, it was also going to co-manage. Forced to invite Solly into their club, the top-tier “white shoe” firms like Morgan Stanley returned the favor by moving into the turf of the distributors. As a starting point, they garnered primary dealer status to establish a strong presence in the Treasury bond market. The timing was perfect; the Treasury bond bull market of the early 1980s was about to begin.

Having stuck a foot in the door through the pristine Treasury market, investment banks moved aggressively into the corporate bond market. Trading desks formed to serve companies that issued bonds and investors who had purchased the new issues. Corporate bonds are more complex than Treasuries for two reasons. First, unlike the U.S. government, corporations can go into default, which means credit risk is added as a consideration to yield curve dynamics. Second, there is far more liquidity risk. There are many more corporate bonds than Treasury bonds, and most corporate bonds only traded by appointment. This led, true to the caricature in Tom Wolfe’s novel *Bonfire of the Vanities*, to bond salesmen whose principal expertise was knowing who owned what bonds (the principal value of having done the bond underwriting) and passing these bonds from one hand to the other for a spread. Sell \$100 million of bonds at a quarter-point spread, and the firm would pocket \$250,000 for making a few phone calls. Success came from building up a client list and knowing what bonds these clients held. Firms used various strategies, ranging from simple cajoling to offering to do large-scale analysis of the clients’ portfolios, to find out who owned what. Then, the next time a client called asking to trade a bond, the corporate bond salesperson knew where to look for a taker. In two or three years every firm had the name of every investor on one of its salespeople’s Rolodexes. Corporate bond salesmen became the trading floor’s version of used car salesmen.

The M&A rush lifted their status. Once pulled in from the dregs of the business night school classes and the back-office ranks, they were now the

prime movers of finance. At Morgan Stanley the fixed income traders were a self-satisfied group with the air of kids who had been picked by the winning team at recess. In their late twenties and early thirties, as self-important as they were clueless, most of them wondered how they had ended up in so lofty a position and took any cue for how they should behave. They emulated John Mack by practicing their golf swings as they stood around their desks. They emulated their image of the big kids—the Salomon traders—by chomping on cigars when they were making money and slamming phones when things did not go their way. In the early stages, just having a seat on the desk was enough to assure profits, so getting in quickly was more important than getting things right. The traders seemed to pretty much make it up as they went along, accumulating on-the-job training and losses through a bad trade here and there. Keep your head above water, avoid major blowups, block other desks from elbowing in to take a piece of your market, and pull in huge bonuses.

These illusions of grandeur would pass quickly as the wave of opportunity moved on to the next hot market: mortgages. Clever structures for splitting up mortgage instruments pushed Salomon and First Boston into the forefront, creating early-retirement opportunities for mortgage traders after a few frenetic years. Other investment firms quickly followed suit, opening up their trading desks to the new, exotic instruments while, like an airplane trying to build its navigation system in midflight, they grabbed up mathematicians and leased Cray supercomputers to try to figure out what they were doing. Each new market reaped a few years of huge profitability before competitive forces drove the spreads down and a new market rose up to become the envy of every freshly minted MBA trainee. Each year fortunes were made as the wave of new products washed past one market and then another.

The breathless reports of Wall Street compensation, while perhaps correct on average, missed (and continue to miss) one important point: The top bonuses may be huge, but the recipients of these bonuses change nearly every year. The big payday for the corporate bond salespeople lasted only as long as they were the hot commodity. After a year or two they were supplanted by the mortgage traders. In the context of an entire career, the lucky ones had one or two years of huge paychecks before settling back down to merely outlandish remuneration for their jobs of making markets or pitching to clients.

A DEMON OF OUR OWN DESIGN

The illusion may have been one of a lottery being won at every hand or the clinking coins and flashing lights of a casino's slot machines, but in reality the prize moved from one desk to the next. The only sure way to create an annuity was to elbow your way into management to get a cut of the action no matter from whence it came.

With so much money to be made and with the clock ticking until others piled on, new products were brought to the market before they were road tested. Basic questions about the security's price behavior or how risk could be managed took a backseat. Those who were late to the party, the ones who would be described as "competitive forces" pushing prices down to try to entice business, entered the fray not to seize profits but to stem the cost of being left behind. For the latecomers it was all downside; if they didn't play the new game they would be seen as unresponsive and lose their clients to other firms that could help the clients try on fashionable new securities. It was a defensive move that did not provide unusual profits.

The 1980s ended with a range of innovations in swaps and derivatives. At Morgan Stanley we marketed and traded through a new entity that spanned trading and investment banking called the Derivative Product Group. With a new decade dawning, though, the pace of innovation slowed. It may not have been so much that there were no new mountains to climb as it was that the payoff in doing so had diminished.

With each innovation, competing firms mobilized and clients caught on all the more quickly. Each new round of instruments was more sophisticated and complex than those that preceded it, although not necessarily understood as well. There was one new act to follow the development of swaps. True to the Marxist dictum for the path of capitalism, the last foray was emerging markets debt.

As with every other firm trying to make hay in the go-go years of the 1980s, we had moved quickly to seize one opportunity after another, never pausing to develop a solid infrastructure for managing all the trading or, for that matter, much of an organizational structure to manage all the traders. Books that were managed on Excel spreadsheets in the nascent days of a trading desk were still managed the same way, but with positions that had grown by orders of magnitude and with one new product thrown on top of another. Inventory was hedged and pushed back onto the heap like so many pairs of old underwear shoved behind the closet door. Money

had been flowing in so quickly that no one was bothered if some sloshed over the sides. But once swaps and emerging markets followed the others into the commoditized world, the froth of profits receded and all the carcasses and rusted hulls that had been ignored rose into view.

John Mack, a strong leader with an appreciation for the organizational aspects of the problem, and, it turned out, an eye on preparing the firm for sale, began to work on his “one-firm firm” program. Mack actually wanted to corral everyone and make Morgan Stanley look like a corporate entity, and not the collection of power bases it had become. Over the course of the 1980s, the trading operations on Wall Street had grown organically, a polite way of saying there were few points of control or accountability. The partnership model evolved into fiefdoms that could grab capital and call their own shots so long as they were passing profits up the line.

SEARCHING FOR LAND MINES WITH THEIR FEET

This lack of control would become increasingly evident as one company after another blew up during the first half of the 1990s.

In 1992, Metallgesellschaft Refining and Marketing (MGRM), the New York-based subsidiary of the German conglomerate Metallgesellschaft, implemented a program to offload oil price risk for companies that used petroleum products by contracting to provide the products at a fixed price for up to 10 years. (If this sounds familiar, perhaps it’s because Enron would try something similar with energy.) The fixed price was set to yield a profit versus current market levels, and the strategy was hedged using short-term energy futures. If oil dropped, the hedge would lose money, but the fixed-price contracts would increase in value. The strategy might have worked but for one detail: Any gains from the sale of oil are realized over time as delivery is made to MGRM’s customers, but any losses in the hedge are realized immediately.

MGRM fell into such a trap in the fall of 1993. Oil prices dropped and the company faced a cash-flow crisis. The cost of rolling over the futures contracts was nearly \$100 million in October and November. MGRM had to obtain funding from its parent organization, Metallgesellschaft, to cover these costs. Its management, either not understanding or not agreeing with the strategy, decided to close out the positions to curtail its losses.

In December 1993, the company cashed in its positions at a loss exceeding \$1 billion.

Meanwhile, in Orange County, California, treasurer Robert Citron had been structuring trades with the help of friends at Merrill Lynch to borrow on the short end of the yield curve to finance positions in the usually higher-yielding intermediate-term rates. Citron's strategy depended on short-term interest rates remaining relatively low when compared with medium-term interest rates. This they did in the early 1990s, so Citron's yield curve bet made money and everyone was happy, with no questions asked. Even in early 1994, when his strategy began to go south, he survived an election that focused attention on his financial management, convincing voters that the criticisms were just so much politically motivated rhetoric. Then the Fed started raising rates in February 1994, and the yield curve started to move the wrong way. Orange County got crushed. A succession of hikes saddled Citron's fund with losses of approximately \$1.7 billion, around 20 percent of its value. The County filed for bankruptcy and then filed suit, reaching a \$400 million settlement with Merrill Lynch for steering Citron toward high-risk, unsuitable securities.

The sudden rise in rates that sank Orange County also took its toll on Bankers Trust, which had been a leader in marketing derivatives to corporate America. Bankers Trust was sued by four major clients—Procter & Gamble, Gibson Greetings, Federal Paper Board Company, and Air Products—big and ostensibly smart companies that asserted they had been misled by Bankers Trust concerning the risk and valuation of derivatives they had purchased. Far greater than the nearly \$200 million in damages from the cases was the damage to Bankers Trust's reputation.

At least P&G could absorb the losses. Other firms were wiped off the map. Nick Leeson, working in the hinterland of Barings Bank's Singapore office, traded unauthorized positions for several years, a simple matter because he was both the trader and the back office for its futures operation. In July 1992 a clerk at Barings mistakenly entered an order to purchase 20 Nikkei futures contracts as a sell. Leeson failed to notice the £20,000 error until that Friday night, and it slipped his mind to correct it when he returned to work on Monday. By the time he finally got around to addressing the error, it had grown to a £60,000 loss. He decided to deal with the error by making a false entry to negate the trade and then entered the loss as an error account until his trading could recoup the loss.

A NEW SHERIFF IN TOWN

Over the course of the next two and a half years, his attempts to recoup this once small loss led to one bad trade after another, culminating in a disastrous position. With Leeson long more than \$8 billion in Nikkei contracts in an account called 88888 (not to mention holding nearly 90 percent of the open interest in Japanese government bond contracts), an earthquake hit Kobe, Japan, generating a flight to liquidity because of a drop in both investor confidence and the anticipated funds that would be needed to repair the damage. The Nikkei dropped 300 points on the day of the earthquake, 1,000 points the next day, and more the day after that. The vibrations were going to be felt in London.

Leeson had been able to conceal his unauthorized trading not only because he managed both the trading and back-office functions, but also because the senior managers at Barings came primarily from a merchant banking background and knew very little about trading. Even in the face of large profits, which should have tipped management off to the fact that substantial risks were being taken, they continued to believe that Leeson held matched positions on the Singapore International Monetary Exchange (SIMEX) and the Osaka Exchange, and hence was making a low-risk profit.

In fact, Leeson was trading derivatives contracts on the two exchanges that were, in some cases, of different types and, in some cases, in mismatched amounts. For example, Leeson executed a trading strategy known as a straddle, with the objective of making a profit by selling put and call options on the same underlying financial instrument, in this case, the Nikkei 225 index. A straddle will generally produce positive earnings when markets are stable but can result in large losses if markets are volatile.

In February 1995, Barings, the oldest of Britain's merchant banks, a 223-year-old institution that had once helped the United States finance the Louisiana purchase, was bankrupted by \$1 billion of unauthorized trading losses. Barings was purchased by ING Group for £1 and disappeared.

Then there was Joe Jett. Jett entered Wall Street with a degree in chemical engineering from MIT and an MBA from Harvard Business School. He enjoyed no commensurate success in the business world. Three years out of school he had been terminated from trading positions at both Morgan Stanley and First Boston. Then in July 1991 he started work as a trader on Kidder, Peabody's STRIPS (separate trading

of registered interest and principal securities) desk. The STRIPS desk takes Treasury bonds and strips apart their coupons to sell as individual “strips” or zero coupon bonds, and also works in the reverse, pulling together zero coupon bonds from various sources to rebuild or reconstitute Treasuries. Jett lost money in his first month trading at Kidder, was close to flat the following months, and received a negative performance review for the year. He could see the writing on the wall for a third failure in his trading career. So he resourcefully developed a trading strategy to improve his performance.

The stripping and reconstitution of Treasuries are noncash exchanges of essentially economically equivalent securities that are easily processed, but for accounting purposes at Kidder a transaction that reconstituted a Treasury was entered as a sale of the STRIPS and the purchase of the bond. As with a real trade, Jett could enter a trade date and a settlement date for the exchange. As there was no real purchase or sale and no price or date set by counterparties, these profit and loss (P&L) figures were purely mathematical constructs that would finally converge to equal one another by the time the settlement date arrived. Until then, there was an apparent profit in one part of the trade. So by entering these exchanges as real trades with a counterparty and putting the settlement date as far forward as the system allowed, he could generate a virtual inventory that showed a current profit.

Jett was losing money most months in his actual trading, but the phantom STRIPS trades more than compensated for those losses. His fortunes exploded with a swiftness rarely seen outside of the world of Wall Street traders. His performance rating went from “poor” to “outstanding” and his boss called him “one of the best STRIPS traders in the business.”

Things were going in Jett’s favor. Even the technology cooperated. Kidder upgraded its computer system, if *upgraded* is the word, so a trader could enter any future settlement date. Now that’s progress: Jett could then extend the lag between trade and settlement, and thus hold onto his profits longer. By the end of 1992, profits for the STRIPS desk were \$32.5 million, more than five times the profit for 1991, even though the real P&L for the desk was a loss of \$8 million. Jett racked up another “outstanding” evaluation in 1992, got promoted, and received a \$2.1 million bonus.

The Joe Jett juggernaut recorded profits of nearly \$150 million between September 1993 and March 1994, even while the real P&L for the desk was spiraling downward to a \$50 million loss. Although it was only one of 12 trading desks in the fixed income division, the STRIPS desk accounted for more than a quarter of the division's profits. Jett was awarded a \$9.3 million bonus in 1993 and named Kidder's "Man of the Year." There was more than enough money to spread around; Ed Cerullo, the division head, claimed a \$20 million bonus largely on the back of Jett's results.

No one really knew how Jett was actually making all of this money (again, think Enron), and Jett rebuffed them, brusque and vague, when they asked him. He would throw out loose descriptions of arbitrage (which is very difficult to do successfully in so liquid and transparent a market), but being on such a roll, he did not have to explain himself to anyone until Cerullo finally decided there were too many dollar bills raining down from the heavens. He initiated an inquiry into Jett's trading not because he suspected fraud but out of a concern that his desk might be taking undue levels of risk. How prudent.

The investigation did not take long to start alarm bells ringing. Jett's strategy was at its foundation a pyramid scheme because as the settlement date from any of his trades approached, he had to add even more exchanges to his book to overcome the accretion of the zero coupon component. By the end of July 1993, he had entered so many reconstitution instructions that in some cases the open instructions provided for purchasing more STRIPS than actually existed. The dollar value of new forward instructions rose from \$50 billion in mid-1993 to nearly \$700 billion a year later.

The investigation naturally led to the concern that these huge positions might be difficult to settle. Because of general balance sheet pressures at the firm and the settlement questions, Cerullo directed Jett to settle or pair off his forward Fed exchanges to remove the positions from the firm's records. At this point the chickens came home to roost. Removing the forward exchanges from the ledger through these pair-offs resulted in a shortfall of some \$300 million. It finally became obvious that Jett's trading performance was an accounting illusion with no economic substance.

Kidder fired Jett, and General Electric, Kidder's parent company, took a charge of \$350 million. Then GE in effect fired Kidder, selling the

company to PaineWebber for \$600 million in late 1994. One more illustrious Wall Street name passed into history.

These and any number of other screwups made risk management a hot topic on Wall Street. Merrill Lynch was early on the case because of a spectacular trading loss on its mortgage desk in 1987, followed shortly thereafter by Bankers Trust. We got into the act at Morgan Stanley in 1993 when Bob Feduniak established the risk management division. Feduniak was one of the “Class of ’84” managing directors. He had been recruited from J. Aron, which became a part of Goldman Sachs, by Lewis Bernard in 1982 to start the commodity trading group. He was thoughtful, even scholarly, with a BS from Stanford and an MS from Berkeley. And he was a great poker player, having placed in the top 20 in the World Series of Poker.

Feduniak asked me to join his new effort and appointed me as Morgan Stanley’s first market risk manager. Most of the market risk was centered in the arcane world of derivatives and products with option-like features such as mortgages, so while it was a plus that I had spent time trading and weathering the 1987 crash, my background as a quant was even more critical in that I understood those instruments. In fact, the first wave of risk managers at most of the other firms also were principals with PhDs and experience with derivatives. The decision to take this position was a difficult one for me, because I had just started running my own proprietary trading book and I was also looking to establish my own trading operation at Paloma Partners Management Company, a firm that set up portfolio managers and traders within its own umbrella of capital and back-office support. In the end, I saw this as a unique time for the risk management responsibility—like being a sheriff during the taming of the Wild West. The opportunity to trade was not going to disappear; I could pick it up again down the road. And with six children at home, becoming a risk manager seemed more prudent than becoming a risk taker.

While I didn’t run into problems as dramatic as those that sank Barings and Kidder, there were a few tremors that served as object lessons for the potential for loss, centered, not surprisingly, in the realm of derivatives, and again, not surprisingly, by rogue traders trying to game the system. It proved disconcertingly easy to do.

The price of fixed income derivatives is based on the offer, but the auditors do their price verification based on the average of the bid and the offer. One trader hid his losses by gradually moving up the offer side of his

volatility to keep his option P&L flat in the face of a dropping market, and then fooled the price verification by dropping the bid lower and lower so that the average of the bid and the offer continued to match the market volatility. It finally got to the point that his bid/offer was .02/.20 when a typical spread would have been something more like .19/.20—and it was a cursory visual inspection of the volatility of the positions that gave the ploy away, but only after it had cost the firm \$20 million.

Another trader who found himself underwater kept his positions from being repriced by simply not doing any trades. The prices of his inventory were marked to market only when there was a new trade, so whenever an inquiry came in from a customer, he would offer the bonds at far above the going price so that the customer would shop elsewhere and the trade would not go through. Without having any new trades, his positions remained on the books at cost despite a major market downturn. The firm lost more than \$100 million by the time someone noticed that the trader hadn't traded for nearly a month. And complicating matters further is that most trading is done through verbal agreements. Is there any other industry where hundred-million-dollar sales occur without so much as even a handshake?

THE APL CULT

These smaller problems pointed out that the first order of business in understanding Morgan Stanley's trading risks was that we had to be better shopkeepers. We had to be able to price the things we were trading and holding in our inventory. This clear objective turned out to be an exceedingly difficult one to meet because of a technical wrinkle in our information technology structure—our computer analytics were based on a programming language that was poorly structured for executing loops.

Morgan Stanley's IT department had become a monastic way station for a cult of adherents to a cryptic but elegant programming language called APL (short for "a programming language"). If APL had been a religion, Morgan Stanley could have claimed a blood lineage to its priesthood, because the son of APL's founder worked in the fixed income research (FIR) division. Ken Iverson, a Harvard professor, had invented APL as a mathematical shorthand notation and was awarded the Turing Prize for this work. His Harvard APL courses might well have been the first

computer science classes ever held. Iverson valued conciseness. APL could be thought of as the result of applying to mathematical notation the advice he gave to a colleague for refining a paper, “If it’s a clause, turn it into a phrase. If it’s a phrase, turn it into an adjective or an adverb. If it’s an adjective or an adverb, omit it. And you apply these recursively.”

The language allows the programmer to use single keystrokes of special symbols to do such operations as invert a matrix or sum a series. For example, summing the elements of a vector that are greater than 100 is simply $+/x > 100$. Or take a more complex problem, one that was once posed to Iverson: Find the sum of all elements of a matrix that are equal to the sum of the corresponding row and column indexes. The expression for doing this in APL takes just nine keystrokes and a bunch of pairs of parentheses. The same commands can be applied on a multidimensional matrix with millions of rows as easily as on one row of numbers.

For the mathematically predisposed, APL is a lot of fun to use. And some of the particularly adept users can provide an endless source of party tricks if you happen to go to parties with lots of quants. Of course, the tricks depend on figuring out the right sequence of commands, but that is part of the fun of it. It seemed that every discussion about developing new models or pricing functions degraded into a discussion among the programmers of who could write the solution in the most succinct way, with a “one-liner” being the prized objective.

There was a reason to put things into as condensed a form as possible. APL is an interpretive language, which means that the computer reads the code and translates it into machine language at run time. The interpretation takes time, so the smaller the set of characters used to execute a command, the faster things run. In fact, sitting near almost every terminal was a dog-eared copy of one code book or another containing page after page of one-liners to do almost every conceivable operation, the APL equivalent to the engineer’s book of integrals. But by the time things get to one-liner Nirvana, the code is impenetrable, so obtuse that the only way for even an expert to figure out what is going on is to construct a simple numerical example and look at how it is transformed with one operation after another.

Morgan Stanley not only used APL; it also turned most of the resources of fixed income research into a cottage industry for creating a new and improved version of the language. Joel Kaplan, who moved over from the Analytical Proprietary Trading (APT) group to run fixed income re-

search sometime after Platt's exit, was an APL fellow traveler. He once said, "My first name starts with J and my last name starts with K, followed by APL. So you know where I stand." (K and J are derivative languages of APL.) He viewed APL as a secret weapon that could allow us to develop more powerful applications and do so more quickly than the competition.

Arthur Whitney had already created a version of APL that he simply called A, written in C with a coding style that was nearly as compressed and opaque as the end objective. This was then extended by a bright and zealous team into A+, a language that continues to run some applications at Morgan Stanley.

In the hands of the mathematically gifted who rotate n -dimensional spaces in their sleep, A+ proved to be a powerful tool. But it had limitations for a class of problems of particular importance for risk management: derivative instruments. For all its intellectual challenge and compact elegance, APL's asset is also its drawback: As an interpretive language, it allows for quick development, but if a program requires a loop, where the same lines of code are reread and reinterpreted thousands or even millions of times, things grind to a virtual standstill, at least by the standards of computer speed for compiled languages.

And almost everything we were doing at that time—at least almost everything that had enough risk to worry about—required loops to solve problems efficiently. For the explosion of new types of derivatives and swaps, getting a good estimate of how the prices of these instruments could vary under different scenarios—a question from Risk Management 101—required positing a wide set of paths for interest rates and then marching the pricing of the instrument along each of these paths. Often the price at any point in time depended on where prices had already been. And prices also varied in nonlinear ways. Plus, for most options the alternative approach to pricing was recursive, meaning that you start with the price at the time of expiration and then move backward in time, which again required moving through many paths, only with the clock going backward rather than forward. And then there were some cases where you had to move both forward and backward. If you program it using loops it is not difficult to do this pricing, and, except for the time it takes, the problem does not look any different if you have 20 paths with 10 nodes each or a million paths with a thousand nodes each, although the result will obviously be far more accurate in the latter case.

Derivative and swap pricing problems landed on the weakest point of APL. I worked with the APL crew posing the problem in different ways. They would try to structure the problem in matrix form, but there was not enough memory to hold a matrix of high enough dimensionality to solve the problems with an acceptable level of precision. These problems were a headache for me and never really did get resolved before I left Morgan Stanley the following year. But they ended up being a disaster for another FIR expatriate.

OUT OF THE LOOP

Sheldon Epstein worked in the FIR insurance group headed by Jim Tilley. He moved from FIR to the trading side of the business at Merrill Lynch, where he was credited with the development of an esoteric financial structure called an index-amortizing swap. An index-amortizing swap is a yield curve swap that pays out the short-term interest rate in exchange for receiving the long-term rate, and has an additional wrinkle in that it amortizes or matures at a rate that is indexed to the short-term interest rate. If the interest rate drops, the swap matures more quickly. The attraction of this instrument is that its pattern of returns, and hence its price, looks a lot like that of a mortgage. Like the index-amortizing swap, a standard 30-year mortgage bond will pay off more rapidly the lower rates drop, because more of the homeowners behind the mortgage bond will prepay in order to refinance at a more attractive rate. But while the index-amortizing swap behaves like a mortgage-backed security, the swap has one very attractive advantage. A swap is not a security, but a contract between two parties, so the bank does not have to pay anything up front when it enters into a swap. If the bank has the choice of reducing its available capital by \$100 million to buy a bunch of mortgage bonds or reducing its capital by zero dollars to enter into a swap, the swap wins.

Epstein easily found willing parties for his index-amortizing swap. But each swap that was put on with a bank added a position to his trading book. If the bank received the long-term rate in exchange for the short-term rate, then he had to take the other side of the contract and receive the short-term rate in exchange for the long. Epstein made his money by charging a spread for the swap. If the long-term and the short-term interest rates were both at 5 percent, he would agree to a swap that would give

the banks only 4.9 percent in exchange for 5. On a \$100 million swap, this differential would be \$100,000 a year.

The problem was that rates change. If the long-term rate suddenly rose to 6 percent while the short-term rate stayed fixed, Epstein would pay out a lot more than the initial \$100,000 spread he built into the trade. So for every swap he put on, he had to structure a hedge that would pay off the same amount he might lose if interest rates moved against him. As with mortgages, structuring this hedge is not child's play. The complication comes from the interaction the long-term and short-term rates have on the maturity of the swap. If the short-term rate drops and the long-term rate rises, then the swap will pay off quickly. And this will happen just when the party who is receiving the long-term rate wishes the swap would keep going on forever, because he is getting a higher long-term rate in while paying a lower short-term rate out. Conversely, if the short-term rates rise and long-term rates drop, then the party at the long end is stuck holding the swap while making higher payments (because these are based on the short-term rate) and receiving lower payments in return. In either case, the amortizing feature works against the party who is receiving the long-term rate—the party who is the equivalent in the swap of the mortgage bondholder. The unattractive feature is designed as intended, because that is also what happens with the holder of a mortgage.

Because of this interaction between the maturity of the swap and the relative level of rates, the value of the swap depends not just on the general level of rates, but also on the potential for shifts between the short-term and long-term rates. That is, the correlation between the short- and long-term rates is critical. It turned out that in pricing and hedging these swaps, Epstein did not take into account the impact of the correlation. This failure was not an intellectual lapse, but a programming one: The same features in APL that caused problems for the recursive calculations also made it difficult to generate correlated price paths for interest rates.

APL was a popular language in the actuary community, and as Epstein had an actuarial background it was natural that he would write his pricing model for index-amortizing swaps using APL. The problem was that, like the recursive derivative problems I faced at Morgan Stanley, building a model to take into account the interaction of the short-term and long-term bonds requires a lot of loops. The interaction can be effectively analyzed only by running many different what-if scenarios of interest rate

paths. Each scenario requires a loop, and within each of these scenario loops there is yet another loop required to follow the path of rates over time. To do a decent job might require looking at thousands of paths, with each path repricing scores of times as the swap moves forward in time toward maturity. This could mean reinterpreting the code millions of times, which is prohibitively time consuming. Short of committing the heresy of rewriting the whole program in another computer language, the only expedient course of action Epstein could take was to ignore the correlations.

Naturally, ignoring a prime source of risk makes for a rosier picture, so Epstein priced index-amortizing swaps far cheaper than anyone else. This did not sit well with his manager at Merrill Lynch, who had enough math skill to know that things were “a little crazy,” and was unconvinced by the economics of Epstein’s creation. So Epstein packed up and headed to the Union Bank of Switzerland (UBS). Like Morgan Stanley, UBS was an APL shop; Keith Iverson moved there from Morgan Stanley and eventually recruited a number of other Morgan Stanley APL stars to work their magic. (APL turned out not to be much of a success at UBS, either. Iverson spearheaded an ambitious project for UBS called CORE, which was intended to link together all of the trading areas of the firm to facilitate middle office functions like risk management. But even with a hundred programmers and a total cost that edged into eight figures, CORE never was fully implemented.)

As the market for the swaps expanded, more and more trading desks made the investment to build two-factor models with both a short-term and a long-term interest rate. Epstein’s one-factor model, which assumed that there was only one interest rate governing bonds of every maturity, had UBS pulling away from the pack. He was winning deals—and instant revenue—at every turn. While the risk of the swaps and their final implications for the firm could be known for certain only after they matured, profits from the transactions were booked immediately—the prices of which were calculated using Epstein’s own model.

Other dealers were mystified by UBS’s aggressive bidding. One competitor was quoted as saying, “Either he has figured out a new way to price options or he is wrong.” I don’t know whether Epstein simply hoped to sneak through or thought that the market would not show the holes in his model. But the market certainly did not cooperate. The surprise Federal Reserve action in February 1994 caused a jump at the short end of

the yield curve relative to the long end, a shift that completely exposed Epstein's model. And it was worse than that. Epstein had hedged the swap positions using options, but had done so in a way that turned out to be less than ideal. Some of his options moved into the money, increasing in value, while others moved out of the money. Just as the rates in different parts of the curve moved in different directions, so did the volatility on the rates. And on each score, it seemed the result was adverse to the hedge. So UBS lost money not only on the swaps, but also on the hedges that were intended—albeit imperfectly—to protect the swap book. The losses were accentuated by another little quirk of the APL code: It did not have a financial calendar that matched the calendar used by the UBS back-office system. The result was that on occasion—especially over holiday periods or when Sheldon was out of the office—options would expire without being exercised.

By the time the bank got around to pricing Epstein's book by something other than his own model and had shown him the door in May 1995, the loss on the book was somewhere in the range of \$100 million to \$250 million, although it is difficult to get an accurate read on these sorts of losses because the traders who take over the desk have an incentive to push the mispricing as far as they can, so that when they finally liquidate the positions more revenue is attributed to them. In any case, this loss was not off the charts by the standards of the day, and it was exaggerated in the press and perhaps unwittingly by the firm itself. But it made a huge impact on UBS, a Swiss bank shrouded in an ethos of conservatism where these sorts of things didn't happen. The Swiss precision was faltering—and there would be more to come. In a few short years, UBS would fall prey to two more risk management errors. The cumulative effect would be the merger of UBS and the Swiss Bank Corporation in 1998 to form UBS AG.

After a few years of Keystone Cops-like crises, with firms tripping over newfangled financial instruments, in some cases abetted by their own traders—Wall Street looked like it had learned its lesson. Every investment bank now had a cadre of risk managers, a flow of daily risk reports, and an organizational structure that included high-level risk management committees. Annual reports started to discuss risk management and extolled the risk management platforms the firms had thrown into place. International commissions brought conformity to risk measures and used them to define acceptable capital levels.

A DEMON OF OUR OWN DESIGN

All of this should have helped, but it didn't. Or maybe it made things better than they would have otherwise been—from my vantage point in the middle of this process I hope that was the case—but whatever it did was not enough. The problem is that no matter how big a team you put on the case, you can only run risk models and pound out reports on things that you can identify and measure. In fact, the bigger the team, the more likely it is that you will get increasing detail on what has already been identified and little on anything new. The types of risks that could be readily measured were better controlled, but those were not the risks that mattered. The real risk is the one you can't see.

CHAPTER 4

HOW SALOMON ROLLED THE DICE AND LOST

One of the dumber things I did in my Wall Street career occurred before I even got started. I turned down Marty Leibowitz in 1984 when he tried to recruit me to Salomon Brothers. He'd gotten wind of the effort by Lehman Brothers and Morgan Stanley to enlist me, and tried to sell me on Salomon. At the time he was well on his way toward his current stature of a legend for his prolific and wide-ranging research. I chose Morgan Stanley for reasons that turned out to be pretty stupid. Then again, I was an academic who was clueless about most of what was going on in industry. He and I would meet frequently over the next decade and he tried to find a place for me at Salomon, realizing more than I that Solly was a far better environment for me than was Morgan Stanley. Once I looked at a position on the client research side. Another time I had a dinner with Eric Rosenfeld, Greg Hawkins, and others about the fledgling

proprietary trading group that would be the fabulously successful precursor to Long-Term Capital Management—an opportunity that for yet some other stupid reason I did not follow up on.

But when Marty called yet again to discuss the role of risk manager at Salomon, somehow I managed to say yes. In 1994, Salomon was the biggest risk-taking firm in the world. It was at the pinnacle of the bond market—even after having its wings clipped by a Treasury-auction scandal that had cost the jobs of Gutfreund, John Meriwether, and several others in senior management. The high-powered, big-brained proprietary trading group had become the envy of Wall Street. John Meriwether and the bulk of his team had just headed off to start LTCM, but the principals who remained continued to trade in the same size.

And along with the pull from Marty, there was a good reason for me to look outside Morgan Stanley. Bob Feduniak, who had started the risk management division, followed his mentor Lewis Bernard out the door, leaving the fledgling division in flux. The company tried to entice Barry Allardice to return to take on that role. Allardice had retired to devote his time to sailboat racing, and I went to Florida as part of the effort to bring him back to New York, spending a day sailing with him on one of his trial runs and trying to keep from sliding off the deck. Amazingly enough, Allardice declined the offer to get off the boat and return to the politics of Morgan Stanley.

Going to Salomon was like moving from a lumbering cargo plane to a fighter jet. Salomon Brothers was not like other firms on Wall Street. In fact it was not even like Salomon Brothers—at least not the Salomon Brothers depicted by Michael Lewis in *Liar's Poker* or envisioned by other trading floors. Maybe things had changed in light of the Treasury scandal, but the trading floor, though bustling, was quiet: no shouting, no swearing, no slamming phones. The atmosphere was reasoned and intellectual. Where discussions at Morgan Stanley seemed to be a concatenation of sound bites, at Salomon things were thought out. While at Morgan Stanley there was a hierarchy that demanded wending down the right path to bring out ideas, at Salomon a vice president who disagreed with the head of a trading desk would just walk up and discuss things. If what he said made sense, that was how it was done. Turf and status were trumped by the primacy of ideas.

HOW SALOMON ROLLED THE DICE AND LOST

Perhaps that sounds Pollyanna-ish, but I saw it happen regularly. At one risk management committee meeting, I watched Rob Stavis, who headed the U.S. proprietary trading group, suggest that the solution to a derivatives problem was to have that particular trading book moved from the client desk to his group. After the meeting, as I headed down to the trading floor, I noticed Rob ahead of me. He walked over to Rick Stuckey, the head of the derivatives desk, and let him know what had been discussed. Stuckey simply nodded and suggested they get together after the market close to discuss it more. And that was that. It would not have gone down like that at any other firm. Stavis had basically put forward an argument to take away part of Stuckey's business, and did it at the highest level of the firm; the chairman and the CEO both were part of the risk management group. At any other firm, Stavis would have kept things close to the vest until it was a done deal. If Stuckey caught wind of it, he would have been livid and looked for an opportunity to ambush Stavis. But at Salomon the question was one of what made the most sense for the firm as a whole, and was analyzed dispassionately.

The lack of fiefdoms was part of Salomon's history as an outsider. Entering the top tier through the servants' entrance as a bond trading house rather than through the white-shoe connections of the old-line investment banks like Morgan Stanley and Goldman Sachs gave Salomon an us-versus-them attitude. People simply did not push their own agendas if doing so would be damaging to the firm. It didn't take me long to see that Leibowitz had been right; I felt right at home at Salomon. Sadly, in just four years Salomon would disappear in all but name.

While the trading was smarter at Salomon, the company's trading size and risk posture had a corresponding downside: The problems turned out to be bigger. Morgan Stanley had had a few errant traders here and there; Salomon was in the throes of major risk problems. If I had come to Salomon a few months earlier, I either would have been a hero or would have been fired for missing any number of time bombs. The firm was just starting to see the effects of the Federal Reserve's February 1994 rate hike.¹ The foreign exchange book, which had generated \$100 million in profit the year before, was suddenly down by that same amount. More critically, Salomon's mortgage book was hemorrhaging.

THE ROOF CAVES IN ON MORTGAGES

Few homeowners realize that when they obtain or refinance a mortgage they are participating in the most complex of all financial markets. Most mortgages are packaged by banks (which don't want to tie up capital) and sold off as mortgage-backed securities (MBSs). The reason for the complexity of mortgages is simple: Homeowners have the right to prepay their mortgages. This prepayment option means the lender has continual uncertainty about when it will receive back its principal. There are a number of reasons a homeowner will prepay a mortgage. Obviously, a home sale will generally lead to mortgage prepayment. Home sales may occur for life-cycle reasons—a growing family on the one hand and retirement on the other, or because of job changes or divorce. These triggers of prepayments are not so difficult to manage, because life cycles have a demographic, even actuarial, predictability.

A far more troublesome and less predictable reason for prepayment is that a change in interest rates alters the attractiveness of refinancing. Prepayments because of interest rate changes work to the MBS holder's disadvantage. The mortgage will get prepaid when rates drop, the very time the higher coupon of the mortgage will be more valuable. Conversely, if interest rates go up homeowners will hold onto their low-interest mortgages for dear life at the very time the mortgage holder wishes it could get its hands on the principal to reinvest at the higher rates.

The mortgage market is made more obtuse because mortgages are carved up and traded in different component parts. This is done to try to create some mortgage-backed instruments that will not be affected by prepayments. Mortgages have little in the way of default risk, so if you can get rid of the prepayment risk you have a bond that will feed a huge market. The sell-side firms created a whole alphabet soup of collateralized mortgage obligations that split the mortgage payments and prepayments up into different tranches. Some of the tranches were very stable and predictable, which of course meant the prepayment risk was swept into other tranches, which were almost radioactive. One of the cleanest and most intuitively appealing ways to break up a mortgage was developed by Richard Roll and his team while he worked at Goldman Sachs in 1985. Roll divided a mortgage into two pieces: a security that paid the holder the principal of the mortgage when received (a PO for "principal

HOW SALOMON ROLLED THE DICE AND LOST

only”) and another security that paid the holder the interest payments (an IO for “interest only”).

The PO is essentially a bond where the timing of the cash flows is uncertain. The holder of the PO will receive a known set of principal payments, the question being whether the payments will come early, late, or on time. Oddly, whether the payments come in sooner or come in later, any deviation is bad news for the PO holder. If interest rates drop, spurring faster prepayment tied to refinancing, the holder of the PO will receive the principal payments more quickly. A rate increase, by contrast, will slow down prepayments and the investor will wait longer to receive most of the principal payments. So at the very time investors would like to have the bonds (when interest rates and therefore reinvestment opportunities are less desirable) they end up with the proceeds of the bonds to reinvest, and at the very time they would like to have the proceeds to reinvest because interest rates have risen, those proceeds are slower in coming.

The other half of the package, the IO, is even more peculiar. While most bonds increase in price as interest rates drop, the IO drops in price with the drop in interest rates. Not only that, it is possible to actually have a negative return on an IO, a rarity for fixed income instruments (excluding defaults). If the prepayments occur too swiftly, the present value of the coupons that are received before the mortgages are paid off may actually be less than the initial investment.

The complexity, counterintuitive pricing, and volatility of the IO make it a prime target for hedge funds and proprietary mortgage traders. The IO essentially contains a concentrated version of the so-called toxic waste that runs through a mortgage security’s prepayment option. Because it is risky, most investors find it undesirable, but because it is an inevitable component in the breakdown of the mortgages, it has to find a home. So the mortgage market provides a dumping fee to find buyers for the IO. If it is priced and hedged correctly, the IO holder will make a profit—a profit that in percentage terms is small but that can be levered up to become significant. But like the trail to the forty-niners’ gold rush, the financial landscape is littered with the corpses of firms that have tried to trade IOs and failed.²

That’s what happened to Salomon Brothers in 1994 when it lost well over \$100 million from its mortgage position because of the Federal Reserve rate shock in February. That same year, for much the same reason, a

major investment fund, David Askin's Granite Fund, blew up, losing its investors some \$500 million. In Askin's case, it was because the firm's highly touted risk and hedging model did not really exist, and the fund was naked to the effects of the rate changes. In the case of our trading desk and those at other firms that were similarly hit, the simplest alternative—selling off the position—was not available. That was because the trading desk's franchise depends on standing ready to make a market in whatever is thrown at it. When things go bad, the trading desks are treated like dump sites, and like real toxic waste, there are only so many places available to store it when it comes trucking in. If you are trying to dump waste at the same time as many others who are holding too much waste, the dump sites fill up and opportunities close down. And you have to stand by while the firm loses a bundle.

THIRTY MILLION OVER TOKYO

Global companies transfer their know-how far and wide, and Salomon was no different. Spurred by the early successes of its U.S. proprietary trading operation, Salomon created similar units in Europe and then in Japan. During the 10 years of its existence, from 1988 to its closing in 1998, the Japanese unit contributed in excess of \$2 billion to Salomon's bottom line. It also contributed two big headaches that cost the firm nearly a quarter of that amount. The reason? You can lay your trading template over a foreign market like Japan, but at some level it doesn't necessarily fit the culture, the technical ability, or the peculiarities of that market. We thought we had a money machine in Tokyo, and for a while we did. But sand had been collecting in the gears, and when the whole thing ground to a halt, they called me in.

For the bulk of Salomon's time in Japan, its fortune was delivered by two strategies. The first was a swap spread trade. The Japanese were hungry for yield, because their government bonds paid piddling rates, and they were willing to take on credit risk to increase the yield of their fixed income portfolios. This meant buying LIBOR-denominated debt rather than Japanese government debt.³ The LIBOR rate is higher than the government rate, because there is always the risk that a company will default, so the additional risk has to be compensated.

Because they already had government debt on their books, the banks would enter into contracts to swap the returns on the government debt for

HOW SALOMON ROLLED THE DICE AND LOST

the LIBOR debt. So with the swap they would pay out an amount related to the lower-yielding government bond and pick up a payment equal to the higher-yielding LIBOR bond. The huge demand for LIBOR knocked down the differential between the government bond rate and the LIBOR rate to very low levels. Often the difference was less than 10 basis points (one-tenth of 1 percent), and on a few occasions the technical imbalances of supply and demand caused the LIBOR rate to drop below the government rate.

Absent some institutional barriers or tax issues, parity between the LIBOR and government rates was simply not sustainable. There had to be a parting of the rates. No one would be willing to buy the bond of a bank that could default for the same yield as a Japanese government issue. So Salomon would put on a swap spread in ever-increasing size to buy Japanese government bonds and sell short LIBOR-denominated bonds when the spread neared zero, and gradually unwind it as the spread moved back up. At times the trade would top \$20 billion, producing a profit of \$20 million for every one-hundredth of 1 percent the spread widened.

The strategy was embarrassingly mechanical; the size of the trade was set as a simple function of the size of the spread: the tighter the spread, the bigger the trading size. This was the recipe for printing money, and it was laid out on a sheet of paper taped on the head trader's quote screen. Each day the trader, Tanaka-san, would look to see where the swap spread was trading, then adjust the position to match the size shown on the sheet.

This trade was not only simple; its risk was well-delineated. Because the spread could not stay negative, the most that could be lost in the short run was the tiny spread. (Contrast that to the type of spread trades that led to the LTCM disaster: the bet there was for spreads to narrow, the risk being in the event of spread widening. The sustainable limit for narrowing is parity; there is no limit for how wide the spreads can become.)

The second strategy was the convertible bond (CB)/warrant book. Warrants are simply long-dated call options issued by a company. If the price of a company's stock exceeds the warrant's exercise price at expiration, then the warrant holder can pay the exercise price and get the stock. If the stock price is below the warrant's exercise price at expiration, then the warrant expires worthless. Convertible bonds are bonds that have warrant-like behavior. If the stock of the issuer goes above a conversion level,

then the bondholder gives up the bond and receives stock instead. If the stock stays below the conversion level, then the bondholder continues to hold the bond.

Japanese companies began a love affair with these two securities in the mid-1980s. They thought of warrants and CBs as ways of issuing cheap debt. Just as they did not appreciate the implications of credit risk, the Japanese companies did not understand that ultimately, once the warrants were exercised and new stock was issued, the debt they had issued would finally end up costing them money. They thought the cost of issuing more stock was close to nil. Only after the fact did they discover, quite painfully, that the stock they were forced to issue to the warrant holders diluted their existing equity holdings, driving down the prices of the companies' stock. The underwriters probably understood this, but they were not talking. The warrants also allowed the Japanese firms to reduce their taxable income by the amount of the warrants. So they got a short-term tax benefit but ate the cost of the issues down the road.

But the Japanese firms issuing these warrants failed to price them correctly. The warrants were cheap when they hit the market. And they flooded the market beyond the natural demand of domestic investors, so they stayed cheap. Armed with the tools of option pricing theory and the capital of the firm, our Japanese arbitrage unit bought up warrants and CBs. The trading desk then went about hedging the market risk of the positions. These instruments, though cheap, could become money losers if the stock of the issuers dropped. To guard against a drop, the trading book shorted the underlying stock. As with other option-related hedges, the amount of the short position varied according to the level of the stock price and the time remaining before the warrants expired. If the hedge was executed correctly, the profit of the warrants and CBs would be locked in, impervious to variations in the fortunes of the issuer, although the desk might have to wait until the issues matured to realize the profit.

Day by day, as the warrants came closer to the expiration date, more of the warrant value could be booked. The value amounted to \$200 million in an average year. Like the swap position, the warrant book was monitored by a single trader, who simply looked at computer printouts each day to determine how much to adjust the hedge. But the CB/warrant book was more complex because it contained hundreds of individual positions on many different companies with varying maturities and exercise prices.

HOW SALOMON ROLLED THE DICE AND LOST

This complexity was masked by considering the overall portfolio as a single entity and using a stock index—either the Nikkei or the Topix—as the hedging instrument.

While revenues roared, the lion's share of the compensation accrued to the unit's boss, Sugar Myojin. From the late 1980s to the mid-1990s, his compensation generally put him at the top of the tax rolls for the city of Tokyo. The rolls were published in the spring, compelling him to take refuge from the press for a week in a set of suites at the Okura Hotel. Sugar lived the sweet life. One year during the betting at the famous Salomon partners' golf outing, he put \$20,000 on a team, declaring loudly in mock bravado, "Twenty thousand dollars means nothing to me!" (Another year, with the arb group rolling in money, the total bet by the partners got up past half a million dollars for teams whose average handicap was upward of 20. Deryck Maughan, Salomon's CEO, fearing negative publicity from such profligacy, cut every bet in half.) Myojin's conspicuous spending was emulated by others in the arb group. He fostered an interest in golf and initiated a Ferrari fetish that flashed through the proprietary trading management team.

By 1997, though, with the business disintegrating, Myojin had other problems than the prying Japanese press. The CB/warrant position inexplicably started to run out of steam in the spring of 1996. The warrants and convertible bonds continued to look cheap and the hedge followed the same pattern it had in the previous years, but rather than posting \$20 million to \$50 million a month in revenue, the strategy was leaking commensurate losses.

No one could figure out why the position's fortunes had changed. The losses were perplexing not only because they seemed to be coming out of nowhere, but also because they were increasing over time. And there was no prospect of simply selling off the positions and calling it quits; the very reason the positions were profitable was that there were not many other investors willing to buy these instruments. Like a ship taking on water in the middle of the ocean, we were stuck with the positions, and if the losses kept on growing, all we could do was watch as the ship foundered.

The firm sent me to Tokyo with the senior members of the U.S. proprietary trading group, Rob Stavis, Dennis Keegan, and Andy Fisher as a SWAT team to try to understand the source of the mounting losses. What almost immediately emerged was a decision tree of potential problems.

The model used to price the warrants was overly simplistic. The warrants were so cheap to buy that Myojin had not seen the value in hiring analytical talent to get things right to the last penny (or even the last nickel). To make things easier for the desk, the hedges were done using index futures rather than the individual stock, so a tracking error could have developed based on discrepancies between overall market behavior and the behavior of the specific stocks. Further analysis showed that the stocks in the warrant book seemed to underperform the index. Not only that, they seemed to underperform other stocks in the same industry, especially as the market rose.

The valuation model, suspect though it was, also did not show them to be as cheap as they had been historically. They had not risen in value, but the model-based measure of their fair value had dropped, and had dropped for a very strange reason. Just as these stocks systematically underperformed, the warrant prices were lower because the stocks underlying the warrants had lower volatility than other comparable stocks. A warrant is worth more when the underlying stock has higher volatility because you get to enjoy any upside surprise while being insulated from the negative surprises. So the drop in volatility had reduced the warrants' fair value and thus the amount of profit that we could expect to extract from them even if the other problems were surmounted.

There were problems everywhere we looked. The models were less than exact, the hedging instrument was less than ideal, the upside moves in the warrants were sluggish, and the margin of cheapness in the warrants had nearly evaporated. With so many moving parts it was hard to know where to start to fix the problem. And the ship was taking on more and more water; by our calculation, in the one week we had spent there the portfolio had lost another \$30 million.

Peeling away layer after layer of the complex position, we homed in on two tracking errors, errors that were harbingers of mistracking in Treasury bonds that would prove fatal to the U.S. arb division and that would act as catalysts for the near-cataclysmic LTCM crisis barely one year later. One was a tracking error that affected the ability to use the stock to lock in the profits on the warrants. The other was a tracking error between the equity index futures and the stocks that altered our ability to use the stock index as a backup instrument for hedging.

The changing competitive field complicated our dilemma, because Salomon was no longer alone in taking advantage of what seemed the clos-

HOW SALOMON ROLLED THE DICE AND LOST

est any trader could get to free money. A number of other—mostly foreign—brokers and banks had bought up warrants and were following the same hedging strategy. The result was that as the stocks rose, we were accompanied by lots of other traders who were selling stocks as a hedge, and similarly, as they fell there were many others who were buying stock back to reduce their shorts. Naturally, the more people sold as the stock rose and bought as it fell, the more the price movement would be dampened. Because this was only being done for the stocks with warrants, they were left behind on market rallies and declines.

Unhappily, the nature of the dynamic hedge was to hedge more aggressively when the stock price rose, so the dampening was more pronounced during rallies. There was yet another negative, which related to the stock float. Much of the stock of the issuers, especially the banks, was locked up and never traded. Many of the larger financial institutions in Japan had cross holdings of each other's stocks for mutual protection. So, only a fraction of the total shares outstanding were available to trade. With so many traders trying to find a particular stock to borrow to short, the price pressure was far greater than it would have otherwise been.

It seemed that it couldn't get more complicated than that. But it did. The price action not only made the stocks underperform the index; this was also the likely culprit for the dampening of the stocks' volatility. The actions of the hedges kept the stocks from going up and down as much as they would otherwise.

What had seemed like a money machine churning out revenue with little intervention was laid open for what it really was: a morass of hundreds of positions that were difficult to model for hedging even in the best of worlds, and that were convoluted by the interactions of market feedback from the unobservable actions of other traders. There did not seem to be a way out; in fact, it would only get worse because, for technical reasons related to the warrants' gamma, as the warrants got closer to maturity the demand for hedging and for changes in hedge positions would become more frequent and exacting. We had to make a choice between a couple of lousy options: hedge using the futures and face the inevitable mistracking of the futures versus the stock price, or hedge using the actual stocks and go into the fray with all the other traders.

Although we did come to terms with the risks and the nature of loss we had identified, there was no well-structured exit plan. Things just

fizzled. Because the choices of staying in the position were to face higher risk on the one hand or the prospect of lower return on the other, a natural third option to consider was to throw in the towel and liquidate. Come to think of it, running away would have been the correct action. In convergence trades like this one, good things do come to an end, and it was clear that the market had awakened to the warrants' prospects.

But liquidating would have been costly, and it would have required getting out of a substantial part of the unit's business. In fact, since the Japan arb unit was built on two strategies, it would have left Myojin with a sizable group monitoring and adjusting a single trade, the swap, which itself was facing a limited future and could be run off of a sheet of paper. Perhaps that played into the inertia. It's hard to know. The Japan arb group did a little of everything, reducing the position as best it could and using futures and stocks to hedge what remained. The play had passed its period of profitability; whatever we did was designed to stem the potential for loss, which was sizable indeed. And when the Japanese convertible market finally broke, it caught a number of firms, including UBS, with huge losses. The actions we took, though less than decisive, spared us from the brunt of it.

These weren't bet-the-farm risks. Far from it. They were typical, and just a warm-up for the next two episodes, which would concern in the first instance the viability of Salomon itself, and in the second, the viability of the financial markets as a whole.

SHOTGUN MARRIAGE

Bob Denham was installed as the chairman of Salomon by Warren Buffett after the Treasury scandal in 1991, the same time Deryck Maughan was installed as the new CEO. A Harvard Law School graduate and partner at Munger Tolles & Olson, the law firm of Buffett's right-hand man Charlie Munger, Denham's primary job was to make sure the Salomon traders stuck to their knitting—a final dollop of oversight for the tamer postscandal Salomon. Tall, fit, and boyish, with a hint of an accent from growing up in Lubbock, Texas, Bob was as even-tempered as his predecessor, John Gutfreund, was erratic. Because he worked in securities law, he was not unfamiliar with the business of Salomon, and he enlisted the help of Eric Rosenfeld and others in the proprietary group to quickly get up to speed on the more sophisticated trading strategies of the firm.

HOW SALOMON ROLLED THE DICE AND LOST

Denham occupied Gutfreund's office on the 43rd floor, complete with Gutfreund's old furniture. There were backlit maple bookshelves on the two side walls, contemporary leather chairs and sofa, and a chrome and leather custom desk on the outside wall. For some reason, Gutfreund's windows were decorated to minimize one of the most spectacular views in Manhattan, looking north toward the towers of midtown from Salomon's headquarters at Seven World Trade Center. The desk itself sat against a wall that had been built to cover one of the windows. Denham's few personalized touches came through a smattering of books and photographs—including a picture with Bill Clinton, who had appointed him to be the U.S. representative to the Asia-Pacific Economic Cooperation (APEC) forum. On the bookshelf near the door stood a tenth-scale model of a red Dodge Viper given to him by a client.

His office was a decidedly uncomfortable place to be in August 1997, when I sat down on one end of the couch, soon to be joined (by phone or in person) by the key traders of the firm: Rob Stavis, head of the U.S. fixed income proprietary trading group; Andy Hall, head of Phibro, the commodities trading arm of Salomon, who was conferenced in from Westport, Connecticut; Costas Kaplanis, Rob's European counterpart from London; and Sugar Myojin, joining in from Tokyo.

The last one connected by phone was Andy Parets, the trader at the center of the discussion. Parets ran the risk arbitrage operation, meaning he took positions on companies that were doing the merger dance and made money when the merger was completed. There isn't much mystery: On the one hand, the prices of the companies involved in a merger nearly always converge when it is completed. On the other, there is always the risk that the merger will not be completed, so the stocks do not match in their relative prices. The key to Andy's trading was to focus on companies whose merger likelihood was higher than the probability implied by the relative prices of the two companies. The overhanging risk was for a deal to fall apart and the stocks move back out of line, which was exactly the threat looming on one of Parets's positions—one that happened to be several-fold larger than any he had put on before.

Earlier that year MCI Communications and British Telecom inked a merger agreement, and Salomon's risk arbitrage group had gradually been building up a trade to profit from it. The risk arb trade was short BT and long MCI, with an expectation that the two stocks would converge.

The obvious risk was that if the merger didn't happen, MCI, untethered from BT, would sink in price. This MCI/BT trade had increased throughout the spring until it had grown into a position of about \$500 million on either side by the summer. Lurking uncertainty concerning the viability of the merger, the problems in executing the two legs of the trade (set as they were in two different markets), or possibly the inability of many investors to do simple math had kept MCI's price under the value it would attain once the BT deal was completed.

A LESSON IN SELF-DELUSION

In Parets's view this trade was a home run. There appeared to be no impediments to the consummation of the merger, yet MCI persisted in having a discount relative to BT. Because the limit for the position was periodically extended, Parets had to appear before the risk management committee to explain the performance, prospects, and perils. It probably frustrated him, since none of the committee members had experience in equity markets, much less in the esoterica of risk arbitrage. But we could count, and by July 10, the MCI/BT trade had grown to a position that was long 11 million MCI shares and short 62 million BT shares.

After the market closed on July 10, MCI dropped a bomb, announcing that because of higher-than-expected costs to enter local markets it anticipated incremental losses of up to \$400 million in 1997 and \$800 million in 1998. The next day MCI's stock dropped 15 percent to \$36. BT lost about 5 percent, to \$7.50. Salomon's position nose-dived from a small profit to a \$35 million loss. The abrupt reversal precipitated the first of a number of meetings specific to the problems of the trade, where we worked through a tree of possible results that led, on one of the branches, to the one we were currently observing.

The tree started with three large branches, outlining three broad future scenarios for this trade. The most optimistic envisioned the transaction would be completed as originally planned and the share prices would converge. The result would be a profit of almost \$6 per MCI share, netting a total profit to Salomon from the July 11 close of \$75 million. The second was a renegotiation of the transaction, with a 10 percent reduction in the price per share. The profit per MCI share would then drop to just under \$2 a share, but the result would still be positive, around \$30 million.

HOW SALOMON ROLLED THE DICE AND LOST

The worst-case scenario—the one that we were now facing—was that the transaction would be terminated. If that happened, the result would be hard to quantify, because MCI could then go into free fall. Given the new, adverse information about its future earnings, MCI stock could plunge far below the premerger price. Parets—calling from his vacation house on Martha's Vineyard because he had suffered a case of appendicitis two days earlier—figured there was potential for a further drop in MCI of just over \$9 a share, and a \$100 million loss from the position's July 11 mark.

Parets argued against the likelihood of the merger terminating on several grounds. First, he said MCI was critical to British Telecom; aborting the merger would signal a severe setback, if not an end, to BT's plans for global expansion. He also downplayed the costs at MCI, huge though they were, as being discretionary. Much of the loss would be attributable to MCI's decision to accelerate its investment in the local loop, a decision that BT, with four seats on the MCI board, must have known was coming. So the expected losses to MCI could hardly have been a surprise to BT, and indeed through its board position BT must have acquiesced to these expenses. It would therefore be strange for BT to now rethink the merger just because the costs had finally become a public reality. Parets's final argument against the likelihood of a termination was that BT could not back out even if it wanted to.

The arb group rarely invested in a prospective merger that was not highly likely, and it saw this merger no differently. The regulatory hurdles had been passed, and the merger agreement specifically exempted "legal or regulatory changes affecting the telecommunications industry generally" from the "material adverse change" clause, the standard way out of any merger agreement. Because the expenses could in large part be attributed to a decision by the Eighth Circuit Court, BT could not apply the material adverse change clause to these expenses.

The nature of risk arbitrage is to win small and bet big. The idea is to extract a small dollar differential per share with a high probability of success. By taking a very large position, the two or three dollars per share you make if things go well still lead to a decent profit. Should things go wrong, of course, the prices can blow apart and the size of the bet multiplies the loss. So it is critical to have a very high probability of success—a lock, as gamblers like to say. Despite the surprise losses at MCI, Parets put that probability beyond 99 percent.

This was something that Hall—who had no experience in risk arbitrage or familiarity with MCI or BT—found nothing short of ludicrous. Hall was a competitor, with a lean, athletic six-foot, five-inch frame. He rowed all four years at Oxford in the sixth position, the power position, of the eight-man crew. Oxford got beaten by Cambridge all four years, which led him to quit rowing for the next 20 years. He began again in his forties, competing internationally in the 35+ age group. He stroked in the four-man, won the U.S. nationals for three years in a row, and then moved to the double scull, winning nationals in that as well. Hall came to Salomon by way of Phibro, which in turn came to be part of Salomon by way of a reverse takeover.

Phibro acquired Salomon in 1981 for about \$600 million in order to help Phibro out of the concentrated and cyclical niche commodity business. Recession and low inflation hurt the commodities business at the same time the huge volatility in the interest rate markets turbocharged profits at Salomon. Phibro was bloated with more than 3,000 employees, so as the revenue moved more and more toward the Salomon side of the new firm, Gutfreund successfully demanded that he run Phibro-Salomon. He sold off operations, cut the personnel by two-thirds, and took Phibro off the firm's name, turning it into just one of several trading divisions.

Hall was an expert trader, especially in the oil markets. His career in oil started at age 17, when he joined British Petroleum, and through his employer got a scholarship to Oxford and later the INSEAD business school, where he won the Henry Ford Prize as the top student in his class. Although a member of the Salomon board of directors, Hall was an outsider. He ran the commodities business within a firm that was focused on fixed income; he made his money through outright position taking where the other proprietary units relied on the finesse of relative value trades; and he ran his business from the remote location of Nyala Farms, a farm-turned-office-park in Westport, Connecticut.

As outspoken as he was physically imposing, Hall became a critical counterweight to the rest of the firm. He had no vested interest in the other trading businesses, and if he thought a trading position was wrong, he pulled no punches in making his view known. He looked at the firm's bread-and-butter relative value trading as little more than a gigantically levered spread trade run through a maze of smoke and mirrors. The arb traders in turn viewed his oil and silver trading as nothing more than a giant bingo game in which he took positions alongside doctors, dentists and

HOW SALOMON ROLLED THE DICE AND LOST

other wealthy amateurs, the only difference being the size of the bet. But year after year he brought a thousand times more revenue into the firm than most doctors or dentists made in a good year.

It was a matter of ego for Hall to match the huge trading size and revenues of fixed income arb. Unfortunately, the physics of the markets made it almost impossible. Commodity markets are far smaller than fixed income markets. But Hall pushed the envelope. During the Gulf War he bought all the oil there was to be had. He even leased tankers to hold it offshore. At one point in 1995 he held a silver position roughly equal to a year's demand for the metal. In 1997 he held more than 300,000 tons of cocoa—nearly a year's supply—and made the news when shorts tried to lobby away their losses by unjustifiably accusing him of cornering the market.

In the MCI/BT trade, Hall's argument against the position was simple: The trade did not make sense because the emerging view was that MCI/BT would be a poor fit, and more to the point the merger would be a net negative for BT. Beyond the specifics, Hall repeatedly made one assertion honed by years of trading experience: There is no such thing as a 99 percent probability of being right. There are too many immeasurables, and any probability assessment is itself subject to error. If you are making a trade that needs better than 95 percent probability for it to make sense, you shouldn't do the trade. To make the point Hall laid out an open wager of \$10,000 with 99-to-1 odds that the deal would not go through. Unfortunately for him, he had no takers.

These arguments notwithstanding, after the close on August 20 MCI announced it was renegotiating. And on August 22 BT announced that it planned to reduce the price it would pay for MCI by 25 percent. That meant that MCI shareholders would now receive 3.75 shares of the merged firm plus a cash payment of \$7.75, compared with the previously announced payment of 5.4 shares plus \$6 in cash. So even if the transaction succeeded, the MCI shares were overpriced and their convergent value at closing would net a loss. The stock price moved to reflect the new terms—and the increased likelihood of a termination. Following the news, MCI dropped from \$35 to \$31. Adding to the loss, BT stock, which we were short, actually increased by about 10 percent.

Given this further loss, the desk had already started to unwind some of its position, albeit in unfavorable circumstances. Because of a greater potential for loss in MCI, the desk's first action was to unwind a disproportionate

amount of the MCI position. MCI's willingness to enter into a renegotiation after so strenuously denying its intent to do so was a red flag. It indicated that a material adverse change may have occurred that could facilitate BT terminating the merger. In Parets's view, a termination could have led MCI to drop as low \$20 a share, adding a further loss of at least \$60 million to our trade. Also, a disproportionate unwind was inevitable for the simple reason that the London Stock Exchange closed at 11:30 A.M. EDT, so the desk could not execute any repurchases of BT ordinary shares. The American depositary receipts (ADRs) traded on the New York Stock Exchange were far less liquid.

Between the losses locked in with the unwind and the unrealized losses on the remaining position, the firm was on the hook for \$130 million as a result of the August 20 announcement. If the deal broke, the figure would move closer to \$200 million. What was forgotten during this crisis was that long before the first shot in July Parets had provided the risk management committee with monthly estimates of potential losses in the event of a termination. The most recent projection, from July 3, was a loss of \$130 million, with an additional \$45 million down if the deal broke because of a material adverse change. He got that one right.

The loss was not huge by Salomon standards. Stavis's group suffered reversals of more than \$100 million from time to time, and his was only one of four proprietary groups in the firm. But for Parets it would be hard to take. It was a number that was several-fold more than his group had ever made in a year. The nature of risk arbitrage is to make reasonable money almost all of the time, and only on occasion take a big hit. The key is to amortize the loss over the profitable years—to realize that even though you have not had a bad draw yet, each profitable trade should be discounted by the expected value of the yet-to-be-realized loss.

Whether viewed in the perspective of the proprietary trading profit and loss (P&L) swings or those of Parets's group, the MCI/BT loss was the first of a wave of shortfalls and pullbacks from trading at Salomon. In less than a month, the sale of Salomon to Travelers Group would be under way, the very activities that had propelled Salomon to success now driving it out of existence.

In any case, with more than \$100 million of losses locked in and the potential for \$100 million more if the deal broke, we were meeting in Denham's office to figure out what to do next. Stavis immediately brought the

HOW SALOMON ROLLED THE DICE AND LOST

issue into its purely analytical perspective. While it was hard to come up with precise amounts, the general view was that if things recovered we would end up making \$25 million, while if the deal ultimately broke and MCI went into a tailspin we would add somewhere around \$100 million to our losses.

Running the numbers under the various scenarios, he calculated that our position would have an expected return of zero if there were a 20 percent probability that the deal would break. That meant, given our transaction costs and a willingness to take a zero profit out at this point, that we should not sell the position unless we thought the probability of the deal breaking was greater than 20 percent. Parets insisted that the probability the deal would break was less than 5 percent, and Stavis was bullish enough given these odds that he wanted to put some of the trade on in his personal account. Myojin piped in through the speakerphone, "Well, you know, I have pulled up a lot of articles on this from Bloomberg, and I have been reading through them. I don't know, it seems like this is something that will go through to me." Myojin was pulling the decision-making process down a rung, to a level of discourse not far from polling friends and family for their views. I interrupted Sugar to get things back to Stavis's analysis. We were dealing with a major event, a possible loss of \$100 million, and we wouldn't get any closer to a decision by trafficking in newspaper excerpts and wire reports. We had no more information than Parets had, and he was the one in the group with the most experience in converting that information into a probability assessment of the event.

There were experts both within and outside of the firm who could provide us with a more informed view, but on the whole they did not seem to add much light to the problem—except for Jack Grubman.

JACK JUMPS IN

Grubman, Salomon's telecommunications analyst, was adamant that the deal would go through, and that MCI would harbor no request for a price concession. Grubman had a scrappy, pugnacious look well earned from his days as an amateur boxer. (His father was a former Golden Gloves boxer, and Grubman remained a fight fanatic.) Grubman believed he had an inside track to the secret workings behind the agreement. And in the other corner on this and many other calls was Jack's rival, Merrill Lynch

analyst Dan Reingold. Dan was the antithesis of Jack in many ways, including using reason rather than connections as the basis for his work. Reingold could not fathom how anyone reading the material adverse change clause could not seriously consider that the deal might break. After all, MCI had just announced that its earnings per share would be cut almost in half. We came out on Grubman's side of the argument, which ended up being dead wrong—although Jack still managed to negotiate a \$25 million pay package the following year. And, unfazed by the \$100 million setback for our firm that had pushed us to the wall in our talks with Travelers, after the smoke cleared on the broken deal, Grubman walked in a few weeks later with Act II, a proposed MCI/WorldCom merger.

As would become clear over time, the entrance of WorldCom into this deal—audaciously so, since by relative size MCI should have been eating WorldCom for lunch—was fueled by Grubman's growing intimate relationship with WorldCom and its CEO, Bernard Ebbers. This relationship was the font of his investment banking success. Ebbers, a gym teacher turned motel entrepreneur in Mississippi, soon expanded his holdings to a half-dozen Hampton Inns and Courtyards by Marriott. His entrance into the high-wire world of telecommunications came when he joined two other motel investors to form a long-distance provider called Long-Distance Discount Service (LDDS). In the early 1980s, discount long distance was a local businessperson's game; Kiwanis Club cadres would buy a switching network and advertise in the local papers to build a subscriber base. Undercapitalized and with limited scope, most of these enterprises were bought up or failed outright. LDDS was on the feeding end, acquiring one entity after another: Advantage Companies in 1989, Advanced Communications in 1992, Resurgens Communications Group, and Metromedia Communications in 1993 to become the fourth largest long-distance network in the United States. In the mid-1990s, LDDS acquired IDB, an international network, and was renamed WorldCom. It then acquired Williams Telecommunications, MFS, and UUNet, an Internet backbone. But all of that paled to what was about to emerge. In the aftermath of the planned merger with British Telecom, which failed even after MCI agreed to cut the \$22 billion deal price by nearly 20 percent, Ebbers made a run in late 1997. Initially valued at \$30 billion, his bid rose to nearly \$40 billion during the ensuing war with GTE.

In any case, the Stavis approach was to decide if, given the probabilities and return estimates, fuzzy though they were, we were comfortable with the

HOW SALOMON ROLLED THE DICE AND LOST

trade. But there were several points that bothered me in this analysis, one that Denham quickly brought into focus. Stavis's way of looking at things was right on point for an ongoing trading operation, but looking at the world from the chairman's seat, Bob did not see the trade as a purely expected value calculation. The losses had become so extreme that there were externalities to Salomon that needed to be considered. Risk arbitrage was a cash cow for an otherwise faltering equity division, and this was a year when the equity division was burning money at a higher than usual rate in its efforts to build an investment banking business in Europe and bidding for talent at exactly the time everybody else was. The division was under budget, a budget that projected negative returns even if everything went as planned.

Equity trading had traditionally been a weakness for Salomon. Parets, along with the derivatives trader Andy Constan, represented the beginning and the end of the real trading talent in the equity division. They were both protégés of Stan Shopkorn, a former Salomon vice chairman and pretty much the only legendary equity trader Salomon ever had. But success with equities involves more than trading ability and risk taking.

THE PROBLEM WITH STOCKS

Profitability in equity trading requires a more complex business structure than is required for fixed income. In the fixed income markets substantial profits can be made simply through the bid/offer spread. For the higher-risk and less liquid bonds such as junk bonds and emerging market bonds, the spread can be as wide as one or two points. Similarly, while the agency instruments in the mortgage market trade with eighth- and sixteenth-of-a-point spreads, the derivative instruments—collateralized mortgage obligations (CMOs), IOs, and POs—can have spreads that are multiples of those.

In contrast to the fixed income market, where a firm takes a principal position, transacts in large volume, and extracts a spread for the inventory and market making service it provides, equity trades generally move through the conduit of an exchange that takes over these functions. The firm receives only a modest commission. Also, as any surfer-turned-day-trader can attest, the cost of entry into the equity markets is low; equities are straightforward financial instruments. So to make any real money in the equity markets the broker-dealer must do more than simply answer the phone, take the order, and call it down to the exchange floor.

Profitability in equity business is tightly linked to investment banking, in particular to equity underwriting and initial public offerings (IPOs). This is the equity business in which some of the other top-tier firms such as Goldman Sachs and Morgan Stanley excel. The profitability is high because companies do not make equity offering decisions on the basis of price. The CEO and other decision makers have their own money and careers on the line, and any errors in equity offerings affect not only their wealth, but also that of their shareholders and board. So those firms that have strong reputations can extract a higher than market price. The cost of the underwriting is of about as much concern to the CEO as the cost of having the best heart surgeon is to the cardiac patient.

A profitable equity business requires a strong reputation in equity underwriting, and this in turn requires a strong investment banking team who can win the respect of senior management of the corporations they cover; a team of equity analysts who, because they have the respect of the institutional investor community, can be successful shills for the investment bankers' product; and a strong sales and trading infrastructure to run day-to-day transactions. Senior investment bankers earn seven figures, as do equity analysts with a strong enough reputation to really matter. And just as each fighter pilot has a ground crew of 20 people in support, each of these senior members requires broad support, covering everything from trading to technology to operations. There are dozens of industry sectors that must be covered, and the globalization of the investment markets has led to demand for coverage spanning not only the U.S. market but also European and Asian markets. Add it all up and it is easy to see that developing a strong equity franchise requires an investment of hundreds of millions of dollars a year. The payoff from this investment is uncertain. Profitability finally depends not just on talent and support, but also on the robustness of the market.

Firms were taking one of two avenues to develop the requisite franchise: recruit young talent and play to the market niches these stars develop, or buy top talent, either individuals or teams, the business world equivalent of bidding for free agents. The second approach met with mixed success. Deutsche Bank's investment banking effort foundered even after paying an exorbitant amount to recruit Frank Quattrone and his technology investment banking team from Morgan Stanley, who then fled en masse to Credit Suisse First Boston; and several years earlier Smith

HOW SALOMON ROLLED THE DICE AND LOST

Barney failed in its attempt to vault into the big leagues by hiring Bob Greenhill and a host of other investment bankers from Morgan Stanley.

As an organization run by traders, Salomon didn't have much patience for the slow and costly process of developing the investment bank and so got into the high-level recruiting game. In a world where astute trading can yield tens of millions of dollars in a matter of weeks, it is difficult to accept years of losses to build a franchise. Nor did Salomon have the free cash to compete with the major banks to buy the talent. Many of those with faith in the firm's trading capability thought we should follow a different approach: use our trading presence and our ability to make markets and take on risk as the core for moving into investment banking. But ultimately we got in line with the lumbering, corpulent likes of Deutsche Bank to bid for bodies.

Salomon had brought in Rod Berens from Morgan Stanley to head the effort to extend and build an equity franchise. This effort involved massive hiring, principally to build up the European equity business. And Berens was buying at the market top. The small and fixed supply of talent was in demand by many other firms—European banks, in particular—that had far deeper pockets and far more staying power than did Salomon. At each year-end partners' meeting Berens would show how the hundreds of millions of dollars of expense would shortly pass the break-even point. "Shortly" seemed like an eternity for the traders.

As Berens focused almost exclusively on building the investment banking component of equity business to create a traditional franchise through conventional channels, he took the high-risk step of relying on the proprietary side of equity trading—the risk arbitrage and equity options books—to provide an ongoing revenue cushion to moderate the losses. And now—as does happen from time to time with proprietary trading—these two businesses were in the red. One point that was not well appreciated about the arb units was that if the business did not make money, no one got paid and the firm had no expenses. Of course, the traders had the option that all traders have, in that if the group lost money, only the firm took the loss, not the traders. But in the case of Salomon's arb units, this was the firm's decision, not the traders'. On a number of occasions they had requested to put their own money into the pot with the firm's trading capital, but were repeatedly rebuffed once the request reached Deryck Maughan. The equities division did not have this sort of pay structure; there were sales-

people who had to get paid and other overhead expenses that had to be met. Parets's MCI/BT position was just too big given the overhead. So long as there was a reasonable chance of losing another \$100 million from the current position, Maughan could not stay with it, no matter what the upside potential might be.

The conclusion of the MCI/BT discussion led to this: Stavis felt that given the small odds of a break in the deal, the position should be even higher; I believed no one had enough experience with these situations to make any assessment regarding the risks and had to rely on those that Parets was providing; Denham's view was that we were facing ancillary risks because of business exposure from the position and the position should be dropped on that basis. Hall thought the position should be scrapped as ill conceived. It was probably bad before, and it was certainly bad now. Parets had already lightened the position over the course of the day. After two hours of deliberation spanning three continents, the risk management committee decided to maintain the position.

Stavis and I joked after the meeting how miraculous it was that with all the dislocations in the market, with all the surprise events and changes in the fortunes of the trade, it turned out that the position we had on at the time of our meeting still just happened to be exactly the right amount. Over the course of the summer everything about the trade had changed. The risks that were thought to be *de minimis* grew large. The strategy that was slated to make a tidy \$25 million had careened out of control, generating losses of over \$100 million. And beyond the trade itself, the equity division now was crippled with other trading losses that mounted to more than \$60 million. The blowup in Asian markets had pushed into the G-7 equity markets, raising volatility to unprecedented levels, and our equity derivatives business had a book that was short volatility and lost \$5 million with each percent rise in option volatility. Yet after all this, our careful deliberations led us to hold not one share less nor one share more than what we happened to be holding at the meeting's end.

It was a phenomenon that I found again and again and that seems to be an innate part of trader behavior: inertia against changing a losing position, and more specifically, inertia when faced with losses coming from unexpected corners. In experimental biology there is a term for this: experimental neurosis. An animal in the laboratory, beset by a strange environment and events that are outside of its past experience, will sometimes

HOW SALOMON ROLLED THE DICE AND LOST

simply curl up in a ball and ignore all of the stimuli. Its reaction to the alien environment is to freeze in its tracks. And, as I will discuss later, this is not limited to the behavior of animals in the lab; it is a phenomenon that arises from the core of how we approach the world.

INTO THE ARMS OF TRAVELERS

Triggered by Bankers Trust's announcement to purchase Alex Brown in February 1997, merger mania was spreading around Wall Street as institutionally oriented firms started to see the value of retail distribution. There was then the merger of Morgan Stanley and Dean Witter in the spring, and in the summer the acquisition of Furman Selz by ING and the purchase of Montgomery Securities by NationsBank Corporation. On Wednesday, September 24, Travelers joined in the fray, announcing its intention to acquire Salomon.

The purchase made a lot of sense from various vantage points. For Travelers' boss, Sandy Weill, it was a natural next step in his ladder of acquisitions. There were already rumors that he was looking to acquire a big bank, but at the time such a move would have met with legal and regulatory hurdles, because Washington often looked askance at commingling insurance and banking. He had already talked with Jon Corzine about acquiring Goldman Sachs, but without success. The acquisition of Salomon would not face the level of regulatory issues that would come with trying to acquire a commercial bank, and Salomon, because of its trading orientation where money talks, was far more likely to come to terms than Goldman.

The merger also looked good strategically. Salomon had Wall Street's elite bond business and had developed a global trading network, while Smith Barney, Travelers' brokerage arm, had retail business and a stronger connection to the smaller domestic businesses. Neither Salomon nor Smith Barney was particularly strong in investment banking, but in combination the two would be among the top five in both U.S. and global underwriting for both equities and debt, and would be No. 1 in municipal underwriting.

And the merger made sense for Salomon's largest shareholder, Warren Buffett, who had been spending years trying to figure out a profitable exit strategy for his Salomon investment. Salomon was being purchased

for more than \$8 billion, nearly two times book value; in reaction, its stock had climbed 80 percent from its recent trading level. Buffett would be able to exchange his control of almost 20 percent of Salomon for a stake in Travelers Group, and with the Weill-led team he would have a firm that was run by the high-quality management he valued.

But the impetus for the merger, what led to discussions in the late summer between Weill and Maughan, was the weakening of Salomon's trading position from the MCI/BT trade. Beneath Salomon's trading power was a perilous foundation. Ever since the trading problems in 1994, Salomon had been teetering on the edge of its investment grade bond rating. The principal focus for our CFO, Jerry Bailey, was to keep the rating agencies abreast of developments with regard to our business and to argue for not dipping below the BBB watermark. The losses related to MCI/BT further weakened this foundation, underscoring our trading vulnerability and making our rating all the more precarious. It also further weakened our equity business, which was already the weakest link in the firm. So having Travelers as a suitor, offering \$8+ billion for a Salomon that seemed all that much closer to the brink, was hard to resist.

Not only did the purchase make sense at the time, but as events unfolded over the next year, had Salomon remained on its own, without the financial support and the controls that Weill's operating genius Jamie Dimon put in place, it is likely the firm would have been blown apart by the fallout of LTCM.

CHAPTER 5

THEY BOUGHT SALOMON, THEN THEY KILLED IT

With the purchase of Salomon in 1997, U.S. fixed income arbitrage, the crown jewel of the Salomon trading machine, passed inauspiciously into Travelers' realm. Deryck Maughan, now co-CEO of Travelers with Jamie Dimon, called me up to his office to offer me a continuing role in risk management at the combined firm. Dimon sat with a single sheet of white paper in his hand, a to-do list with several columns of minuscule handwritten notes. As we discussed various issues, he vigorously crossed out one line here and another line there, then folded the paper when we concluded. I would learn more over time about his remarkable appetite for detail and his penchant for condensing things onto one sheet of paper, as we would construct a risk report with an almost unreadably tiny type size to compress all the relevant risks of the firm onto one 8½ by 11 inch page.

Despite his assurances to the contrary, Sandy Weill was not enamored with the high-stakes proprietary trading embodied by this group. That antipathy did not stop him from buying Salomon, the biggest trading house in the world—a firm that, even after many attempts to broaden its revenue base, still made all of its earnings through proprietary trading. To try to control the unit, he first pushed to have both the head of U.S. fixed income arb, Rob Stavis, and his European counterpart, Costas Kaplanis, report to his son, Mark Weill. Both men made it known to Sandy that they were not particularly impressed with Mark and offered to resign instead. So much for first impressions. For a group that was used to letting its revenues do the talking, it was just the start of the journey from the unfenced Western territories into the canyons of the corporate world.

Following this rebuff, Weill set up a series of meetings so that he and the new senior management team could try to become more comfortable with the arbitrage unit's business approach, trading strategies, and risk management controls. Over the course of the sessions, the three proprietary trading heads—Stavis for the U.S. unit, Sugar Myojin for Japan, and Costas for Europe—were summoned. Maughan and I were the other Salomon representatives. From Travelers came Weill, Dimon, and my new boss, Heidi Miller, who was Travelers' CFO. We gathered in the anteroom that lay between Weill's and Dimon's offices on the top floor of Travelers' building on New York's Greenwich Street. The building was a prize from Travelers' merger with Shearson American Express. It was about as out of place on the western edge of TriBeCa as a mountain in the middle of Kansas, and if there had been any "build it and they will come" theory behind its construction, it had yet to be realized. At that time there was basically nothing but low-cost housing and warehouses for blocks in either direction. (Today many of the warehouses are chichi loft residences.) At night the Town Cars lined up to carry the late-working investment bankers back to their homes or business dinners while an endless procession of panel trucks funneled into the side streets after completing their delivery runs for the day.

The anteroom was lined with bookshelves. They held a few books, most of which included some inscription or note of thanks for Weill's many charitable contributions, but more space was given to framed pictures of him with other notables, interspersed with awards and plaques. There were pictures with Arnold Palmer and Jack Nicklaus on the golf

course, with presidents from Carter through Clinton, and with a range of other celebrities and luminaries including Isaac Stern and Bill Cosby. He was clearly a man who had made peace with his ego.

Stavis came into the meeting with a number of counts against him. One was declining Weill's suggestion that he report to his son. The other was simply that he made so much money. And Stavis was something of a second-stringer; he got the top job after the free-agent stars had left to start Long-Term Capital Management—all true, but a bad rap nonetheless. He might have been young, but he had remarkable ability. He was just a few years out of college when he started work in the arb group, while Eric Rosenfeld, Larry Hilibrand, and the others had essentially already completed hypercharged careers either in teaching or in trading on the client side by the time they moved into the fledgling group.

Stavis was both incredibly bright and almost unbearably articulate. He would win any argument, most of the time because he was right but sometimes because he stated his views so clearly and persuasively that it seemed that he should be right. This could be mistaken for arrogance, a trait Stavis was often viewed to lay claim to, but in fact he considered any trading and risk management discussion as a search for the best answer. And in that regard his overpowering talents at debate, argument, and persuasion could lead to a skewed result.

Each of the presentations to Weill was straightforward enough; Myojin, Kaplanis, and Stavis in turn went through their approaches to trading. Stavis was last to present because his positions were the most complex. The others provided Weill et al. with an opportunity to ride the learning curve by first digesting some of the basics. Then it was Stavis's turn. The Travelers crowd listened quietly until he completed his remarks, and then they rose to leave.

It was all very orderly until Stavis made a huge gaffe. Nearly shouting above the after-meeting buzz, he asked everyone to please return the books he had prepared for reference during the presentation. He explained that they contained sensitive trading information, so he did not want them to leave the room. This request may have been instinctive for him; the proprietary unit had always jealously guarded its positions out of fear that they would be leaked to competitors, or even the client side of the firm, who could then trade against them (hence the name "proprietary"). When I received detailed position or risk

information from his group it often came on a special paper also used by the U.S. government for sensitive documents: dark red with a zigzag of fine black lines that made the text almost unreadable but prevented it from being photocopied.

But the group was no longer at Salomon and this was a distinctly non-Travelers moment. In most corporations, business units do not ask the chairman and the CEO to return information. Dimon simply tossed his book back onto the middle of the table. Weill tucked his under his arm and headed out of the room without so much as a glance in Stavis's direction. Heidi Miller, Travelers' CFO, asked politely if they would mind her pulling out a few of the pages where she had taken notes and then gave a purse-lipped glance in my direction. She had made it known to me more than once that she did not like this group, but now they had slighted her mentor and boss, and the time would come for payback.

THE ARBS LOSE FACE

Political miscues from the arb group would have been a lot more forgivable had they not been compounded by something that wasn't as easily dismissed: poor performance. Stavis's unit in particular had been struggling for most of 1998. Their bet that spreads would narrow was being undone by the continued uncertainty looming over the emerging markets—especially in Asia and Russia—and by the growing number of relative-value hedge funds that were making the same bet. In a textbook example of efficient markets in action, the profit differentials of the relative value trades, pursued aggressively by a number of well-capitalized and sophisticated firms—Salomon, LTCM, of course, but also Goldman and upstarts such as Convergence Capital, founded by Andy Fisher, formerly co-head of Salomon's U.S. arb unit with Stavis—were being bid away. And beyond the trades, the advent of euro-denominated debt was leading many of the markets themselves to disappear.

The Salomon risk arbitrage group's mortgage position was also facing losses because of an unexpected drop in interest rates and subsequent rise in mortgage prepayments. Even for the lower level of interest rates the amount of prepayment was much higher than expected, due to a number of factors: Electronic applications made it far easier for people to refinance, the cost of refinancing had been cut in half, and the increase in

perceived wealth due to the rise in the stock market increased prepayments. The result was a vast underperformance by the interest only (IO) securities in the portfolio. Others got nailed as well, from Goldman Sachs to LTCM, Michael Vranos's Ellington Fund, and Convergence Capital. At first the depressed prices in the IO market were viewed as overly cheap; for prices at that level to be justified, the rise in prepayment rates would have to be sustained over many months—something that none of the traders thought likely. The drop in the relative value of the IOs and the related mark-to-market loss on the positions were viewed as an opportunity, because if the IOs were attractive before, they were now even more so. No wonder groups like LTCM added to their positions in the face of the mortgage losses.

For the Salomon arb group these kinds of mortgage losses were bearable. And even in their worst-case assumptions—allowing prepayments to continue at their present rate for the next quarter—the IOs were at best fairly priced and still likely to be cheap. So the position was deemed to be recoverable, the losses stemming from market rather than model error.

Staring into the abyss was neither new nor scary to Salomon. The situation was similar to one that occurred in 1993, when Larry Hilibrand and most of the future LTCM team were still there. At that time the U.S. fixed income arb unit faced a loss of \$300 million in its mortgage position. Hilibrand made a presentation to the board and explained that the mispricing in the market was tied to the liquidation of Kidder's position. One of the most investment-savvy boards ever amassed, including Warren Buffett, Charlie Munger, and senior Salomon traders, agreed with Hilibrand. The board allowed him not only to maintain the position but also to increase it. Before the year was up Hilibrand and other senior members of the team joined "Salomon North" (LTCM) in Greenwich, Connecticut. That left the responsibility to two vice presidents, Rob Stavis and Andy Fisher. Nervous senior managers called Dennis Keegan back from his post as co-head of London trading to babysit the new U.S. fixed income arb team. It was probably the best move of his life. The mortgage position Hilibrand had so forcefully argued to maintain and even increase snapped back from its extreme mispricing, and by the end of 1994 it had amassed a stunning \$1 billion in trading revenue, which led to bonuses of \$30 million each for Stavis, Fisher, and Keegan.

A DEMON OF OUR OWN DESIGN

While this was business as usual for Salomon, in the merged world of Salomon Smith Barney losses upwards of \$100 million were not commonplace, and the style of convergence trading that required models to track positions too complex for the human mind and that saw losses as an opportunity for further gain was totally foreign.

But the most perplexing problem for the U.S. fixed income arb group—and the one that finally did them in—was not political arrogance or lack of profits, but an inexplicable tracking error in their yield curve trades.

GOOD-BYE TO QUANT HEAVEN

Since its inception in the mid-1980s, the fixed income arbitrage unit sat in the center of Salomon's cavernous fixed income trading floor at Seven World Trade Center, not far from the client-based Treasury and derivatives desks. The proximity helped perpetuate the fiction that the group was actually linked to the rest of the firm, a fiction that became increasingly hard to maintain as the arb group's compensation levels rose first by multiples, and then by a factor of 10 beyond the compensation of those who sat across the aisle doing the tedious, day-in, day-out customer flow business, buying and selling orders for bonds that came in from institutional clients. In 1996, the fixed income arb unit moved from the fixed income trading floor on 42 to its own private domain on the 35th floor. This separation allowed the arb unit to get more service from other dealers, who understandably were shy about dispatching information when the Salomon market makers were just a few feet away. It also gave them a chance to freely express their personalities in the work environment.

Silicon Valley's roll-your-own office had nothing on these guys. Certainly, they had the standard trading machinery: Two rows of trading desks were laid out back-to-back by the windows facing the World Trade Center towers, with conference rooms beyond them at one end and the research offices at the other. Behind the trading desks, though, stood an adult playpen, a large open space filled with various toys, including a Nintendo, a chess set with clock, a cappuccino machine, and a putting green with a computer-controlled mechanism that could alter the slope and pitch of the surface. There were also scattered electronic gadgets. "We all take

THEY BOUGHT SALOMON, THEN THEY KILLED IT

turns seeing who can bring in the most obnoxious toy,” observed Jeff Rosenbluth, the unit’s co-director.

The desks were covered with catalogs of conspicuous consumption: auction books on wines, brochures for custom-made silver belt buckles, and announcements about Ferrari shows. Every Thursday afternoon part of the conference room was cleared for a masseuse, who brought in her table and provided massages for the group. Inlaid in the center of the aisle, between the rows of desks, was a worn relic that had also adorned their area when they were on the fixed income floor: a maroon two-by-three-foot prayer carpet that John Meriwether had bought years earlier.

These silly outward trappings belied the real spirit of the group: The arb unit was the most intellectually intense place I have ever run into. The lights burned a little bit brighter there. It was a concentrated remnant of the Salomon of old.

Rosenbluth, a roommate of Stavis’s at the University of Pennsylvania, was by 1997 running the unit jointly with him. Jeff was one of Amazon.com’s first and best customers, his desk buried under a mass of math books. He took master’s-level classes in mathematics at the Courant Institute of Mathematical Sciences, and after retiring in 1998 studied there for a PhD. His previous avocation was chess. The chess set was his; he had trained with a Russian grandmaster and helped finance a chess Internet site that was also sponsored by Anatoly Karpov, who stopped by to meet with us once in the partners’ dining room.

There was plenty of time for games because most of the spread trades had long holding periods, sometimes a year or more. They required only occasional adjustments to hedge out unwanted risk factors. If these guys had to spend their time simply monitoring the trades, it would make watching a freezer defrost seem like a movie on fast-forward. Instead, the trading day was spent positing different scenarios and testing their impact on the positions, devising reports to better elucidate risks, and refining the trading models for any anomalies that were observed.

The models became ever-refined repositories for the aggregate experience of the traders and the dozen or so PhDs who did the analytical and programming work. They also engendered a close working relationship among the traders, who more than anything were risk managers. In most Wall Street research areas, the programmers and researchers sit in cubicles off the trading floor and do the traders’ bidding within a clear

caste structure. As one client-side trader put it, “I stop down there every now and then and throw some raw meat over the side of their cubicles.” By contrast, in the U.S. fixed income arb unit everyone was a researcher. The day was spent challenging assumptions and proposing new tests. The models were not black boxes that spit out trades. They were the summation of market analysis put in a form where they could be refined and tested and, had the world been different, could have even survived the group.

This was the intellectual Mecca at the foundation of Salomon’s reputation, a world that was replicated at LTCM, where any academically trained finance professional would have loved to work. Over the spring and early summer of 1998, I spent increasingly more time in this world, working with the arb group to try to explain their strategies to a hostile and less intellectually inclined Smith Barney management—principally my boss, Heidi Miller—and to try to ferret out the origin of the new losses.

CIRCLING VULTURES

Stavis was facing an inexplicable tracking error in the U.S. yield curve trades. The trades are designed to be market neutral; that is, they should not be affected by changes in interest rates. But for most of the winter and spring, the trades of the arb unit could not seem to immunize their interest rate exposure. With interest rates dropping, this unintended interest rate bet posted more than \$100 million in losses. Taking outright interest rate exposure was outside the mandate of the arb group, and however unintended the exposure was, it provided ammunition for Smith Barney honchos, who had never understood, and therefore had never been comfortable with, these sophisticated strategies.

The problems with interest rate exposure were a major concern for the U.S. arb group throughout the spring. As rates dropped, the trades moved against them. Though the mounting losses were a clear cause of concern, the group persisted in holding the positions, because their initial working hypothesis was that they were hedged. In the long run, the various instruments on either side of their trade had to converge. For the moment, they were experiencing a mark-to-market loss, mere scorekeeping, rather than an imperfection in their model. They regarded the correlation of their losses with the changes in interest rates as some type of mar-

ket inefficiency. The market *had* to move back to match the mathematical equivalence—except that it didn't.

To make matters worse, the supposedly smaller minds were minting money. The Treasury desk on the client side shared neither the long-term trading horizon nor the theoretical tools of the arb group. It held positions based on client demand, and could shift from being net long to short in days. Also, it was not given the leeway by management to hold large inventories and stand back as the day-to-day price movements led to P&L swings of hundreds of millions of dollars. Charlie Parkhurst didn't have time for the market to come to him. Parkhurst, the head of the client Treasury desk, relied on a historical rather than a purely theoretical model for determining the hedge between the futures and cash Treasuries. That is, he ran a regression over the recent past to see how the two classes of instruments tracked, and then hedged on that basis. The recent history included this mysterious correlation between interest rates and spreads, so his hedge took it into account.

Sometimes it is better to be lucky than smart, and with this simple-minded hedge Parkhurst was making money holding the same strategy as the arb desk while the arb group was seeing its P&L sink by tens of millions each month. The fact that the client desk was earning profits with rudimentary tools while the supposedly elite arb desk could not extricate itself from the interest rate bind despite all of its PhD firepower brought the group down a notch or two from its pedestal.

Stavis, Rosenbluth, Parkhurst, and I had had discussions about the possible reasons for the mistracking and the residual losses. One was that it was simply a spurious correlation. The trade might have just happened to be getting cheaper at the same time interest rates were dropping. Markets either go up, go down, or stay the same, so if you are losing money in one there are bound to be others that will be losing at the same time. That does not mean the two are functionally linked. Another possibility was that the arb model did not pick up all of the factors affecting interest rates. The model was a proprietary yield curve model dubbed the "two plus" because it looked at the yield curve as two factors, plus a parameter to signal the effects of Federal Reserve policy shifts.

The two-plus model was the citadel of intellectual capital for the group. It was a closely guarded secret, although, despite the group's best efforts, it found its way to a number of other firms as talent was periodically bid away.

The model was developed by Bill Krasker in the mid-1980s shortly after he came to Salomon from a brief stint teaching at Harvard. An exceptionally well-trained mathematics major from Princeton, he was dubbed by Maughan “the world’s highest-paid computer programmer.” (This was before the Internet bubble.) Like Eric Rosenfeld, who had also flown the coop to LTCM, Bill had been a classmate of mine at MIT, and in the second year of our program he gave a set of lectures during the winter break entitled “Functionals in Banach Space” where he spun out page after page of definitions, theorems, and proofs without the benefit of any notes.

All the risk arb group’s analysis indicated that adherence to their well-tested two-plus model would keep them hedged to interest rate levels, shifts in the slope of the yield curve, and even changes in yield curve convexity, or twist. They hypothesized that there must be some new residual factor—which they dubbed the behavioral factor—that still kept them exposed to interest rate changes. There may have been some solace in giving it a name, but it remained an enigma.

In the eight months since the merger, there was already an uneasy relationship that magnified the need to isolate and eliminate the source of this tracking error: Smith Barney bosses considered the arb group to be arrogant and impertinent. And in an attempt to pull out profits in the face of uncooperative markets and skeptical management, they had already been caught straying from their relative value mandate. The most visible example was an outright yen trade that had grown to better than a billion-dollar position and had been obscured in reports. When it finally came out in a risk management committee meeting, Dimon was furious. Jeff wrote a memo trying to explain why their position made trading sense, but that missed the point. The issue for Smith Barney was control. Most of those in Travelers’ management felt they had purchased a bunch of immature cowboys, and while they didn’t quite know what to do with them and couldn’t understand them, they did know that the way Travelers made money—besides acquiring assets in a rising market—was tight control. In fact, after the merger the only control function not taken over by Travelers was my risk management role.

THE JULY FOURTH MASSACRE

Whatever the root of the losses, they had become significant and persistent. In the end-of-quarter analysis of the arb unit’s performance I pre-

pared for Dimon after June month-end, the tracking error and related losses were painfully obvious. When Stavis and Rosenbluth got a copy, they hit the roof. They immediately called me to argue the analysis: Should it be done on first differences versus levels? Was the sample period large enough? Did it include spurious correlations? Were all the positions being included? But the computations were simple and incontrovertible: In the four months since March, they had an interest rate exposure that would lose \$3 million for each basis point drop in rates. Over that time, the 10-year Treasury bond had dropped by 35 basis points, so rate exposure could explain about \$100 million of their losses, more than half of their total of \$170 million. The rate exposure had increased over time, but had been significant for nearly a year; since late in 1997 it had averaged \$1.5 million for each basis point move.

I was upset by the strong reaction my memo had engendered, but that was because I had remained naive to the troubles that were mounting for the arb unit. I started to get more than a hint of what was coming, though, when Heidi Miller called me shortly after receiving her copy of the report. Her questions made it apparent that she was gathering ammunition in her hunt to bag proprietary trading. I recounted for her the tracking error problems that had been the target of what was thus far an unsuccessful analysis, to which she simply replied, "Well, that's your view."

While there had been ongoing discussion of the tracking problems, it was easy to present an alternative spin that rather than a major miscalculation in their hedge, Stavis and Rosenbluth had deliberately made an interest rate bet to try to salvage their business. Such a conclusion would mean that they had disobeyed the direct injunction of Dimon not to take outright directional risk. Even the most benign interpretation was mildly damning; for all the apparent sophistication of the arb group, they still seemed to need remedial help with basic bond hedging. In the face of the earlier losses and the seemingly arrogant, aloof posture of the group, the implications of this memo were enough to push things over the edge. In a meeting over the July Fourth weekend the decision was made to close down the unit.

This was not the first time the sights of the firm had been drawn to the proprietary trading unit. Warren Buffett had occasional thoughts of getting rid of the arb unit, depicting the proprietary/client setup at Salomon as "a casino with a restaurant out front." But he could never seriously

consider doing so because the arb unit made half the firm's revenue and pretty much all of its earnings. Travelers, however, could take such thoughts more seriously, because proprietary trading represented only about 10 percent of its revenues. By July, with the Citibank merger looming on the horizon, it would not be long before the arb unit would be little more than a footnote on the income statement, irrelevant in terms of earnings or contribution to the retail franchise.

The next question was how to engineer the closing. The key factor in this decision was—or at least should have been—to minimize transparency. The U.S. fixed income arbitrage unit was sitting on billions of dollars' worth of positions that spanned from the mortgage market to U.S./Canadian bond spreads. These positions took months to build and, even in a cooperative market, would take just as long to sell off. Some of the positions were huge; others were both huge and in specialized markets; still others were in complex, interleaved strategies where executing a smooth unwinding was like planning the uneventful demolition of the Eiffel Tower. The U.S. fixed income arbitrage unit's stockpiles were matched by closely linked books in Salomon's European and Japanese arbitrage units. And other firms held similar positions. For all the attempts at secrecy, it was inevitable after going into the market year after year and having members of the group move on to other firms that the major strategies not only would be known but, given the spectacular success and revered status of the group, would also be emulated. Goldman Sachs's proprietary operations mirrored many of these positions, as did a host of hedge funds. And at the head of the pack was LTCM.

Most investors do not appreciate the issue of transparency because it is only when positions are so large that their liquidation can affect the market. But when you have to put billions of dollars to work and there are only so many good ideas to execute, the individual positions can take days or even weeks to work into the market and later liquidate. If other funds know that you have a large position and are in the process of closing it out—and especially if the reason for closing it out is that the price has been going against you—they will start to sit on the sidelines or even trade the other way.

Of course, this can end up being a game that can be played various ways. In the early nineteenth century, Nathan Rothschild, who ran the London office of the family banking empire, executed a famous case of gaming transparency. The story is often told that he made a killing in the sovereign debt market based on his early knowledge of the outcome of the

THEY BOUGHT SALOMON, THEN THEY KILLED IT

battle of Waterloo, where in 1815 an allied European army ended the global ambitions of France and Napoleon Bonaparte. Like most such stories, it contains a kernel of truth—Rothschild indeed had the scoop—but it is what he did with it that is the real history lesson.

ROTHSCHILD AND WATERLOO

By the end of the final hour of battle of Waterloo on June 17, 1815, just 10 hours after first contact, a quarter of the Duke of Wellington's troops lay dead. The French losses numbered nearly 30,000. Within the space of half a day, Waterloo claimed more casualties than any other battle in history. Given the communications limitations of the period, Great Britain could not immediately know about the carnage of Waterloo or the swift victory. The last word that had come across the Channel had been of a Prussian defeat at Ligny, British reverses at La Haye Sainte, and then a massing of French troops that would have put Wellington at a numerical disadvantage. Wellington's envoy, Major Henry Percy, was dispatched to send the news of the victory to the War Office in London, but he and his horse were affected by the physical toll of the battle. Even with his best efforts, he did not arrive until late on the night of June 21. Until that time, all of Britain waited in suspense—all of Britain but one man, Nathan Rothschild.

There are many legends about how Rothschild received the early news of the victory at Waterloo. Some say he used carrier pigeons; others even have Rothschild, a rotund though expert horseman, riding about the battlefield and then racing at full gallop to a waiting boat and crossing the Channel in a raging sea at the peril of his life. The actual path of the information was probably much less dramatic: As Europe's bankers, the Rothschilds depended on speedy information and reliable couriers for their business as a matter of course. They had an agent and messenger in Amsterdam who could easily have been kept stationed near the battlefield. With the news in hand and fresh horses at the ready, the messenger could have been in Dunkirk after a day's ride, crossed the Channel in a waiting boat, and dashed to London. In any case, Rothschild knew of the results by the early hours of June 20.

Rothschild's first stop with the news was a visit to the prime minister, Robert Banks Jenkinson, second Earl of Liverpool. He arrived after the PM had gone to bed, and so returned in the morning. Whether this call

was out of jubilation, arose from a sense of duty, or served another strategic purpose cannot be known, but surely by the time he had returned to Downing Street word must have leaked out that the Rothschild network had information about the battle.

Rothschild's next stop was the London Exchange. Dressed in a coat and top hat, he took his usual position in front of one of the pillars of the Exchange, a spot so familiar that it was known as Rothschild's Pillar. It was there that he initiated a strategy that he would repeat, in less spectacular fashion, many times in the future. The primary trading instrument on the exchange was British consols, bonds that provide coupon payments in perpetuity. Once the victory was known, the consols would rise in price. But rather than buying, Nathan started to sell. As the price dropped, he sold even more. With the Battle of Waterloo on everyone's mind, his aggressive selling made it clear that the firm did indeed know the outcome. And they were selling; the Battle of Waterloo was lost.

Rothschild's very public position-taking made the information behind his actions transparent. And the market then provided him no quarter. His selling was met by others who piggybacked on his information and sold on their own. Prices fell in an avalanche without him doing any more. Of course, the information that was conveyed was the wrong information, and once the market had reached bottom, he started to buy. The buying, done quietly through a number of agents while he stood expressionless and dispassionate, was far less transparent to the market than was the initial selling. The result was a fortune for Rothschild and ruin for those who sought to ride on the back of his information. Although this was the most dramatic and quick strike by the Rothschilds, it was just an asterisk on the enormous fortunes they amassed over the course of the Napoleonic Wars. The wealth and station that would hold them in stead for generations came over the course of the war through various other channels too: by acting as agent to the German Prince William in securing and then speculating on his wealth, by gathering William's tax receipts ahead of the French army, and by smuggling gold through France to finance the British effort.

FOOLING JUST ABOUT NOBODY

If Rothschild's strategy sounds like the inspiration for the closing scene in *Trading Places*, where the villains get wiped out on the commodities floor, it

probably is. But such conclusions work better in movies than in real life. More to the point, with the positions of Salomon and LTCM, there was little doubt about the direction the holdings were heading.

Because the positions were no longer opaque, it was imperative that the winding down be kept as quiet as possible. If the market sniffed Salomon's selling, especially if it knew the selling was based on marching orders from senior management, then the liquidation would turn into a fire sale. The best course of action was to gradually and silently take the positions down. This was already happening in any case, both at the behest of management and because of a paucity of good trades.

Dimon then gave me a devastating countervailing view: The death of the arb unit was a material event and could not be kept from the public. He told me that you can't keep quiet about closing the unit with the largest concentration of trading talent and experience in the firm, one that had contributed the majority of Salomon's earnings over that past five years. From CFO Miller's perspective, a public announcement might serve another purpose: It would foreclose backpedaling; it would be the stake through the heart that would keep the U.S. fixed income arbitrage unit from rising again.

Dimon called a meeting of the arb group for Monday afternoon at 5 P.M. of the long July Fourth weekend. He walked into the room punctually and announced: "The following will be released to the press in five minutes." Then he read a tersely worded memo that began, "We have decided to restructure and significantly decrease the risk profile of the U.S. fixed income arbitrage group and integrate its people into other units of Salomon Smith Barney." Game over. The members of the group were told they had a month to decide if they wanted to stay at Salomon Smith Barney in another capacity or take a severance package. With that Dimon departed. One of the managing directors of the group rushed out of the building to call his wife, only to discover that his company calling card had already been canceled.

On Tuesday morning, as I got ready to return to work, I pulled in the *New York Times* from outside my door and sat down in the living room to glance quickly through it. A succinct article inside the business section stated that Salomon had announced that it was shutting down its U.S. fixed income arb business unit. The stake had been driven. The headline of a later editorial by Floyd Norris of the *Times* said it all: "They

Bought Salomon, Then They Killed It.” The spirit of proprietary trading had been exorcised from the body of Salomon Smith Barney, but ultimately at a huge cost.

Back at work, dumbfounded and guilt-ridden, I did not feel I would be welcomed in the arbitrage domain to commiserate. By not doing so, though, I ended up alienated from some in the group, who, as they would later remind me, had just a few months before staunchly defended me in the face of attempts by Miller to dislodge me from my role. Within the firm, the announcement caused a 180-degree shift in the view of the arb unit. The group that was touted by Salomon as being unsurpassed was brought to earth: “Anyone could do their strategies.” “The client business had been doing the same thing all along with better results.” “You don’t have to be too smart to lose \$1 billion.” The envy the client-driven traders had for the freedom and huge payouts of the arb unit had been unleashed.

I could not even persuade the new Salomon Smith Barney to spend the time to debrief the group and secure the use of their models. Despite my urgings, an unmatched repository containing the codification of a decade of trading experience and insight that even with the recent setbacks had been the main engine of profits for the firm was tossed aside. I managed to play both sides of the battle badly. The firm was suspicious of me for thinking too highly of the arb group while the arb group thought I had become the quintessential organizational man.

Outside of the firm, the announcement changed perceptions of the markets. Every trader around the globe could expect Salomon to unload its huge positions. No one could say when, but no one wanted to be on the same side of the market when the floodgates opened. For the rest of the summer, the U.S. swap market was ominously quiet; volatility dropped as traders stood back. It was like a bunch of teenage boys milling around the edges of the dance floor to see who would be the first one to ask a girl for a dance, or like an old war movie cliché: “The market is quiet—too quiet.”

It was in the quiet months after the arb group’s demise that the genesis of the next crisis, LTCM, would emerge. Certainly, the LTCM debacle involved the interplay of many factors: It employed absurd leverage, hundreds of billions of dollars of position with borrowed money, and a mere \$3 billion of its own capital. And LTCM contributed to its own problems by straying far afield, moving into risk arbitrage and equity volatility, strate-

THEY BOUGHT SALOMON, THEN THEY KILLED IT

gies where the principals had no expertise. Its prospects were also becoming dulled by the entrance of others into its relative value trading reserve.

But ultimately, LTCM's collapse stemmed from a crisis of liquidity. Following the Russian default in August, LTCM needed to sell assets, but the market simply was not there. And the reason the liquidity disappeared can be traced to a single event: the public announcement on July 6 that Salomon Smith Barney was closing its U.S. fixed income proprietary trading unit. If the Russian default was the match that set off the conflagration that consumed LTCM, the tinder was dried and stacked over the July Fourth weekend by Travelers' chairman, Sandy Weill, and his co-CEOs, Jamie Dimon and Deryck Maughan.

With U.S. fixed income arb closed, everyone knew it was only a matter of time before the next shoe would drop. The thinking in the market was that if Salomon was closing its U.S. proprietary group, it probably would do likewise with the London and Japan offices as well. The assumption was that Salomon's London inventory would be unloaded first. The effect on liquidity suppliers was just like that caused by the downturn of the equity markets on October 16, 1987: Everyone knew there would be an overhang, and no one wanted to be on the other side of the first trade. Who wants to buy the first \$100 million of \$10 billion of inventory knowing another \$9.9 billion will follow? Japanese proprietary trading held huge convertible bond and warrant positions and, at times, had \$20+ billion in swap spreads. European fixed income arb held positions that closely mirrored LTCM, with swap spreads, mortgage bond spreads, and volatility positions scattered across Europe. Without a trader making a move, without another word being said, the announcement of the closing of the U.S. fixed income arb unit led all these markets to freeze in their tracks. Clueless to the standstill, LTCM chugged along full steam ahead.

LEVERAGE AND THE ROOTS OF CRISES

The very ability to liquidate—clearly a desirable attribute of an investment portfolio—is, ironically, at the root of the liquidity crisis cycle. To see this, consider why the spiral of a crisis is common for financial institutions but not for corporations facing bankruptcy. The reason is that financial institutions are liquid and corporations are not. When the run-of-the-mill corporation cannot meet the terms of its creditors, no attempt is made to

throw assets into the marketplace, because there is no ready market for the assets. The trigger of the liquidity crisis is unique to financial institutions because their assets are liquid and can be marked to market. The steel mill that runs into trouble will not face this sort of spiral into oblivion even if its assets are heavily levered because they do not have liquidity. No one can pull an up-to-the-minute quote for the value of the steel mill; it cannot be monitored for default in real time. No one can think about liquidating it at a moment's notice even if it is in default.

Of course, because a corporation is not readily liquid, no one will lend it money on the same favorable terms that they will lend to those who borrow to buy financial assets. Liquidity means you can determine prices at every instant and you can demand cash on a moment's notice, which means you are more willing to lend money against collateral. Liquidity allows for ready leverage, but it also creates the means for crises.

During a market shock, the liquidity suppliers—the market makers, broker/dealers, and bargain hunters in the trading community—will recognize the risk of a cascade of liquidity demand and be reticent to be the first ones to take on supply. Indeed, some of the liquidity suppliers may actually find themselves in the position of demanding liquidity.

HOW TO PREVENT A LIQUIDITY CRISIS

If we could estimate the price impact from the flow of the margined fund's liquidations, we could chart the full course of the liquidity cycle for the fund. For example, suppose that when the fund sells 10 percent of its initial assets in response to the price shock, that sale lowers the market by 1 percent. And suppose further that this relationship is linear, so, for example, a sale equal to 5 percent of its initial assets will drop the market price by 0.5 percent. The cycle will continue to work itself out by the subsequent 1 percent drop leading to a sale of 10 percent of the fund's remaining assets. The price will then drop by only 0.9 percent because now the fund's remaining assets are only 90 percent of its initial portfolio. This price drop will lead to a margin-induced sale by the fund of 9 percent of its still remaining assets, which in turn will trigger further price drops by the market, inducing more margin sales, and so on.

Depending on the price elasticity of liquidity supply, we could end up without convergence because the price would have to move so far to elicit

sufficient supply that the fund would have to sell an ever-increasing proportion of its remaining assets. The cycle would become a market crisis as the drop in the market price led to similar needs to liquidate at other firms. The flood of liquidation would accentuate the price drop at every turn. The point of no return would come when the effect of liquidation elicited greater demand for margin for the remaining fund position than the amount of cash it would raise from the liquidation. That is, suppose that to raise \$10 million to meet a 10 percent margin requirement, the fund must offer down the prices to an amount that causes the prices of its remaining positions to decline by more than \$10 million. The fund then faces yet another call to meet the resulting mark-to-market loss. It is caught in an ever-widening downward spiral and cannot satisfy its creditors' or investors' demands no matter how aggressively it sells. Indeed, the very need for aggressiveness in liquidation becomes the root of the problem.

In this situation, although poor investment decisions or adverse market events may act as a catalyst, once the losses move beyond a critical point, the crisis becomes self-sustaining; it feeds off the need for liquidity, and that need for liquidity does not come from the market but from the demands imposed by the fund's creditors and investors. To understand the risk of crisis, we need to understand each stage of this cycle—the risk of a large loss, the risk that the loss will force liquidations, and the risk that a forced liquidation will adversely affect market prices.

As a case study for this, it is hard to beat the death spiral of LTCM.

CHAPTER 6

LONG-TERM CAPITAL MANAGEMENT RIDES THE LEVERAGE CYCLE TO HELL

Jamie Dimon wasn't buying any more of the Russia story. The Salomon Smith Barney co-CEO held his hand up in the air, his thumb and forefinger nearly touching. "By our next meeting," he ordered, leaning toward Mark Franklin, "I want our Russia exposure down to this."

Franklin, relatively new to the arb group, had a portfolio loaded with hundreds of millions of dollars of Russian bond exposure. And Russia was heading toward default. The rule of law was absent in Russia; conducting ordinary business could be life-threatening. Even the most rudimentary analysis showed that the Russians were spending money and not taking any in. The country had become politically unstable, a kleptocracy with rampant corruption. While many saw the writing on the wall, the arbitrage group was still holding on. They and a number of other trading desks remained enamored with the huge yields on Russian bonds and kept their

positions with the hope the country was “too big to fail.” And not just “too big”; the Russians had nuclear weapons. So, some argued, the West would have to bail them out. Dimon was not going to take that bet; he had pushed for months to get positions down, but got little cooperation. Finally, toward the end of June 1998, after having listened to Franklin’s optimistic assessment yet one more time, Dimon simply presented the decree. I was the one who ultimately had to track the arb group’s positions, so to make the message clear he swiveled toward me, held his measured hand in the line of sight between his eyes and mine, and repeated: “Down to this.”

Everyone at Salomon wanted to hate Dimon for interjecting himself into our culture, but over time he won widespread respect, even admiration. The fact is that he could have done almost everyone’s job better than they could. His dogged pressure to get out of the Russia position would have earned him a Salomon Trader of the Year award in an earlier era. On August 17 Russia simultaneously devalued its currency, defaulted on most of its domestic government debt, and declared a moratorium on the payment of principal to foreigners. There was a run on its banks, a precipitous fall in the exchange rate, and an acceleration of inflation that boded ill for future Russian economic reform. From that point in early June when Dimon issued his decree until the Russia crisis hit in late August, our gross trading inventory of Russian debt dropped from nearly \$3 billion to a tenth of that amount and our Russian ruble exposure was cut by 80 percent.

We had dodged a bullet by getting out of Russia. Long-Term Capital Management had not. The Russian default acted as a trigger for a global crisis that would send LTCM into a fatal liquidation spiral, and come close to taking Lehman Brothers and Merrill Lynch with it. Although Salomon wasn’t as damaged, it still suffered more than \$1 billion in trading losses, which ultimately led to Dimon’s departure and Maughan’s demotion. In fact, of all the members on Salomon’s risk management committee, representing some of the top trading experience and talent on Wall Street, only one, Tom Maheras, the head of the fixed income division, would remain at Salomon Smith Barney three months later.

RISKY BUSINESS

Salomon’s weekly risk management committee meetings were the highest-level gathering in the firm, attended by the heads of the major trading

LTCM RIDES THE LEVERAGE CYCLE TO HELL

units as well as the CEO and the chairman. After Salomon's purchase in 1997 by Travelers, the trading heads were joined by the new co-CEOs of the merged Salomon Smith Barney, Dimon from Travelers and Maughan from Salomon. These meetings used to take place in a nondescript room off the fixed income trading floor, but after the merger with Smith Barney they were held in a grab bag of venues. Salomon's offices were at Seven World Trade Center, and Smith Barney's were in the Travelers building, about a 15-minute walk to the north. Both Maughan and I were housed at Salomon, while Dimon retained his office in the Travelers building. Fixed income and U.S. proprietary trading were at Salomon's building, while equity trading was at Smith Barney's. As part of the bridge-building process, the meeting alternated between the buildings, and we seemed to pick a different room each time.

For the first meeting of the third quarter of 1998, in the first week of September, we were in Salomon's soon-to-be-deserted boardroom on the 43rd floor. The room was paneled in mahogany. A huge oval table stood in the center, surrounded by two semicircular tables for the support staff. A brass stand next to Chairman Bob Denham's chair held a red phone. Double doors led out to the polished black-and-white Italian marble floor of the foyer, and another door led to the partners' kitchen. The windows were covered with heavy curtains treated to shield the room from electronic eavesdroppers. The boardroom seemed to be a throwback to the long-gone, hushed, clubby days of investment banking; but it was both an anachronism and a façade. The room was seven years old, dating to when Salomon moved to its new building, which was originally built for Drexel Burnham Lambert and then purchased by Salomon after Drexel's bankruptcy. And Salomon never was part of that investment banking club in the first place.

Even with everyone in attendance for the risk management committee meeting, the room felt eerily empty. It echoed the passing of a great risk-taking trading house into a lukewarm purgatory of retail client coverage. There would be no more Salomon board meetings. The weightiness of Warren Buffett and the acerbic critiques of Charlie Munger had passed into history. Indeed, this was one of the last meetings that would be held in this lofty space before it would be dismantled, along with the trading floors below it. Television monitors dotting the periphery of the room tied in Costas Kaplanis, the head of European proprietary trading, from London.

But for all the monitors, there was only one camera, and it made the meeting a bit awkward. At first, half of those in attendance were seated outside its range. Even once they were reseated in range, the camera was set oddly so that if you tried to address Kaplanis by looking at the monitor the camera caught you in profile. From his view, you were not just avoiding eye contact; you could have been addressing someone else altogether.

We were reeling from losses related to the LTCM debacle, and this meeting was focused on each area of loss. Losses of nearly \$900 million in our European fixed income arbitrage unit were spread more or less evenly across swap spreads, yield curve, and a European-U.S. yield volatility trade. Added to that was another \$100 million loss from an obscure relative value trade that the European arb group had put on between Scandinavian inflation-linked bonds and ordinary Scandinavian bonds. (The trade was not that obscure, it turned out; LTCM and a host of others had managed to navigate their way into the same trade). In U.S. arbitrage, our bellwether unit and usually the area of greatest risk taking, we were about \$200 million in the red. Travelers had closed it in early July, and the bulk of the positions had been liquidated. This action saved the firm a loss that would have certainly exceeded \$1 billion, yet perversely it was what set the whole crisis in motion in the first place. The swap spread trade in Japan that had been the mainstay of the Japanese arbitrage unit's profit chipped in another \$100 million in losses.

Given the magnitude of losses that were occurring elsewhere, and given that the Japanese swap spread strategy had added a couple of billion dollars of profit over the years, this loss—similar in size to the one I was brought in to troubleshoot a few years earlier—was little more than a hiccup by comparison.

On the fixed income side, our losses mirrored those hitting LTCM, which was not surprising. Both outfits were famous for strategies that would bring them both to grief: relative value trades. The principals of LTCM had pulled up stakes at Salomon in 1993 and 1994 and set up shop in what became known as “Salomon North” in Greenwich, Connecticut. The original fixed income arbitrage group had been founded in the mid-1980s by John Meriwether, who was forced out of the firm to make peace with the Federal Reserve after the 1991 Treasuries scandal. Meriwether & Co. landed at LTCM largely intact, a cadre of MIT-trained professionals, many with PhDs and former careers as university professors. Meriwether

LTCM RIDES THE LEVERAGE CYCLE TO HELL

has been celebrated for his coolness under pressure, but that is not an uncommon trait among top traders. What was bold and innovative was that he recruited a team with analytical talent to run a trading operation rather than just support it, and then set them loose with substantial capital and a set of new trading approaches that were ideally suited to their skills.

THE DECEPTIVE CHARMS OF RELATIVE VALUE

Relative value trades have inherent appeal. They do not depend on market direction, they seem logical, and they appear to keep risk under tight control. The problem with these trades is that while the profit opportunity might be easily captured, the resulting profit is very small. A good opportunity for a convergence or spread trade is a spread differential of 20 or 30 basis points. If you have a billion-dollar hedge fund and put all of the capital in such a trade, if all goes well it will return only \$2 million to \$3 million. To get real returns out of these strategies—returns of 20 or 30 percent—the trades need to be leveraged 20, 50, or even 100 to 1. But if you are leveraged 50 to 1 and things don't go well, and the trade moves 30 basis points the wrong way, you will be out \$150 million—a loss that on a \$1 billion fund may put you out of business. The very act of using that much leverage increases the likelihood that things will go against you, and the larger the hedge fund the worse it will be.

And spread trading has its problems. As Charlie Munger once put it in a Salomon board meeting, "A spread trade is like a stick with s...t on both ends." To listen to the arb group, a spread or relative value trade is far more appealing than an outright position because when you are holding a similar instrument long and short there are many risks that will affect both sides of the trade similarly and will therefore be hedged away. A less positive view is that a spread trade matches one type of instrument against another, so the chance for something to go wrong doubles. If either market faces a liquidity problem or a falling-out with investors, the spread will be in trouble. Because the spread demands the proper hedge ratio between both sides of the trade, a shift in volatility in either market relative to the other will affect the hedge and leave the position, on net, exposed to one market or the other.

The large position size required to extract reasonable returns out of relative value trades can also cause a number of problems. First, as was

made clear with LTCM, these very big positions are hard to liquidate, and the newer, less liquid markets are usually the very ones that exhibit the spread discrepancies. 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, so he may have a long holding period.

The third risk is that you might screw up the model. 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, a model error of just 1 percent adds up to a lot of money. Yet the best opportunities are in the more complex markets—like the mortgage market—that are fraught with potential for error, and in the nascent markets where there is little liquidity and/or stability in relative relationships and little experience to back up the models.

The integrity of the hedge and the accuracy of the perceived risk, however, are only as good as the hedging models and the assumptions that go into them. Of course, no model is perfect. Most firms are confident of their swap and Treasury hedges, but only the foolhardy would consider a hedged billion-dollar mortgage derivative position to have low risk.

The accuracy of the models determines the level of confidence in the hedges. Even beyond the market risks, large balance sheet positions like those necessary for spread and relative value trades transform the secondary risks inherent in trading into devastating ones. An unanalyzed credit risk, an operational or clearing problem, a sudden collapse of liquidity—with the related inability to update a hedge—and undetected ambiguity in the legality of a financial contract are all that more troublesome when the inventory is \$5 billion than when it is \$50 million.¹

HAGHANI STEPS OFF THE CLIFF

At LTCM Victor Haghani kept buying Russian bonds after Salomon threw in the towel, and it wasn't because he had a bad model. He had *no* model. Far removed from the mathematically intensive, risk-controlled relative value trading that the team had developed at Salomon and nurtured at LTCM, Victor had started to roll the dice on trades, moving into markets and strategies where he was a novice, albeit a novice who could throw around hundreds of millions of dollars. He was betting that somehow his information

LTCM RIDES THE LEVERAGE CYCLE TO HELL

and contacts on Russia were right and that the market, which continued to sell the bonds down, was wrong. Everyone could see the odds on this bet were worsening even as prices became more attractive, but Haghani, along with some others, was caught in a trap and refused to sell out.

It's a behavior with an evolutionary pedigree. In Southeast Asia, monkey trappers stake a small wooden box to the ground with a hole just large enough for the monkey to slide in its hand. Inside the box is placed a lychee nut. The trappers wait in hiding until a monkey comes by and reaches in for the prized nut. As the trappers emerge from their hiding place, net in hand, the monkey howls and screeches, and tries to pull the nut out of the box and escape. All the monkey has to do to free its hand is let go of the lychee, because the hole is too small for the monkey's hand if it is clenched around the nut. The monkey, too focused on the reward and ignoring the danger, refuses to give up the prize, while the trappers descend with their nets.

LTCM and its compatriots who persisted in holding Russian short-term bonds, the GKO and MinFins, through the summer of 1998 were captives of the same psychology. The specter of failure was clear, yet they rode their Russian position into the August default. How clear? In early 1998 Russia had already suffered two bouts of exchange rate instability. Nonpayment of taxes was notorious; by May Russian companies owed more than 250 billion rubles—around \$50 billion—to the government's budget.

In July and August, capital outflows were continuing in spite of currency defense by the Russian Central Bank and an International Monetary Fund (IMF) package of \$20 billion. By mid-August the first tranche of the IMF package had been spent, and the Central Bank lost upwards of \$4 billion in reserves defending the currency. As the prospect of default loomed, Russian bond prices fell and interest rates on Russian debt rose. But for every extra month the positions were held, LTCM and others pulled in more than some hedge funds make in a decent year, and many of them simply could not let go of the prospect of a continuing stream of apparently found money from a simple short-term investment.

When Russia finally defaulted, many firms with holdings similar to those of LTCM posted losses and were forced to lighten up on their non-Russian positions. As they exited, prices spiraled downward, causing further liquidations. All of this had a magnified impact on LTCM because their positions were held on margin or through repos.²

A DEMON OF OUR OWN DESIGN

The problem for a firm like LTCM in using repurchase agreements is that they come due at different times for different bonds, so keeping the loan going is like keeping a bunch of balls in the air. As one bond comes back, another goes out, with the terms of the repo constantly changing based on the going interest rate and on the credit risk faced by the party making the loan. And the lenders are notoriously risk averse. They view the repo as a safe way to add a few extra basis points to the bonds they are holding. They come to the table with their shoes still on and staff in hand, ready to walk away at the slightest whiff of credit risk, so the banks increase the haircut they demand unilaterally and without warning. And if things get really bad, repurchase agreements can dry up in a flash. The lenders will not do the repo at any price and the door to the loan market will be slammed shut.

As the Russian bonds dropped, the banks' collateral for LTCM's loans also declined; they required LTCM to put up additional margin, but LTCM did not have cash reserves available to post it. The company had been running its operation full throttle; earlier in the year the partners had decided to cut their capital base by 50 percent, returning nearly \$3 billion to investors so they could deploy their own investments in the fund with more leverage. To get the money for the Russian haircut, they had to sell off some of their other positions. When they entered the market to do so, they discovered they had stepped off the ledge.

THE COMPETITION STICKS THE KNIFE IN

It is hard to keep a secret when credit departments at every major investment bank are running to their trading desks to reevaluate haircuts. So it took little time for word of LTCM's problems to get out. That added to the negative momentum. LTCM had been facing small losses for several months, which were reported to its investors and quickly leaked to the marketplace. The fund's positions were also well known. While Meriwether and company made an elaborate show of their secrecy, they traded in huge size and trumpeted their acumen with panache. Now they were like a hunted elephant trying to hide in the knee-high grass. Most dealers and large hedge funds knew LTCM was short of cash and would need to sell, and they knew what it would sell. No one wanted to be first to take the

LTCM RIDES THE LEVERAGE CYCLE TO HELL

other side of those sales. In fact, some started to move out of their positions ahead of the expected stampede.

As a result, LTCM had to drop prices dramatically to find buyers. By the time it raised enough cash from selling assets to post the higher margin required by the creditors of its Russian positions, the prices had dropped so far on the rest of its portfolio that the remaining inventory was worth substantially less. The banks providing margin for that inventory, in turn, also demanded larger haircuts. LTCM was caught in an irreversible downward cycle driven by its high leverage. Meriwether and crew were past the point of no return: Any attempt to sell required a drop in prices that lowered the value of their remaining positions by more than they could raise from the sale.

The prospects of a lethal liquidity squeeze had likely never passed through the minds of the principals at LTCM. In their former lives at Salomon, liquidity had never been an issue; there was always more capital to be had if things got tight. The senior management at Salomon understood the game; if the downturn was due to a liquidity event, there was the expectation that the arbitrage unit would get the added capital to tide it over. And if a position ever had to be sold, it could be turned over to the client desk, where it would be put in inventory or unloaded without anyone knowing whether the position was from the firm's proprietary trading desk or from outside clients.

In the freewheeling Gutfreund days, no one ever kept score on how much capital Meriwether's arb unit used. If they had a good trade, they would present it at the risk management committee, where the CEO and the chairman had broad authority to allocate more of the firm's balance sheet to the trade. If they had a good trade that hit the skids, Salomon would usually provide the breathing room to stay in. For one trade, a huge relative value trade in the mortgage market that was down over \$300 million in 1993, the arb group member who owned the position, Larry Hilibrand, went as far as presenting the arguments to Salomon's board. Convinced of its merits, the board authorized him to add to it despite its sizable mark-to-market loss. Hilibrand left not too long after, heading to LTCM later that year. The new co-heads of the U.S. arb group kept the trade on through most of 1994, reaping a profit of \$1 billion for Salomon and a bonus of \$30 million each for themselves.

WHAT WERE THEY THINKING?

The partners at LTCM presumed an ever-ready reserve of capital as a given, and apparently did not fully appreciate the implications of a corporate lifeline when they went off on their own. And why should they? As they were building up their leverage, banks were ready with cash and credit at every turn. Facing the capital squeeze, Meriwether would later try to do with his investors what he had reflexively done on numerous occasions with Salomon's board, and then later with bankers. On September 2, he sent out his famous "Dear Investor" letter:

In August, many [of our relative value trades] diverged at a speed and to an extent that had not been seen before. LTCM thus believes that it is prudent and opportunistic to increase the level of the Fund's capital to take full advantage of this unusually attractive environment. . . . Many of you have asked to add to your investment in the Fund. Since it is prudent to raise additional capital, the Fund is offering you the opportunity to invest in the Fund on special terms.

The attractive opportunity found no takers.

Without the cushion of capital and distribution they had enjoyed in earlier years, and without investors willing to take on that role, Meriwether's panicky partners now had to look for loose change throughout their positions.³ When you can't sell what you want, you have to sell what you can. They had amassed a huge position in a U.K. swap spread trade. As they unwound the positions, buying back the bonds and selling off the bank paper, the interest rate differential, or spread, between the two rose from 35 basis points to 120 basis points, the highest level in history, exceeding the previous high point in spreads by more than 40 basis points. Attempts to sell a similar position in the U.S. swap spread market led spreads to rise by a third, from 60 to 80 basis points. The severity of the moves is especially notable in context; typically these spreads vary by only a few basis points from one month to the next.

The LTCM partners prided themselves on their broad diversification, but now that it mattered everything that should hedge didn't and everything that was supposed to be uncorrelated moved in lockstep. Scandinavian inflation-linked bonds, which usually moved closely with the Scandinavian ordinary government bonds, fell to a 10 percent discount.

LTCM RIDES THE LEVERAGE CYCLE TO HELL

On the morning of August 21, the U.S. swap spread widened by 19 basis points (20 times its usual daily move). This move was mirrored in the U.K. swap, where the fund was also short. LTCM's risk arbitrage position in Ciena Corporation/Tellabs blew apart because Tellabs canceled the shareholder vote on the acquisition. By the end of the day LTCM had lost \$550 million, nearly a sixth of its capital. Pick a market—Danish mortgage bonds, Japanese convertible issues, the Mexican Brady bond/Eurobond spread—if LTCM held it, or more precisely *because* LTCM held it, the price plummeted.

And in the midst of the dislocations and exploding market volatility across the globe, in one market, and one market alone, interest rate volatility was declining. German interest rate option volatility actually fell, dropping by greater than 30 percent to near-historic lows. Why? Because LTCM had to sell option positions on German interest rates, pushing the option-related volatility levels down.

Although LTCM was in bonds and stocks in Europe and the United States, doing a dozen different types of strategies, its diversification was undermined because almost all of its large trades had two types of risk in common: credit risk and liquidity risk. LTCM had exposure to credit risk because its large bets typically were constructed with a long position in a more risky asset and a short position in a less risky one—often bulletproof government bonds. In a credit event, these spreads would widen. LTCM would lose money on the bond with credit risk as it dropped in price. And it would also lose on the short side of the trade because the Treasury bonds would appreciate as investors fled to quality. If the market perception was that the world was a riskier place where defaults were more likely, the vast majority of LTCM's positions would be hit. And the firm's instability triggered that fear. Not only did its need to sell out of other markets cause liquidity providers to stand on the sidelines; LTCM's actions also led the market as a whole to view any position as being more risky. Prices dropped accordingly.

LTCM faced liquidity risk because it had become a major liquidity supplier to the market, and was providing the other side of the market for those who wanted to hedge their credit exposure in various instruments. Just as it was generally long instruments that had credit risk and short instruments that did not, it also tended to be long instruments that were less liquid and short instruments that were more liquid. LTCM's short liquidity

position even extended to its exposures in the Treasury market, where it was long the less liquid off-the-run bonds and short the more liquid on-the-runs.

Looking at the portfolio this way, the reason LTCM made money during the previous years was not as mysterious and complicated as it had seemed. The firm profited because it offered a service to the market: It supplied liquidity and took on credit risk other investors wanted to unload. To do this it justifiably received a return for its efforts. Indeed, its returns may well have been fair market returns, considering the credit and liquidity risk it took.

However, in bearing the risk burden, LTCM—and its investors—seemed to have focused only on the return, missing the potential for downside. Once the downside showed itself, it came as a surprise that had not been anticipated and became increasingly difficult to manage. LTCM, in the end, may have been no different from the short volatility trader or credit risk trader of many other firms. In other operations these traders would make a little, make a little, then lose a lot. But because the banks were willing to provide LTCM so much leverage, they could turn this formula into one of “make a lot, make a lot, then lose it all.”

MIRROR, MIRROR, ON THE FALL

Long-Term Capital Management was neither alone in these trades nor alone in facing unexpected losses. The strategies had become well known, and what seemed like a good profit opportunity to LTCM became attractive to other trading houses. Indeed, it often looked good—and almost free of risk—precisely because LTCM was in it. Both Goldman Sachs and Salomon had large positions that mirrored those of LTCM. For Salomon, this was understandable, because all of the LTCM principals came from Salomon’s proprietary trading group. Losses mounted for hedge funds, brokers, and banks across the world.

While other proprietary trading desks found themselves with market losses, the implications for LTCM’s major creditors were disastrous. Some of the most stalwart institutions shook to their foundations. They lost not only from market positions but also from the possible default on loans to LTCM. In lieu of the billions they had lent LTCM, they would be left with only the positions as collateral, and those positions were now trading at

LTCM RIDES THE LEVERAGE CYCLE TO HELL

pennies on the dollar if at all. For a period of three days in early October 1998, Merrill Lynch had a net capital deficiency ranging between \$1 billion and \$3 billion. It cut its fixed income trading unit by nearly 40 percent. (It would later admit it might have “cut some bone with the fat” and try to hire key people back.)

Goldman Sachs was forced to delay its initial public offering, leaving its partners as mere millionaires until it succeeded the next year, and leading to the resignation of co-chairman Jon Corzine, who would later try his hand in the U.S. Senate and then as governor of New Jersey. Lehman Brothers, whose stock plunged by two-thirds over the course of the crisis, skirted bankruptcy. Salomon prepared to pull its trading lines with Lehman, and relented only out of fear that doing so would initiate a cascade from other firms that would turn a potential for default into a certainty. UBS's strategic relationship with LTCM, which made it the hedge fund's largest investor, piled onto other senior management missteps, forcing the venerable institution into a merger with Swiss Bank. The proprietary trading business at Salomon lost \$1.2 billion in August and September, more than it had made the previous year. The bulk of the loss was in Europe, approaching \$900 million.

To a large extent, the losses at Salomon were the result of inexperience and faintheartedness on the part of the new Travelers-based management. The financial markets were not in ruin; there was no reason for the incredible drop in prices. It was a classic liquidity-driven market crisis. LTCM was cash-strapped; Meriwether had daily margin calls and needed to sell off his portfolio to come up with the cash; Salomon Smith Barney did not. With its deep pockets, Salomon/Citigroup could have weathered the storm, but a sense of panic pushed senior management to reduce positions almost in lockstep with LTCM. Still, for all of its losses, Salomon came out relatively unscathed. The firm had no Russian positions at the time of the default, and little in the way of counterparty commitments and related credit overhang to LTCM, and had limited its exposure with a decision to close the U.S. proprietary trading group in early July. Had those decisions been different, the losses might have doubled or tripled.

Nonetheless, the political carnage did not bypass Salomon Smith Barney. The former Salomon CEO, Deryck Maughan, who tried to distance himself from the crisis with a politically astute, studied aloofness, was demoted to the catchall title of vice chairman of Citigroup. A friend of Sandy

Weill's, he would hang on at Citigroup to bounce back in a number of senior-level roles, and even be knighted by Queen Elizabeth in 2002. Sir Deryck finally would pay the price in October 2004 when embarrassing problems in Citigroup's Japan operations surfaced. Maughan's co-CEO, Jamie Dimon, was forced out even though he had been the most vocal in reducing risks, perhaps a lingering result of his contentious relationship with Weill's daughter, Jessica Bibliowicz, or perhaps in a move intended by Weill to cut off any future competition for the chairman's job at Citigroup. Weill had scaled the pinnacle of the financial services industry. Shored up on all sides by a corporation of gargantuan proportions, one more fellow climber was cut from the safety line. Dimon would later return to New York by way of Chicago to take the CEO job at Citi's New York archrival, JPMorgan Chase, shortly after it acquired Bank One.

A critical decision point for LTCM actually happened a year before the firm's failure, when it decided to give back client capital so that the partners could take more of the trading revenue for themselves. This move was hotly contested. Two of the most aggressive and outspoken partners, Larry Hilibrand and Victor Haghani, wanted to build up "dynastic" wealth—apparently, over-the-top wealth didn't quite cut it—and to do that they wanted to lever their holdings as far as they could. Bob Merton argued that the firm could maximize its value by leveraging its name rather than its capital. LTCM had a mystique that extended even beyond its record. If it took its trading ability and married it to large funds, it could make small returns, but on a far larger capital base, and do so with much less risk. And having the stability of a huge capital base and marquee value, the company could go public for a windfall far in excess of what its trading alone could make.

Adding to the persuasiveness of this approach was a reality the partners knew better than most: The party was over for relative value trading. Competition had seen to it. At first just Salomon and LTCM chased relative value opportunities. But by 1997 a host of other hedge funds and investment banks were converging on convergence and spread trades, slicing into the returns for both LTCM and Salomon's proprietary trading. LTCM had earned upwards of 40 percent in both 1995 and 1996, saw its returns cut by more than half in 1997, and by the spring of 1998 started to see substantial negative returns.

LTCM RIDES THE LEVERAGE CYCLE TO HELL

As usual, Hilibrand and Haghani ruled the day. Eric Rosenfeld, who often cast the tiebreaking vote in these matters, was a peacemaker who acceded to their view. Meriwether eschewed his position of leadership and declined to confront his partners.

LOUSY WITH LEVERAGE

Why is it that leverage crises occur frequently in financial markets, but not in the broader business world? The reason is that the very ability to liquidate—clearly a desirable attribute of an investment portfolio—is, ironically, the root of the liquidity crisis cycle. Had LTCM been a corporation where prices were not reported by the second, where inventory could not be put into the market to be sold at a moment's notice, and where the nature of the inventory was not broadly known, the downward cycle that ushered in the demise of the firm and spread market panic across the globe would have been averted. But then again, if it had been such a corporation it also could not have levered to the hilt.

For all of its sophistication, LTCM used simple historical analysis to assess the trading opportunities for unleashing this leverage. The partners identified market relationships that they believed had long-term stability, such as the interest rate differential between Treasury bonds and high-grade corporate debt. They then compared the current level with the historical levels. Predicated on their conviction that the relationship had long-term stability, they would take positions based on the assumption that it would return, or converge, back to that historical value. What they did not appreciate was that they had changed history: There had never been someone trading hundreds of billions of dollars in the middle of this relationship before.

LTCM had scrupulously modeled and monitored its market risks. It prided itself on having broad diversification across the globe and across markets. The firm's risk models looked at past price variability and provided the partners with assurance that they could survive the financial equivalent of the 100-year flood. The problem was that their models assumed they were in a "game against nature" where their decisions did not alter the playing field. In a normal market environment, with small players, this is a reasonable assumption. But as the largest player in a world of

looming illiquidity, this worldview was naive at best. Their actions did change the game, because the decisions of other traders would change depending on the actions LTCM took, or was perceived to take. The partners looked at their risk as if they were playing a game of roulette, where the possible outcomes were unaffected by what was bet and how much was bet. The market turned out to be more like a game of poker, where the outcomes depended on the behavior of the other players, and whose behavior in turn would change in response to their opponents’.

The one risk Meriwether missed was that his firm could cause the flood. As he later noted, “The hurricane is not more or less likely to hit because more hurricane insurance has been written. In financial markets this is not true. The more people write financial insurance, the more likely it is that the disaster will happen because the people who know you have sold the insurance can make it happen.”

This feedback of positions on the market price exists even with small positions and trades, but only becomes evident with larger positions, and only becomes problematic when a crisis dries up liquidity or forces trades to be done quickly at any price. Or, worse yet, when large positions and low liquidity combine with having the bulk of the market participants leaning over your side of the boat. In this circumstance, as Victor Haghani observed, “it was as if there was someone out there with our exact portfolio only it was three time as large as ours, and they were liquidating all at once.”⁴

Over the years during which LTCM held its preeminent position, Hilibrand parceled out the financing of positions across banks in a way that kept each unaware of his overall strategy. One bank would finance one leg of a spread trade, while another bank would handle the other leg. The objective was to prevent the banks from seeing the strategies. This might have bordered on seeming paranoid, except for the fact that it was merited. Everyone on Wall Street was trying to figure out just how LTCM did it. Protecting the process was easy within the recesses of Salomon’s trading floor, but LTCM needed outside financing, and the firms that provided it insisted on understanding where their money was going. LTCM, therefore, had to open the kimono. Even with Hilibrand’s efforts, over time it became increasingly difficult to execute a trade without that trade revealing the new opportunity, and then before long the opportunity would vanish.

LTCM RIDES THE LEVERAGE CYCLE TO HELL

While opaqueness may have actually been beneficial in normal times, it was a different story when the firm was on the ropes. Short-term lenders have a stunted sense of risk-return trade-offs. Unlike commercial banks, whose creditors can look to the Federal Deposit Insurance Corporation (FDIC) or to the “too big to fail” doctrine, securities firms have no declared sugar daddy to deter runs. It is not a matter of simply paying a higher price if lenders perceive that their capital is at risk. In fact, waving a premium rate in front of them can be counterproductive; it makes them suspicious.

Since no bank knew the other side of the position they were financing, they treated the position as an outright trade, and required multiples more in margin than they would have required if they had known they were financing a spread. Hilibrand had gotten it completely wrong. LTCM was opaque to the very institutions that could save its bacon. Simultaneously, the trading desks at the major investment banks, having dealt with LTCM for years, could see right through it, and the fact that LTCM needed to sell made it more expensive to do so—so expensive, in fact, that each sale precipitated a need for further liquidations.

The liquidity providers that had the ability to take on the firm’s positions, and might have done so in a less charged environment, elected to sit out. And worse still, several firms with capital and sharp teeth hunted LTCM like it was an injured animal, trading against LTCM’s positions, racking up profits each time LTCM tried to cover its margin needs, and further precipitating its need to do so. For these outfits, LTCM had gone from being the master of the game to becoming the game.

THE JAPANESE TAKE UBS TO THE CLEANERS

The failure of LTCM and the resulting turmoil in the markets shook other firms like Salomon Smith Barney that had been following LTCM’s lead in relative value trading, and it reached the banks that had been extending the credit to LTCM that had allowed it to lever its book.

One firm that was swept over the edge was UBS. (Although UBS continues, it is now the product of a forced merger into Swiss Bank, which then decided to retain the name.) In the land of precision time-pieces and unerring train schedules, where fastidious control and attention to detail are national characteristics, UBS had a series of lapses in

risk management, culminating in a transaction with LTCM that served as the coup de grâce.⁵

UBS managed to get itself hopelessly tangled in the two signature financial crises of the decade, LTCM and the Asian currency crisis. The tale began innocently in the mid-1990s when a number of Japanese banks, decimated by a languishing stock market and expanding credit woes, seized on an innovative structure to increase their capital base. Called preference shares, these were convertible bonds with a twist. Like conventional convertible bonds, they could convert into stock, but rather than the conversion depending on the stock appreciating, these bonds converted to stock at a set time in the future no matter what the stock price.

UBS bought into this wacky scheme, which was yet another Japanese stab at inventing financial structures to provide free capital—akin to building a perpetual motion machine. Initially, the issuers failed to realize that issuing more stock, whether in the present or in the future, would inevitably dilute the value of existing shares and drop the stock price an amount proportional to the capital they received. If a company doubles the amount of shares it has outstanding, the value of each share is thereby cut in half.

This mistake was compounded by the features of the offering. If the bond simply mandated a purchase of a fixed number of shares at some time in the future, it would be nothing more than a forward sale of the stock. That would be too easy—it would just appeal to those who would buy the stock today and would transparently dilute the stock price. And it would be risky, because the buyer of the bond would be subject to variations in the stock price between the present and the exchange date. So as an alternative, the preference shares promised the holder a fixed yen amount of stock. If the stock doubled in price, the holders would get half as many shares; if it dropped 50 percent, they would get twice as many shares. The payoff, then, though paid in shares of the company, was fixed in yen terms, just like the payoff of a regular bond.

The investment banks that cooked up this idea realized there was a problem in the structure if the stock dropped too far: The exchange can potentially be infinitely dilutive. For example, if the stock were to drop by 90 percent by the time of the exchange, suddenly the exchange would net the bondholders 10 times more stock than was initially forecast, much to the detriment of the existing shareholders. As a safeguard, a cap was put

LTCM RIDES THE LEVERAGE CYCLE TO HELL

on the maximum number of shares that would be exchanged. This cap meant that if the stock dropped too far, the bondholders would become subject to the risk of the stock price. The cap was far below the current stock price—the stock might have to drop by 60 percent or more before it would kick in—so as far as things went, this feature seemed like a mere footnote in the overall terms of the issue.

The Asian crisis that started in Thailand in July 1997 turned the footnote into a headline. The shares of Fuji Bank and Sakura Bank, two of the largest issuers of these bonds, dropped more than 60 percent from the start of July to year-end. As the stocks dropped, the option feature of the bonds became more and more critical; rather than being a bond, the preference shares were starting to look like stock, at least on the downside. If the shares dropped further in price, the bondholders would lose one-for-one with the drop.

This risk made the preference shares illiquid, and no one had more of these sitting in inventory than UBS. No one would take them off the bank's hands; the only alternative was to try to hedge out the stock risk by finding a position that would move opposite to the preferred shares, rising if the stock price declined. Because UBS would lose money with a drop in the stock price, the obvious choice for the hedge was to sell stock short. Naturally, selling stock—whether it is selling stock you already own or selling stock short—puts downward pressure on the stock price. The attempt to hedge depressed the stock price further, increasing the losses and requiring yet further selling. The net result was a longer-term, slow-motion version of the portfolio insurance-induced market crash in 1987.

The situation was complicated by the institutional structure of the Japanese markets. The first step to execute the hedge for the ill-fated preference shares was to find stock to borrow. This was easy most of the time, because the preference shares were issued by large banks with plenty of stock outstanding. The problem was that in Japan companies must hold the stock in their portfolio two times a year, at fiscal year-end on March 30 and at the half-year point in September. All stock that is borrowed must find its way back home during these times, so all short sellers have to cover their positions. At both of these points, twice a year, the stocks would rally as the shorts bought back the stock to return to the lenders.

Then came the double whammy. In September 1998, credit spreads went through the roof because of the LTCM crisis. And the crisis fell near

the midpoint for the Japanese fiscal year, so the stocks needed to create the short hedge could not be borrowed in any size. The stocks hit bottom on September 30. Tokyo-Mitsubishi dropped by nearly two-thirds, to 800 yen from more than 2,200 yen a year earlier, and Sakura declined by some 75 percent, dropping from 900 to under 200. When the stocks became available again for shorting, the actions of the hedgers buttressed the stock prices and actually pushed them up from their lowest levels. But by that point the shorts had locked in their losses. It cost UBS more than \$500 million.

RAMY GOLDSTEIN TAKES UBS TO THE CLEANERS

The master of this disaster was Ramy Goldstein. Goldstein was a rugged, high-energy former Israeli paratrooper who did graduate work in finance at Yale before landing for an assault on Wall Street. He made sizable profits from taking advantage of the sheer complexity of UBS's financial organization by what amounted to arbitraging UBS's internal financial labyrinth. UBS had about 40 different financial entities, each of which provided financing based on its own capital structure. An entity in Latin America might have billions available to lend, money that, because of various tax or capital restrictions, it could not repatriate to the parent. Another entity might need to go into the market to borrow to finance its enterprises. Goldstein used his trading desk as a switching station for these capital variations, where he would essentially borrow from one part of UBS and lend to another, with his desk pocketing the spread. (This strategy was employed in different guises by a number of trading desks on Wall Street.)

Taking a page out of Sheldon Epstein's index-amortizing swap book on how to trade at UBS, Goldstein also made profits within the obscuring complexity of derivatives. He used his own models to value the securities he was selling, and UBS paid him based on the profit these model prices implied. The losses approached those attributed to Epstein. But Goldstein was just getting started.

With new blood constantly moving into Wall Street, the half-life of lessons learned is short. Goldstein came into UBS and set up a global equity derivatives division with a new take on portfolio insurance, the engine of the 1987 market crash, as its primary product. Rather than having the customer take the risk of the dynamic hedge misfiring, he created long-term

LTCM RIDES THE LEVERAGE CYCLE TO HELL

options tied to government bonds that guaranteed the payoff that the customer desired. The hedging—and the risk of mishedging—was taken directly by UBS. The call option that Goldstein supplied was long dated, with maturity usually around five years.

The typical product that Goldstein's group sold after he had his group up and ready to trade in 1992 gave the investor the upside of the index like the FTSE Index (the market index in the United Kingdom) at the end of five years, and guaranteed that if the FTSE fell the investors would at a minimum get their money back. Of course, even the return of the principal implied a loss given the time value of money—if the funds had been invested over that period they would have ended up with more than their initial investment—but nonetheless it was very attractive: At the best you make money with the equity market and at the worst you end up with your money back.

Goldstein also offered long-dated options on individual stocks. He executed this business by buying up convertible bonds, “stripping” the option embedded in the bonds, and selling it to the equity desk to then pass along to clients who wanted long-term options on individual stocks. In theory, this made sense. The convertible bond did indeed have a call option that kicked in if the underlying stock price rose high enough. The convertible bondholder would end up getting stock in exchange for the bond. The problem was knowing how to value that option and knowing how to hedge out the ancillary risks of the convertible bond, most notably interest rate and credit risk.

The basic mathematics and procedure for executing these sorts of trades were well known. It was also well known to those who lived through the crash just 10 years earlier that things did not always work out, and that if you had to sit on the risk for five years there is all the more chance that something can go wrong. For one of Goldstein's primary products, what went wrong was a change in the U.K. tax code in 1997 that negated a tax break critical to Goldstein's pricing. He had priced the product assuming the break would remain intact over the full maturity, going out five years or more. Its elimination adversely affected the cost of hedging.

With the tax break gone, UBS was saddled with a loss of nearly \$100 million going out of the gate. And years still remained for Goldstein to continue to delta hedge and wait for other problems to occur—which didn't take long.

When the Asian economies dropped, one after the other, a linkage between these markets and other equity markets became evident. The result was an increase in volatility in Europe and the United States. Volatility had been at a low point for several years, but now it suddenly moved up. This posed a problem for UBS. UBS needed to use shorter-dated options to effectively hedge the longer-dated options imbedded in the structured products it had issued. The cost of buying these options had increased because of their higher volatility.

Coming on the heels of the U.K. tax change and rampant rise of equity volatility yet a third problem loomed. Goldstein, searching for broader markets and products to expand his franchise, had stretched his business into more exotic types of instruments, including structured products that provide a payoff based on the correlation between different stock markets. These options were particularly difficult to price and monitor because correlations from one equity market to another are not very stable. This instability was exacerbated as the volatility of the markets increased. The correlations changed unfavorably; under further scrutiny by researchers in Goldstein's group it was discovered that the underlying models he had used at the outset were flawed. Basically he had been selling these products below their fair price even in the best of times, and now as disaster struck, this mispricing became worse at every turn.

True to form, UBS did not have any independent method of pricing the options or monitoring their risk. Just as it had with Epstein's trades, UBS took Goldstein's models as the basis for booking immediately the profits of strategies that would not be fully resolved of risk for many years. Nor was Goldstein's compensation so dependent; he earned a bonus of upwards of \$10 million in 1997 alone. In a way, UBS was lucky that so many problems came as early as they did. Had the business continued without event for a few more years, the book and related losses could have been multiples of what was already a sizable deficit.

LTCM TAKES UBS TO THE CLEANERS

With the UBS traders so active in generating huge losses, it was only a matter of time before the senior management of UBS took up the sport, this time under the tutelage of none other than LTCM, which offered to provide the company with insight into its trading and risk management practices.

LTCM RIDES THE LEVERAGE CYCLE TO HELL

The partners in LTCM not only ran a highly levered fund for their investors; they doubled up on that leverage by borrowing money themselves to put into the fund. Hilibrand borrowed \$24 million from Credit Lyonnais; Hans Hufschmid, formerly the head of foreign exchange trading at Salomon and a more recent addition to the LTCM partnership, personally borrowed \$14 million, and a few others with either less wherewithal or more prudence borrowed less.

Myron Scholes, a Nobel laureate and latecomer into the LTCM hedge fund partnership, championed another approach to pile leverage onto their positions: a warrant that would pay off as the fund increased while limiting the partners' liability if the fund value dropped. They shopped the warrant to Merrill Lynch and to Chase, and both firms demurred. Spurred on by the prospects of further riches—and in the case of the warrant, apparently tax-advantaged riches—the warrant was an ongoing imperative for the partners. But if viewed rigorously by the counterparty, it would be seen as a scheme that could expose the seller to huge, uncontrollable risks.

With most warrants, the seller can hedge the risk by taking a dynamic position in the underlying asset. This could not be done with a warrant on the future earnings of LTCM, because that was not a traded security. Even worse, there was the problem of moral hazard. LTCM could change the nature of its strategy or leverage at any time; the risk of the fund could be increased at its whim. LTCM would benefit if that increased risk bore fruit, while the warrant writer would be left holding the bag if things blew up. The very issuance of the warrant would in fact increase the incentive for the hedge fund to do so, because it was increasingly gambling with other people's money. And for taking this risk with a possible loss of \$200 million—the seller would get a premium of only \$15 million.

In fact, the rise in leverage to the detriment of the warrant writer is exactly what did occur. And when it did, it was not Merrill or Chase on the other side, but UBS. LTCM hit pay dirt by inviting Mathis Cabiallavetta, the UBS CEO, to its Greenwich offices. For the starstruck senior management of UBS, it was like a homely wallflower being asked out to the prom by the captain of the football team.

UBS put itself on the hook for \$800 million. The warrant changed from its earlier incarnation not only in size, but also in some of its terms. LTCM agreed to pay UBS \$300 million, and in return UBS agreed to pay

the LTCM partners an amount equal to any upside return that would accrue to an \$800 million investment in the firm over the following seven years. If the hedge fund value dropped, then UBS could convert its holding into a seven-year note—although the payment of that note still required that LTCM be solvent at the end of the seven years, so UBS had merely replaced the market risk in LTCM for pure credit risk in the hedge fund.

UBS could have taken its \$300 million premium and just sat back, but then it would have been liable for the potentially huge returns that LTCM—if it continued to run true to form—might rack up. Instead UBS chose to buy an \$800 million stake in the fund. UBS was then hedged against any upside gains, since its investment exactly offset its liability. But while hedged against performance on the upside, that investment now led UBS to be exposed on the downside. If LTCM failed, UBS would be left with only its initial \$300 million premium; it would be out \$500 million.

Its \$800 million stake in the fund made UBS LTCM's largest investor. Then, oblivious to the risks, UBS doubled up its bet by taking the \$300 million payment from LTCM and then turning around and investing it in the hedge fund, too. This made it the largest investor in LTCM by far, and left it facing the potential of a billion-dollar hit if the fund failed.

But knowledgeable or not, this masterstroke resulted in a loss of more than \$600 million for UBS when LTCM hit the skids. After having stumbled from one loss to the next, UBS was history. More from an attempt to hide the embarrassment of the loss than anything else, UBS merged with Swiss Bank. Cabiallavetta was out, and the reins of management moved to Swiss Bank, the smaller of the two banks. Six months later Cabiallavetta emerged as a vice chairman of Marsh & McLennan, a company with a senior management of castaways. Its chairman, Jeffrey Greenberg, the son of American International Group, Inc. (AIG) chairman Hank Greenberg, left his father's firm in 1995. And Cabiallavetta was soon to be joined in the vice chairmen's ranks by Heidi Miller, who boarded Marsh's insurance brokerage division to cap off a bout of career moves that started with jumping ship from Citigroup only to land adrift at the struggling Priceline.com, Inc.

SALOMON SMITH BARNEY LOSES ITS NERVE

We weren't quite finished losing money, either. By September 1998 the usual monthly risk management meetings moved up to once or twice a

LTCM RIDES THE LEVERAGE CYCLE TO HELL

week, and the march between the Salomon building at Seven World Trade Center and the Smith Barney building at 388 Greenwich Street quickened its pace. The meetings were spent enumerating most of our proprietary trading losses and what had been done to reduce our exposure.

One of these was a U.S. Treasury spread position, an on-the-run versus off-the-run trade that had moved from being a sure thing to generating a loss of \$100 million or so for the client desk. Unlike the fixed income arbitrage trading units, for whom risk taking and staring down large losses is all in a day's work, the client desk is supposed to do nothing more than move the merchandise. It earns money by buying bonds from clients who want to sell and then turning around and selling the bonds to others who want to buy, pocketing a fraction of a point on each turn. Call it the used bond department. I could see that this position was entirely salvageable—if we could somehow get management to play along.

This wouldn't be easy, because the rules for convergence and other relative value strategies defy financial convention, and by now the Salomon Smith Barney culture wasn't simply risk averse, it was risk repulsed. The usual rule in sizing trades, drilled into every aspiring young trader, is to cut losses and add to winners. But it's the opposite for yield curve and relative value trades. You actually add to the losing positions and cut the exposure to ones where profits have piled up. The reason is that these trades are reversionary: They are executed because some relationship is out of line, and they cash in when things go back to normal. If the strategy is correct yet still losing money, it means the prices are getting even more out of line, and thus more attractive. Perhaps investors do not know where the price really should be, or perhaps they know what the price should be but are constrained from acting on the mispricing because of taxes or institutional limitations. In any case, the mispricing will usually be a short-term phenomenon. The smart strategy is to fasten your seatbelt and continue to hold or even add to the position.

When people finally get it, prices will converge. Even if some mispricing persists, the position will make money through what is called the carry of the trade: The income from the short side of the trade will exceed the cost of financing the long side of the trade. In the case of convergence, all the profit comes in one fell swoop from the price adjustment. In the case of ongoing mispricing, the profit just takes more time in coming as, coupon payment by coupon payment, more money flows into the strategy

than needs to be paid out in financing costs. It can be a hair-raising game of playing chicken with the market, putting on a bigger and bigger position when the trade is going the wrong way. But with faith in the model and reserves of capital to sustain the position, that's what it takes to get the most out of a relative value trading opportunity.

The problem is that this faith can never be total. An alternative explanation for losses on the position is that there is an error in the model. Perhaps it is missing a critical factor that is driving interest rates or it is using a misestimated parameter. Perhaps what is thought to be a hedge position is not really hedged after all, and the losses are not due to persistent errors on the part of other traders but due to an undiscovered, unintended factor exposure—a leak in the system. It happened to us in the Salomon U.S. fixed income arb group and it happened to LTCM.

Market aberrations mean more opportunities; model-based aberrations mean more risk. It is a tension that is not limited to just model-driven yield curve trading. It is even more endemic in the mortgage market, where prepayment can depend on factors ranging from media coverage of lower mortgage costs to changes in the technology for applying for refinancing. Or the model could be using the wrong functional form. The portfolio insurance models in 1987, for example, failed to incorporate the feedback from the portfolio insurance hedging into market prices, while in the U.S. fixed income arb group, the failure in their model was never found, even after months of frantic searching.

A convergence trade is made between two different securities that tend to become more similar over time. Because they differ from one another, there is day-to-day risk in the trade—mark-to-market risk. But over the longer term the prices will tend to move closer and closer to each other because it will become increasingly obvious that the two securities are similar. Convergence trades usually take time—months, perhaps even years. If the convergence in price were imminent, then the market would be likely to recognize the price divergence. Things can get worse before they get better, and stay worse for months at a time, so for these trades to be successful you have to have a high degree of confidence and the staying power to stick with them.⁶

The Treasury bond spread trade that the client desk had devised fit the profile, and I proposed to the other members of the risk management committee that we take an aggressive position. By my calculation, the pair

LTCM RIDES THE LEVERAGE CYCLE TO HELL

of 30-year Treasury bonds in this trade had a relative mispricing of more than 80 basis points, nearly 1 percent. That is massive in this kind of bond, yet the two were virtually identical up to the last of their coupons, which would be paid 20 years hence.

This trade was a variant of one of the simplest, though potentially painfully long, convergence trades. It involved taking offsetting positions in on-the-run and off-the-run Treasury bonds. The U.S. government issues 30-year Treasury bonds on a semiannual basis. The most recently issued bond, termed the on-the-run bond, becomes the standard for trading. When traders enter orders to buy or sell the 30-year, the on-the-run bond is the bond they mean. This trading demand means the on-the-run bond enjoys a liquidity premium—that is, it trades at a slightly higher price than its less interesting older brothers, which, though once on-the-run bonds themselves, are now beyond their six months of fame. These bonds, termed off-the-run bonds, have less liquidity and are not in as much demand; devoid of the liquidity premium, they trade at a lower price. But their coupon payments and principal repayment are just as secure as the current on-the-run, so if the price of the on-the-run and any of the off-the-run bonds strays too far, a convergence trade can be structured by buying the cheaper off-the-run bond and selling the higher-priced on-the-run bond. And these two had definitely strayed off the farm.

It was all but inconceivable that this disparity in pricing could persist. And, equally important, the trade met an old Salomon/LTCM hurdle that a position can be justified only if there's a reason for prices to be out of line, with "the market is just stupid" not being an acceptable one. There was an economic justification underpinning the disparity. What is more, Citigroup unquestionably had the capital to hold the positions for however long it might take to move back to fair pricing.

In the simpler Salomon days, when all the committee members, from the head of the trading division to the chairman, were sitting in the room together, the point would have been easily made. But we were now a few months into the Citigroup reign, and with the added layers of that firm those with more trading acumen got the point while those with more of a political focus got confused. Charlie Scharf, Salomon Smith Barney's CFO, looked at me like I had three heads. How could I seriously propose adding to a position that had already posted losses of over \$100 million? The markets were in turmoil, sure, but you could see Scharf thinking,

A DEMON OF OUR OWN DESIGN

Here we are trying to demonstrate the prudence of the committee to the top management and the head risk officer wants to go farther out on the ledge? No thanks. My proposal did not sit well with the new corporate-minded structure that was being erected, layer after layer, over the Citigroup behemoth. Politically, no one was willing to take ownership of the risk, because if the trade did fail, all the others around the table in the expanded and less fraternal Citigroup risk management committee might be more than happy to use that for their own gain.

Ignoring the bank's clear capital advantage, the Smith Barney contingent not only rejected my idea of adding to this trade; they insisted we follow the lead of LTCM and others who went lemminglike over the cliff. So we confidently locked in hundreds of millions in what would have been transitory losses, rather than make a similar amount of profit. In contrast to Salomon, where risk taking—researched, structured, and imaginative—was its lifeblood, Citigroup was a company of organization men. And, beyond any measure of doubt, the organization had won, although it would see its own problems follow.

CHAPTER 7

COLOSSUS

The tagline of ads heralding the creation of Salomon Smith Barney within Travelers was “We’re Just Getting Started.” In his meeting with the managing directors to announce the Travelers-Citibank merger, Jamie Dimon played off that line to signal that the summit had now been reached by flashing a slide that said “We’re Done.” It would be a tall order indeed to top the deal that created Citigroup. And it would be a creation that required total managerial focus just to maintain.

Sandy Weill’s assemblage of the Citigroup colossus was his second act after having bootstrapped his securities firm, started in 1960, into Shearson Loeb Rhoades, which he then sold in 1981 to American Express, a move that catapulted him to the president’s office. A Jewish kid from Brooklyn, he chafed against the Brooks Brothers bureaucracy of Amex and left in 1985, loser in a power play. Act II started the next year, when, at

age 53, with a young Jamie Dimon on his team, he bought Commercial Credit, a second-tier Baltimore consumer lending operation. They leapfrogged from one company to the next, Weill using intuition and connections to find cash cows with inflated expenses that Dimon's disciplined management style could trim to profitability.

Commercial Credit bought Primerica and its Smith Barney brokerage firm, and then Travelers insurance. Through Travelers, Weill in 1993 bought Shearson back from American Express and four years later bought Salomon Brothers. This conglomerate of companies and cultures is the final entity that Weill merged with Citicorp to form Citigroup. The result is a company with around 300,000 employees and operations in 100 countries.

There was a lot to address in my little patch of this giant forest. I came out of the merger with the same set of responsibilities for market risk management I had at Salmon and then Salomon Smith Barney. For all practical purposes there were unlimited resources unleashed for measuring and monitoring risks, which was great if the problems arose from previously identified or easily identifiable risks. But the most dangerous problems come from unexpected corners. My great concern was that the sheer complexity of Citigroup would add so much structural uncertainty that it would become nearly impossible to react to events that were not already on the radar screen.

My anxiety was not relieved by the many-layered structure that took shape. At Salomon I never had a staff of more than 10 people, and my budget was less than \$5 million. The risk management committee totaled eight. The lean structure meant less detail but more attention. Less could fall through the cracks. The Citigroup risk organization swelled. Some 200 people made up just one of several risk groups, Risk Architecture—those who would build the risk measurement and reporting systems. The budget for risk management climbed past \$40 million just for the fixed income division's share. More than 20 people sat on the risk management committee. Meetings required up to four video hookups. If the objective had been informed discussion and decision making, the conversation would have been totally unmanageable.

The complexity of an organization of this scale has obvious consequences. More layers of management put distance between data gatherers and users in senior management. Information gets lost in the compression

and transmission of the data. (I had already seen this problem when I tried to get Citibank to take on Treasury positions once Salomon became part of the Citigroup conglomerate. What would have been a natural “one-click” decision if it had been a lean firm was stymied by layers of management and the attendant politics at Citigroup.) And having more people in the process inevitably means junior people, who have fewer skills and less authority and are further from the action, unable to ferret out anomalies that are beyond the reach of the existing reports.

I also believed that the consolidation would have adverse consequences for liquidity—consequences beyond a reduced willingness to take on proprietary risk. While the market structure afforded close-to-instantaneous execution, consolidation diminished the ability of the investment units to absorb that speed. When two firms merge, they do not end up with a trading floor equal to the sum of the size of the premerger entities, or two of anything: They’re left with one 10-year bond trader in the merged firm. And the capital allocated to the trading and the risk taken by the desk is not the sum of the two firms, either. In fact, it might be smaller than in either of the firms before they merged. My concerns about the complexity of Citigroup were quick to find validation. But it came, as do most risks of real import, from an unanticipated source: not market or credit risk, but from regulatory and legal missteps linked to the firm’s organizational complexity.

SANDY STEPS IN IT WITH BOTH FEET

For Sandy Weill, 2002 was a year when nothing seemed to go right. True, he had triumphed in forming Citigroup in 1998 and had vanquished Citicorp chairman and CEO John Reed two years later to take unquestioned control of the behemoth. His status as a business star had been cataloged in a book, *The King of Capital: Sandy Weill and the Making of Citigroup* (John Wiley & Sons, 2002), stacked in every bookstore’s windows within throwing distance of Citigroup’s Park Avenue headquarters. And he had maintained his perennial position in the society pages of the *New York Times*. Yet he seemed to get a leg tangled in every scandalous vine that surfaced in that tumultuous year: For Ken Lay, Jeff Skilling, and Enron Corporation, the faux new economy model, Citigroup was a leader in providing the off-balance sheet structures that were passed off as debt. For WorldCom,

Bernie Ebbers' house of cards built on scores of mergers, Citigroup's investment banking arm, Salomon Smith Barney, was the engineer. Meanwhile, Citigroup's Associates First Capital Corporation had to pay a \$200 million fine for its predatory lending practices. And representing the improprieties committed by analysts that had cost shareholders billions of dollars, there was Salomon Smith Barney's Jack Grubman, who acted as a shill to keep clients clamoring for increasingly questionable WorldCom stock.

Grubman became the poster child for everything wrong with investment banking and stock analyst ethics. His distorted if not disingenuous view of WorldCom had helped push Salomon into the arms of Travelers, and thus Citigroup. His advice on the failed MCI/BT merger contributed to a disastrous trade for the equity division; he then doubled down with the MCI/WorldCom merger. Grubman remained loyal to WorldCom to the bitter end, when, in the wave of bankruptcies in the Enron era, WorldCom started collapsing under its own weight. He remained unswerving in his buy rating—his highest recommendation—and scored WorldCom as having just a medium level of risk as late as March 2002. In about 50 research reports following the MCI merger, when WorldCom went into overdrive, he maintained a target price for WorldCom that was often 40 or 50 percent above the current level. As the price descended from above \$85 a share down into single digits, he boldly forecast that the stock price would triple. In Grubman's world, this stock was "a must-own," the "single best idea in telecom." He viewed the stock as "dirt cheap" and wrote that investors should "load up the truck"; he declared that money managers who passed up the opportunity might want to consider a different line of work.

Steadfast in adversity, even after the Securities and Exchange Commission (SEC) started its inquiry into WorldCom's finances, Grubman stonically held on to his buy/medium risk rating, reconnoitering with Ebbers to mold the analyst pitch to keep the company above water. When he did retreat, it was only to amend the rating to buy/high risk. Was he taking a stand or were there other motives at play in this dogged determination, such as \$100+ million in revenues that accrued to Salomon Smith Barney in representing the well-touted WorldCom in dozens of deals?

In this, Grubman was just the supersized version of what had by that time become a pervasive problem: Stock analysts got rewarded richly for playing to the investment bankers rather than providing marginally better

analysis to investors. For the nearly 1,200 companies covered by Salomon Smith Barney, there were hardly any sell recommendations during the early 2000s, few that were even so negative as “underperform.” Even firms in the throes of bankruptcy couldn’t get a thumbs-down. One Salomon Smith Barney analyst related a conversation with an institutional investor who “just thinks that we make ourselves look stupid by recommending names right up to the point of bankruptcy.” Or, as Grubman himself put it, “We support pigs.”¹

STANDING TALL

Swimming in the middle of this mess was Dave Bushnell, who had my old job as the head of risk management when all of this was occurring. Bushnell looked the part of the corporate risk manager, bald and bespectacled, with thin lips and sallow skin that lent him the aura of a Dickensian bookkeeper. He moved into this position after having run Salomon’s finance desk. The finance desk of a trading floor is like a car’s radiator: No one thinks much about it unless it stops working. It focuses on the very short end of the yield curve, beginning with overnight financing. When the long-bond traders or others on the desk take a position, they turn to the finance desk to get the funding and the leverage for the trade. The finance desk sweeps up the loose change, shaving a basis point or two off the funding rate, saving one trader a little bit here, another a little bit there. Bushnell’s experience at Salomon made him the furthest you could get from a cowboy while being the closest you could get to a trader. With credit and market risk equally critical to Citigroup, his familiarity with the ins and outs of the trading desk made him a natural choice to pilot the risk management effort in this new environment.

Citigroup’s Enron-related loans passed by Bushnell for approval. As the Enron transactions became more creative and suspect, he found himself drawn into the fiction. In memos, he sugarcoated the implications of the Enron transactions. An unabashed description of what was going on might have been “We’re making this up, and if anyone finds out, we’re in big trouble.” Bushnell’s version was that the “accounting is aggressive and a franchise risk to us if there is publicity.”

The memos, and shortly thereafter Bushnell himself, made their way to the U.S. Senate’s hearings on Enron, and the similarity between the two

forms of expression was not lost on the committee when they questioned Bushnell.² Things went from bad to worse for him, and I couldn't help thinking that had events taken a slightly different course, I might have been sitting in his place. His testimony had a bit of a Watergate "I cannot recall at this point in time" ring. Responding to grilling by Maine Senator Susan Collins about who knew what at Citigroup, Bushnell responded with "I wish I could recall" and "I can't recall it," and then, after the baton passed to Michigan's Carl Levin, "I can't recall verbally saying . . ."; "I wish I could recall that; I really do"; "I can't recall the nature of the conversation"; "I wish I could remember it." Unfortunately, his testimony piqued Senator Levin's ire: "Finally, when David Bushnell was asked whether he agreed that it is the responsibility of a financial institution like Citigroup not to participate in a deception, believe it or not, Mr. Bushnell said, 'it depends upon what the definition of a deception is.' I guess that's what is meant by 'standing tall.'"³

But of all Weill's problems, it was an imbroglio with Grubman that struck closest to home. It ensnared Weill in a series of news articles containing increasingly embarrassing, almost preposterous revelations that had the effect of making the pair appear simultaneously malfasant and pathetic. The New York tabloids couldn't get enough of it.

Once again the starting point was one of Grubman's stock calls. At Weill's behest, Grubman moved his recommendation for AT&T from glum to rosy just in time to get Salomon Smith Barney in as one of the lead underwriters for the AT&T Wireless Services IPO, the largest in history. Citigroup did pocket \$40 million, and if you are going to pay an analyst \$10 million to \$20 million a year, it's just good business to pull out all the stops for the firm when the need presents itself. Yet by some accounts, the goal of Weill's request and of Grubman's rating turnaround wasn't as high-minded as making money for Citigroup's shareholders; it was base political ambition: ingratiating Weill to AT&T's chief executive and Citigroup board member Michael Armstrong to secure his support in the boardroom battle ensuing between Reed and Weill—"to nuke Reed in a showdown," as Grubman put it in one e-mail.

The intrigue took a comic turn when, as the backroom machinations hit the press, Grubman's version of damage control was to state that he had invented the whole Weill-Armstrong link "in an effort to inflate my professional importance and make an impression on a colleague." The ex-

planation didn't get far; one colleague, on hearing this variant of "my dog ate the homework," threw up his hands, exclaiming, "That guy needs his head examined."

KINDERGARTEN CONFIDENTIAL

A final incredulity was the revelation that the quid pro quo for Grubman helping with his boss's boardroom infighting was for Weill to take a million bucks from the coffers of the world's largest financial corporation to give Grubman's twin two-year-olds a leg up getting into nursery school.⁴ And not just any nursery school. Grubman had his sights set on the 92nd Street Y, a bastion of privilege in New York's Upper East Side, with a roster of billionaires on its board and a flock of their children (including a sprinkling of Bronfmans and Lauders) in its school. Based on the numbers of applicants and the hoops a prospective student has to jump through even to be accepted into the application pool, it is harder, as Grubman observed in a memo to Weill, to get into the 92nd Street Y's nursery school than it is to get into Harvard. And it surely must be a lot harder for a nursery school to pick out that special toddler from the pack than it is for an elite college to size up high schoolers. What can possibly differentiate one two-year-old from another?

To be one of the lucky few who gets to spring for the \$20,000-a-year tuition, parents hire private tutors to prep their toddlers for their admission "interviews." Consultants polish the parents as well: take an interest in the arts, head off for exotic travels with social import, be broad of interests and culturally aware. And of course, as anyone at the school will tell you, while "not at all considered in making the admission decision," underwriting a few test-run fund-raisers might not hurt. Nor, as Jack Grubman surmised, would making a well-directed contribution.

The competitive nursery school phenomenon is a New York singularity that is difficult for normal folk to fathom. Many of those parents fiercely competing for spots for their children are simultaneously pushing for diversity, affirmative action, and social equality on many fronts; it's just that the public school down the street is not one of them. We have had friends implore us to donate to school fund-raisers they were sponsoring so that they could demonstrate their worth as a valuable asset for the "parent body"—the first time I had heard that term.

These nursery schools are also a feeder for the next rung up the New York private school ladder, which also wrangles the affluent parent body. After all, these are precertified donors; they have already demonstrated their willingness to pony up, and so their nursery school graduates have an inside position competing for kindergarten spots at the elite private elementary schools (\$20,000+ annually), whose administrators know that, as Grubman put it, there are “no bounds for what you do for your children.”

The 92nd Street Y connection gained such notoriety it was echoed in a “ripped from the headlines” episode on *Law and Order* called “Kid Pro Quo” where the Grubmanesque parent was a porn king and the Weillist fixer was a mafioso who owned a cement company. As absurd as nursery school admission may seem as a motive for this sort of thing, it is not as far removed from the workaday business world—the Weill and Grubman sort of workaday world, anyway—as it might at first appear. Getting your kid into a school like the 92nd Street Y—or a handful of others in Manhattan—has little to do with education and everything to do with executing business by other means. The school itself is largely beside the point. No one seriously thinks that a three-hour playgroup can be transformed into a singular educational experience, no matter how much money and culturally aware parentage is thrown at it. The focus is on entering the exceptional parent body and the privilege of paying \$20,000 or more a year to connect to the high-society play date underground.

THE CONSEQUENCES OF COLOSSUS

The solution to these problems? Well, after Grubman’s electronic postings found their way into the newspapers, Weill posted a sign on his desk with the word *e-mail* slashed in red ink. It might not have been surprising that when it seemed things could not get any worse, they did. In 2004, Citigroup had to pay nearly \$3 billion to settle securities claims related to WorldCom, then pay a \$70 million fine for its abuses related to mortgage loans made to low-income/high-risk borrowers (the largest fine imposed at the time by the Federal Reserve for consumer lending violations), then suspend two executives in China who had presented false information to Chinese regulators, and then watch British regulators investigate a \$13.5 billion bond trade that implicated the company in market manipulation.

Then the most damage of all to the firm's reputation occurred: Japan closed Citigroup's private bank because of fraudulent transactions and regulatory violations.

Citigroup's Japanese private banking arm had been caught, the first time a non-Japanese company was called to task. Management flew over to apologize, in the Japanese custom. Three senior members of Weill's management team, including the great survivor, Deryck Maughan, were forced to resign. (Maughan later reemerged as the head of Asian operations for Kohlberg Kravis Roberts.)

All of this left some with the impression that Citigroup simply could not be properly managed. The new CEO, Chuck Prince—a lawyer no less—was utterly exasperated. “I never thought before that you had to say to people, ‘You’ve got to make your numbers, and, by the way, don’t forget not to violate the law.’”⁵ In response, he developed a road map to improve internal controls and employee training. The Federal Reserve's reaction was something like “It’s about time.” In March 2005 the Federal Reserve let Citigroup know that it expected the firm would “devote the necessary attention to implementing its plan fully and effectively and won’t undertake significant expansion during the implementation period.”⁶

The Citigroup that Weill created was not built to be managed. Weill's modus operandi had been, basically, bargain hunting. He would scour the market for merchandise selling at a bargain-basement price, snap it up, and restore its value through a disciplined focus on the bottom line. Expense control is fine initially but it is not the same as management, so the steps that led to the creation of Citigroup did not necessarily contribute to a smooth-running operation. Weill was ultimately a world-class deal maker, not a corporate manager. Sadly, this meant that, as one analyst put it, he built up “a reputation as a brilliant strategist and he’s going out as someone who couldn’t manage everything he has under one roof.”

Consolidation of the kind that created Citigroup is a natural result of the commoditization of financial information and markets. Weill was one of the first to see that. The informational ether—with access to real-time prices; market commentary flowing as a continuum from screens at home and office, airport terminals, and teller lines; stock recommendations spilling out of countless Web sites—reduces the market advantage of investment firms. With little left to differentiate themselves, their best of-fense comes from economies of scale and cross marketing. But if there is a

point at which an organization can get too large to manage effectively, Citigroup probably found its way there. And if there is an industry in which size can become cumbersome, it is probably the financial industry. With the speed of decision making and the rapid flow of information being key to success, any trader can make decisions that will put the firm on the hook, and the stakes are so high that there is a strong temptation to cut corners. With large organizations, it can become difficult to determine who is making the decisions, and momentum can take hold and move the process on with a life of its own.

THE NUMBERS ARE THE ISSUE

The organizational dysfunction that seems inevitably to creep in with size and complexity was joined with one other facet that made Citigroup's problems more difficult. Citigroup, and for that matter all of the financial profession, is saddled with a set of tools and measures that more often than not miss the mark of their intended objective.

While it became a popular pastime to poke fun at Jack Grubman, he was only a symptom of a deeper problem. He simply took advantage of the deeper failure of the financial world: information that is so obtuse and ill-designed that investors end up relying on others' expertise to unravel it and divine its implications. If the fundamental data of the company were laid before an inquisitive investor, would it have been so hard to tell that WorldCom or Enron was in trouble? Could someone like Grubman have overcome challenges to his statements of corporate health and his projections for WorldCom?

The problem is that the data are not available. They are locked up in the corporation, combed through and then summarized by a staff of accountants as financial statements that follow form rather than function, a form that relies on methods that are centuries old, using information that in earlier times was penned in a ledger by guild members and merchants, and later codified at a time when the business issue at hand was tangible assets like steel mills and rolling stock.

The objective of statistics is to take a mass of data and come up with aggregates of those data that extract all of the relevant information—what are termed in mathematics “sufficient statistics.” For example, meteorologists cannot track the path of every air molecule in the atmosphere. In-

stead, they monitor statistics that provide information about all of those collisions, statistics such as temperature, wind velocity, barometric pressure, and humidity. While they summarize trillions of events into just a few numbers, for most purposes these are sufficient to answer—as best as we can answer—the questions of weather. For traders, the raw data set is the trades and the bid-offer spreads with their associated volumes over the course of the day; the statistics most of us care about are the closing price, and perhaps also the open, high, and low prices, along with the total trading volume.

Accounting has failed us as investors and managers. Accounting should be about generating sufficient statistics to assess the state of a firm. It is more than simply taking numbers and putting them in the appropriate bin based on generally accepted accounting principles (GAAP). In an ideal world, the role of an accountant is to take the mass of all of the transactions of a company—from the purchase of materials to the sale of products, from money that is borrowed to wages that are paid—and boil them down to a set of statistics that provide a picture of the company's performance sufficient for the equity holders to make investment decisions. And the accounting statistics should not only have a high correlation with the company's value and prospects; they should also be standardized across firms for comparison purposes. Accountants have not been doing a very good job of this, and if anything they are marching backwards, doing worse and worse over time.

One reason for this failure is the well-known problem of conflict of interest. Accountants have a financial incentive to be on the company's good side so they can keep their mandate and garner additional consulting work. This conflict was the main reason for the erosion in the quality of financial reports over the course of the 1990s. While Arthur Andersen's failure was the most visible, with the restatements and liability arising from Enron, Waste Management, Sunbeam, the Baptist Foundation, and Global Crossing, the accounting firm may have not so much been an aberration as the most unlucky.

During the 1990s, the number of restatements, which is generally viewed as a factor that is correlated with the number of instances of fraud, rose rapidly. In the early part of the decade, the number of restatements for publicly held corporations averaged less than 50 a year. In 1997 the figure rose to nearly 100, then to 150 for 1999 and 2000, and held at a level

over 200 for the next three years.⁷ Even if you fix this, the problem still remains because accounting is mired in historical methods and standards that are increasingly less relevant to the way businesses operate.

PACIOLI RUNS THE NUMBERS

It is telling that the man who is still considered to be the “father of accounting” lived more than half a millennium ago. Luca Pacioli was born in 1445 to a poor family in the small commercial town of Sansepolcro in central Italy. He might have led an unremarkable life but for a love of mathematics and the friendship of a famous painter and master of perspective, Piero della Francesca, who had a studio and workshop in the town. As Piero’s protégé, Pacioli learned the principles of perspective and associated with the artist’s patrons. (Pacioli also posed for the painter, and is immortalized as the figure of St. Peter the Martyr in “Madonna of the Egg.”) Piero introduced Pacioli to Leone Battista Alberti, a writer and scholar as well as a noted architect, who arranged for him to pursue the life of a scholar in Venice. The pinnacle of Pacioli’s scholarly work was his treatise *Summa de Arithmetica, Geometrica, Proportioni et Proportionalita*, which laid out a comprehensive exposition of much of the collected knowledge of mathematics and the mathematical basis for proportion in drawing. The book became a standard, in part because Pacioli wrote in Italian rather than the less accessible Latin. It was translated into German, French, and English, and was one of the first books printed using the Gutenberg method.

Leonardo da Vinci admired this work and requested that his patron bring Pacioli to Milan to tutor him in mathematics and proportion. The two worked together for more than a decade. Pacioli’s teachings on perspective preceded Leonardo’s painting of the Last Supper. Leonardo provided illustrations for Pacioli’s work, *De Divina Proportione*, which applied the golden ratio, a topic that later became the object of study in his famous Vitruvian Man.

As a favor to one of his patrons, Pacioli included in the *Summa* a discussion on the Venetian method of bookkeeping. The Venetian method recorded every business transaction in two parts, a debit (an account that received funds) and a credit (an account that provided funds). If that sounds familiar, it is; this is the source of double-entry bookkeeping, which makes up the core premise of accounting 500 years later. The propagation

of the Venetian method through Pacioli's work was a key ingredient in the voyages of discovery of the sixteenth century. It provided a reliable system to record investments and returns, essential in attracting wealthy merchants to provide the backing required for these voyages, which opened the way for trade with the New World and the Far East.⁸

DA VINCI'S ACCOUNTANT IS STILL KEEPING OUR BOOKS

This approach has remained the foundation of accounting to the present, first blossoming in the United States with the building of railroads. The railroads depended on public financing and needed a means of assuring investors that their money was at home under a positive stewardship. They had accountants run through their books and then made their finances known through business journals. (Lots of railroads went bankrupt anyway, victims of overcapacity—the rails were the Internet stocks of their day.) Standardization of accounting methods also came about through the demands of railroad accounting. In 1887 the Interstate Commerce Commission took control of the rates that railroads could charge. A fair return was determined by looking at costs, revenues, and capital invested, all reported through standardized accounting. After the crash of 1929, the Securities and Exchange Commission (SEC) mandated that these standards be applied to all public corporations. This mandate persists today, but as often as not, the SEC is mandating the reporting of irrelevant information.

The orientation that accounting took for railroads was, not surprisingly, focused on the assets of the company—the track and rolling stock, along with the depreciation of capital—rather than earnings. Value was defined simply as the cost of the assets less the depreciation of the assets over time. This orientation carried through smoothly to other industries of that era, principally manufacturing and transportation, where cost could be used as an index of value.

For real assets such as physical plant, assembly lines, machinery, and real estate, valuation in terms of costs is logical: A business can reproduce the enterprise by simply going out and buying each of the component parts that constitute the production process. But the relationship between the cost of assets and the value of the enterprise does not work as well for companies with intangible assets, and these increasingly form the basis of

economic value today. Intangible assets—ideas, patents, proprietary software, brand names, trade secrets, trademarks, and copyrights—have values that cannot be extracted from their costs.

The reach of intangibles is extensive; as Charles Leadbeater has said, “modern corn is 80 percent science and 20 percent corn,” alluding to the extensive lab development behind hybrid corn seed. By some estimates, intangible assets now make up 80 percent of the value of the S&P 500. They are what provide companies with their franchise value, sometimes bordering on monopolistic market position. Intangible assets are the product of imaginative people who walk out the door every night; others are formulas locked in a vault. And in many cases, once they have been created and the intellectual property has been claimed, they cannot be reproduced at any price.

One simple indication that the current accounting conventions do not reflect the actual value of the enterprise is the disconnect that has appeared between market and book value. In the industrial era of the railroads, market value was all but defined by book value. If market value moved above book value, you would simply create the same enterprise for less money by replacing it brick by brick. The market-to-book ratio stayed near one-to-one through the 1970s, but since the 1980s has slowly moved up. The ratio in the mid-1990s was on average about three-to-one, and shot up to six-to-one by the end of the decade. The extreme is in part due to the euphoria of the Internet bubble, but the ratio has been out of its classical balance for the better part of two decades, more than can reasonably be ascribed to a market disequilibrium.

Like the diner who complains that the food is awful and the portions too small, accounting methods not only summarize business information into irrelevant measures, but are also excruciatingly slow at doing so. In a world of instantaneous and continuous data, the quarterly 10-Qs are little more than a historical record, and the annual financial statements are archaeological. How well a company did in the past, reconstructed based on the fragments made available by the accountants, is of limited value in meeting the key demands of the market: understanding future performance. The limited relevance of financial reports is evident in the fact that many well-paid analysts are employed to pick these reports apart and reconstruct them, like archaeologists piecing together the clues of a dig, trying to discern the real story. The raw material they receive to do their

analysis is slow to arrive and aggregated in a way that may have nothing to do with what they are after.

There is no reason to think this exercise of tearing apart the accountants' aggregation and then trying to reaggregate it into a meaningful form can be successful. And it certainly is not the ideal. The ideal is not to take something that is pieced together incorrectly and then redesign it; the ideal is to start with the raw materials, the actual transactions themselves, and build from there. For example, beyond the standard accounting numbers, statistics that might be helpful for companies with nontangible assets are the cost of acquiring new customers and the retention rate for those customers—think of the insight these would have provided into America Online (AOL) in its years of burgeoning growth—sales backlog, contracts received versus proposals made, training expenditure per employee, revenue from new products compared with revenue for old products, the proportion of business that is done with existing customers, and the time it takes for a new product to recover its development cost. Some of this information dribbles out: Construction and defense companies, for instance, typically list backlog.

The ideal is to move accountants out of the data aggregation business. With modern data markup languages and Internet tools, companies can make the raw transactions data available for anyone to analyze. In the spirit of open source software, analysts will aggregate the data into whatever form they find illuminating, and the market will then determine what statistical representations of the raw data are the most useful. The best statistics may vary from one type of business to another. The role of accountants will be to verify and expedite the process, to make sure that all the data are made accessible and that they are described correctly.

MY DEPARTURE FROM CITIGROUP

For Citigroup, by whatever measure you get to it, the proper accounting focus was on the market valuation of its business segments, and in market value terms the high-multiple retail business is king. Trading falls by the wayside because trading revenue has a low multiple, does not contribute to the franchise, and is often regarded by the market as not sustainable. (Indeed, part of the game for the large firms is to find ways to expropriate value away from individual talent. The multiple is never going to get very

high if the market's perception is that your assets not only might fail to deliver but also might walk out the door at any time.) Salomon Brothers' proprietary trading revenue amounted to nearly half the firm's revenue (and almost all of the earnings); once Travelers bought Salomon, it was part of a much bigger pie and represented only 10 percent of revenues.

After the Travelers/Citigroup merger, whatever Salomon once did was the equivalent of a rounding error. And as that unit's significance dropped with consolidation, the corporate political risk for being identified with trading revenue, especially propriety trading revenue, grew. It is off the radar screen most of the time, so any blowup seems to come out of nowhere. No one has time to understand trading or see it in context. Management usually just looks around and axes whoever is closest to the scene of the accident. To no avail, I argued that we should adopt the old Salomon model in which those running the risk were more integrated into the risk management process. I didn't get very far down this road before my time at Citigroup was over.

An often annoying but on rare occasion indispensable aspect of the Wall Street world is the headhunter. I would field calls from headhunters a few times a month. Most operated with coded messages that attached a certain dignity to the process, although why they cared is beyond me. They would call to ask for assistance in a search. Half of the time they really did need my help to find the right person, but the other half of the time they had me in mind and were hoping that as they went through the description I would say so. I tried to stay cordial with them because occasionally I had colleagues, present or former, who were trying to find a better position, and less frequently I was looking around myself. In early November 1998, the aftereffects of LTCM were wending their way through the political corpus of Citigroup. It was clear that over time Citigroup would become an ever more political beast with ever less appetite for what I did, so when I got a call from a recruiter about a risk management position at Moore Capital Management I took the bait.

This was fortuitous because Heidi Miller was gunning for me. Miller, who was the CFO of Travelers and my immediate superior, was miffed that her protégé Bebe Duke had been passed over in the merger and now was my deputy. Given that the vast majority of the risks of the combined firm rested with Salomon's trading desks, and given my experience in the role, it would have been strange if Maughan's and Dimon's decision had gone

COLOSSUS

the other way. That said, I was the only one from the Salomon side to have any oversight role in the merged firm.

There were times when I had the feeling Heidi and Bebe seemed to be spending more time figuring out ways to edge me out than getting their arms around the risks of the new enterprise. Sometimes what they did was just funny, like producing an organization chart that had me at the top, and then Bebe, and then everyone else under her. In other cases I was unaware of their machinations until someone told me after the fact that Maughan or Dimon or one of the trading desks had jostled with Miller over my operation. In late 1998, when Weill forced out Maughan and Dimon as co-CEOs, the coast was finally clear for Miller, and I was out of Citigroup a few weeks later. Shortly thereafter, I joined the exodus from what had become of the white-shoe sell-side world into the new frontier at one of the world's largest hedge funds, Moore Capital Management.

CHAPTER 8

COMPLEXITY, TIGHT COUPLING, AND NORMAL ACCIDENTS

Complexity is a byproduct of today's interrelated markets. It is not always benevolent; it is at times catastrophic and is always helped along by the organizational jumble of firms like Citigroup, as well as by the host of derivative instruments that have come to dominate the financial landscape. These derivatives can be customized to meet specific needs, often with unintended consequences. Examples abound: Bankers Trust marketed a low-cost derivative contract to hedge interest rate risk, but nobody seemed to realize that the low cost masked a feature that would turbocharge losses if rates suddenly shot up, as Procter & Gamble found out \$150 million later. Working in the bank's Asian hinterlands, a young trader for Barings engaged in derivative cross-trading, amassing losses of more than a billion dollars and leading the bank into collapse. When the foreign exchange desk at Salomon Brothers wrote put options

in a bet that the yen had peaked, no one noticed that the nonlinear feature of the put options magnified the loss nearly fiftyfold when the yen move 10 points higher rather than just one or two points. Granite Partners hedge fund earned a yield substantially above comparable corporate and agency bonds by specializing in esoteric mortgage derivatives, but neither Granite nor the investment banks that sold it these positions could figure out how they would behave in a changed interest rate environment.

In hindsight it is difficult to understand how anyone could have been blind to any of these risks. But the same can be said for the two shuttle disasters and for Chernobyl and Three Mile Island. Failures have causes, and after the fact these causes are often easy to explain. The problem is complexity itself. We cannot prepare for every thread of causality through every interaction; in the speed of the event we find there is no time to make adjustments.

THE TIES THAT BIND

The complexity at the heart of many recent market failures might have been surmountable if it were not combined with another characteristic that we have built into markets, one that is described by the engineering term *tight coupling*. Tight coupling means that components of a process are critically interdependent; they are linked with little room for error or time for recalibration or adjustment. A space shuttle launch sequence is a tightly coupled process because each step—the ignition, the liftoff, the clearing of the tower boom—all must proceed at precise intervals and cannot be interrupted without scrubbing the whole operation. A process that is tightly coupled can be prone to accidents even if it is not complex. Think, for example, of the Charlie Chaplin caricature assembly line in *This Mechanical Age*, where a mishap leads one pie after another to pile up in a heap. Each process of the assembly line is tightly coupled to the next, each action immediately triggers the next, and there seems to be no point to intervene successfully. When things go wrong, the error propagates, linked from start to finish with no emergency stop button to hit.

The tight coupling in financial markets comes from the nonstop information flow and unquenchable demand for instant liquidity. Information spurs trading, and the trades are entered and executed without a pause. Tight coupling is accentuated by leverage, itself a direct result of liquidity.

Leverage and margin are simply loans that use securities as collateral, and the willingness to lend against this collateral is directly related to the lender's ability to quickly sell out the securities if the margin is not posted. The more liquid the securities, the better the leverage terms will be. So tight coupling means higher potential leverage.

For financial markets, the tight coupling born of liquidity feeds right back to the source of complexity. Liquidity is the lifeblood of derivatives; unlike the underlying securities, derivatives are created on the assumption that they can be hedged on an ongoing basis, and so make continuous demands on liquidity. Without liquidity, derivative markets die.

The interplay of complexity and tight coupling that comes from combining liquidity with its derivative and leverage offspring is a formula for disaster. If all the eventualities cannot be anticipated (which is the case for complex systems) and if there is no time to rework the process before the problem is propagated down the line (which is the implication of tight coupling), then when things go wrong a crisis will be unavoidable. Things will go bad, and when they do, they will quickly move from bad to worse before the cascade can be stopped.

The Long-Term Capital Management debacle is a prime example. It started with a relatively minor loss for the hedge fund, a loss that required LTCM to liquidate positions to meet demands for margin. A simple exercise in raising cash through a security sale generated a downward cascade when the liquidated securities sucked prices down, causing the overall portfolio to lose value. The drop in the portfolio elicited further demands for cash to cover declining positions. These in turn precipitated further liquidations. The market reacted more to each wave of selling, and the next cycle of demands of the creditor banks then followed immediately. The crisis came to an end unnaturally, when the Federal Reserve strong-armed the creditor banks to break the cycle, effectively unlinking the tight coupling.

The cascade of failure that occurs in systems that share tight coupling and complexity gives rise to what are called normal accidents.¹ As ironic as the term sounds, normal accidents are accidents that are to be expected; they are an unavoidable result of the structure of the process. The more complex and tightly coupled the system, the greater the frequency of normal accidents. The only way to reduce the number of accidents is to reduce the complexity or add some slack to the process. This cannot be

done for the launch phase of a space shuttle or a nuclear power plant spinning toward criticality, although, Charlie Chaplin notwithstanding, it can be done for an assembly line. And it can be done in financial markets, but only by pulling back from the quest to create a real-world analog to the academic model of what the markets should be.

THE REGULATION TRAP

The natural reaction to market breakdown is to add layers of protection and regulation. But trying to regulate a market entangled by complexity can lead to unintended consequences, compounding crises rather than extinguishing them because the safeguards add even more complexity, which in turn feeds more failure. Trying harder means sinking deeper into the quicksand. Yet regulators and institutions can't stand idly by in the face of potential crisis; something must be done. They may not know the mechanics behind the crises and they certainly cannot know where or how the next crisis will arise, but they learn from their mistakes and develop new or revised regulations and safeguards. And in a world of increasing complexity, those safeguards add yet more to the complexity. And so on.

For example, the first line of defense for regulators to prevent banks from failing is placing limits on the amount of risk they can take relative to their capital base. The idea is to create a self-dampening reaction to crisis: If a bank loses capital because of a bad investment, or if a bank finds increased risk in its loans, it will have to offload risk by selling off assets or recalling loans. This will reduce exposure to further losses.

The problem is that selling assets or recalling loans doesn't occur in a vacuum. If the bank's actions take place in the region under pressure, it adds fuel to the fire. Its forced sales will drop prices further, reducing the bank's capital base, which in turn requires further liquidation. At the same time, the increased selling will raise the volatility of the market, implying greater riskiness for the bank's positions and signaling greater demand to liquidate. Trying to control the risk ends up creating the liquidity crisis.

If the banks elect to reduce exposure by selling off assets in other regions, the effect may be contagion, turning a regional crisis into a global one. Banks in other regions of the world will see their market plummet for no apparent reason. And, indeed, the only reason the crisis is transmitted is because the bank that was under pressure happened to also be holding

assets in that other region. This unintended and wholly uneconomic consequence is all the more likely to result because banks often have similar concentrations of assets and loans, so the actions of each bank will magnify the actions of others. The 1997 Asia crisis started in Thailand, then swept across the Asian Tiger economies because of the regulatory requirement that banks adjust their exposure. It then spread to Brazil for no real economic reason, but simply because most of these banks also happened to have large and liquid Brazilian exposure. Just where the global reach will extend is unpredictable; it all depends on which country is having a crisis, which banks are involved in that region, and which other assets those institutions happen to be holding in their portfolios at that moment.

In the idealized market, the starting assumption is that the market should run cleanly and transparently. We are faced with more pernicious problems, however, in attaining these goals. When the market ideals collide with the real world, with individuals who are not in control of full information, with institutions that do not act quickly or necessarily in anyone's best interest, the result is like taking a race car for a spin off-road.

In the face of progress and technological advances that have resulted in stability on many fronts, financial markets, designed to provide a mechanism for managing and addressing economic risk, have developed a structure that has made them inherently more risky. The irony is that this structure has features that at face value are desirable, in some cases approaching the essential elements of the ideal. As with many ideals, its origin is in academia, in this case a theoretical framework that underpins a half-century of work in financial economics called the perfect market paradigm.

To understand more concretely the structures and resulting interactions that can lead to unavoidable crisis, here are two case studies drawn from industries in which failures may be measured not just in dollars and cents but in loss of life: nuclear power and airlines.²

Three Mile Island: The Dangers of Safety Systems

The objective of a nuclear power plant is simple: to boil water to produce steam. The steam is then used to turn a turbine, a process that dates back at least two millennia to Hero of Alexandria. The complication with nuclear power is in how the water is boiled. The nuclear reactor contains

thousands of pellets of enriched uranium that are maintained in a slightly supercritical state; that is, on average a little more than one of the free neutrons from the fission of any uranium atom ends up hitting another uranium atom, leading it in turn to release a neutron. Heat is released in each reaction. There's not too much heat released per reaction, but there are a lot of atoms and a lot of reactions, so a pound of uranium, about the size of a baseball, can generate as much energy as a million gallons of gasoline.

The fact that the mass of uranium is "slightly" supercritical means that the heat from the reaction is containable and can be directed to generate the steam. But left to its own devices, the reaction will be self-sustaining and the heat will accelerate to the point that all the uranium and anything in contact with it will melt together. The melting event will suck in more and more of the uranium, and the more the uranium melts together, the greater the heat generated and the more radiation released as a byproduct. The end result is a white-hot ball of radioactive metal that will melt through its steel containment chamber and the concrete containment building surrounding it. The China Syndrome is a reference to a molten blob burning its way through the ground on its way to China. This burrowing blob is not the real problem, however; it is the huge emission of radiation that escapes containment along with the molten metal, releasing poisonous, airborne gases. Given the large amount of fissionable material in a nuclear reactor the size of the one at Three Mile Island, the radioactive fallout from a meltdown would dwarf the amount released in the atomic bomb dropped on Hiroshima.

To prevent any such disaster, the simple process of boiling water is accomplished using the most complex plumbing known to man. Systems and backup systems pump water through the reactor to capture the heat and whisk it away. Another system with yet other backup systems then takes some of this heat for boiling the steam and carries any excess to a cooling tower. Accompanying each system is an array of control valves; the function of each is monitored by arrays of status lights in the cavernous control center. The failure of one of these lights was at the core of the Three Mile Island accident.

The cooling water for the reactor at Three Mile Island Unit 2 was cleaned of impurities by a set of devices called polishers. Each of the polishers held a supply of resin beads that needed to be replaced every four

weeks. The cleaning, a routine maintenance process not even regarded as part of the safety system, had had periodic problems with leaking seals. At about 4 A.M. on March 28, 1979, some of the water from a leak entered the pneumatic system that drives some of the reactor instrumentation and interrupted the air pressure controlling two pumps that provided water to the reactor, forcing the pumps to come to a stop. With these pumps out of operation, emergency pumps were automatically engaged to run water from a storage tank through a secondary emergency cooling system to keep the reactor core temperature under control. Unfortunately, the pipes sending water from the storage tanks were closed off; someone had shut the valves several days earlier in the course of some other reactor maintenance.

This problem was not immediately evident. The controls indicated that the pumps were running, but did not indicate that water was not being delivered to the generators. The operators discovered the actuality after a few minutes, but by that time a sequence of events had already unfolded to create a malfunction that led to the disaster.

Human error and mechanical failure conspired unwittingly. Within the short period that the primary and backup cooling systems were out of commission, the pressure in the core built to the point that it triggered a relief valve. The relief valve is a little like the weighted valve on the top of a pressure cooker. If the pressure gets too high, it opens to let out some water and steam, and shuts again. In this case, however, the relief valve failed to close. This was anticipated; the valve was expected to fail about 2 percent of the time. For that reason, a light had been added to the control room panel to signal if the relief valve was open or closed. Incredibly, both the valve and the indicator light failed. (According to some accounts, the light was functioning but was obscured by a yellow maintenance tag). Thirty-two thousand gallons of water, one-third of the capacity of the core, went down the drain before the source of the problem was discovered.

If the workers in the control room were perplexed by the rising temperature in the core, they were probably even more baffled by the array of instrumentation they had to navigate in their unsuccessful attempt to uncover the problem. Nearly 750 alarm lights were spread out among the control room consoles, not necessarily close to the function they monitored. Sometimes they were above the relevant control, sometimes below or to the side, and sometimes across the room. The major warning lights

were interspersed with advisory lights, and even lights monitoring the building elevator. Along with the warning lights were hundreds of on/off indicator lights for the equipment. Most were simply identified by labels with code numbers or letters. Some were color coded, a system that at first seemed reasonable if not a little confusing: red for equipment that is running, valves that are open, and circuits that are closed; green for equipment that is stopped, valves that are closed, and circuits that are open. Besides the obvious problem of keeping all of this color coding straight, during normal operation some valves should be open and some should be closed, so without thoughtful study it would be impossible to glance at the lights and know what, if anything, might be wrong. A few minutes after the initial incident, the reactor coolant pumps started to shake violently. More lights started blinking. Three different alarms sounded, adding to the confusion. Despite the distraction of the alarms, the operators did not dare turn them off, because doing so would also shut down some of the warning lights.

The light for the relief valve was checked several times; so were a host of other controls. If the light had been working, the workers could have immediately activated a block valve to stop the outflow of the water, thus averting the disaster. If the light had a testing switch to verify it was operational, they could have at least been alerted to the fact that the warning light was not operating, and have used an alternative method to check the problem.

The operators resorted to what is termed high-pressure injection to force cool water into the core. This raised issues of its own, as the sudden injection of cold water into the superheated core could crack the structure. But the results of this operation were impossible to determine up front; one dial showed the pressure inside the vessel was dropping, while the other showed that it was rising to dangerous levels. One dial measured pressure in the system and the other in the reactor; they should have been moving in tandem as they always had in the past because the reactor and the system were connected. One of the dials might be wrong, but if so, which one? The confusion emanated from the original source of the problem: the stuck relief valve. The rise in pressure was occurring because the core was becoming uncovered as water was released through the stuck valve.

Neither the control room nor the set of instrumentation within the reactor had been designed with emergency operation in mind. The most

critical question in the hours after the accident was the temperature within the reactor and the path of the reaction that was occurring within the core. There was, by luck, a set of thermocouples in the reactor. They were not there by design, but had been put there temporarily as part of an experimental analysis of core performance. The problem was that they had been set to monitor the core during its normal operation. If the temperature rose above 700 degrees or fell below the normal 600-degree operating range, or if the thermocouple was not functioning properly, the readout was simply a set of question marks: “???????” As the operators struggled to determine whether the injection had worked, the governor of Pennsylvania was contemplating the evacuation of the city of Harrisburg, situated near the plant.

The failure of Three Mile Island was officially registered as an “incredible” event. But it was a normal accident. What occurred was one of an almost countless number of sequential failures. Any one sequence of these could lead to a catastrophic failure. The probability of any one event is small enough to be dismissed, but taken together, with so many possible permutations and with combinations beyond comprehension, the odds of one or the other happening are high. Nuclear power plants are the archetype of normal accidents.

ValuJet: The Human Factor

In the early afternoon of May 11, 1996, a ramp agent at the Miami airport loaded cargo into the bins of ValuJet flight 592. Passenger bags were loaded in bin 2 first, and once it was filled, the remaining bags were put in bin 1, along with 60 pounds of U.S. mail and some company-owned material that was headed back to ValuJet’s parts and components department at its Atlanta headquarters. This material included three airplane tires—two inflated—and five boxes marked “Oxy Canisters—‘Empty.’” The ramp agent loaded one of the two large tires first, laying it flat on the floor with the small tire wedged upright inside it. The boxes, each weighing about 50 pounds, were then positioned around the smaller tire to keep it upright. Then the third tire was put into the bin upright against the compartment wall and leaning over the other tires and the five boxes. As the agent was stacking one of the boxes, he felt its contents move and heard a clinking sound.

A DEMON OF OUR OWN DESIGN

Flight 592 to Atlanta pushed back from the gate at 1:40 P.M. and was cleared for takeoff at 2:03 P.M. Seven minutes into the flight, the captain heard a sound like a chirp along with a simultaneous beep on the public-address system. The flight-data recorder later showed that this coincided with a sudden pulse of high pressure—it was likely one of the tires exploding in the hold. Almost immediately, the instruments indicated an electrical failure. The captain radioed to the controller, “We got some electrical problem,” followed by a rapid succession of increasingly desperate transmissions: Five seconds later she radioed that “We’re losing everything,” two seconds later “We need, we need to go back to Miami”; then three seconds later came background shouts of “Fire, fire, fire, fire,” followed by an unidentified male voice shouting, “We’re on fire, we’re on fire.” The jet turned around to head back to the Miami airport, but in just over a minute disappeared from radar. Two witnesses fishing in the Everglades described the twin-engine DC-9 descending in a steep right bank, the bank increasing until the nose dropped to near vertical before impact. They rushed to the accident site, but all that remained were engine parts, scattered papers, and other debris. The Everglades had swallowed the rest of the plane whole.

The five boxes contained oxygen generators that were all beyond their expiration dates. They had been removed from three used MD-80 jets ValuJet had recently purchased and was refurbishing. Each generator supplied emergency oxygen to two or three of the masks above each seat of the aircraft through a chemical reaction; sodium chlorate is converted into sodium chloride, releasing oxygen as a byproduct. The oxygen passes through a number of filters and then to the masks. The process is initiated when a retaining pin is pulled, allowing a spring-loaded hammer to strike a percussion cap that has a small explosive charge that ignites the reaction. Tugging firmly on the mask when it appears over the seat, the instruction familiar to all air passengers from the preflight safety review, pulls the retaining pin.

The chemical reaction that produces the oxygen is exothermic, meaning that it liberates heat. When the generators are properly shielded and ventilated to dissipate their heat, the temperature still rises beyond 500 degrees Fahrenheit. If they are placed in a confined space without this ventilation, the temperature can exceed 1,000 degrees. Investigators determined that a pin fell out of at least one of the canisters, igniting a smoldering fire

that was fed by the abundant flammable material—the cardboard boxes, mail, and, most importantly, the tires. The pure oxygen generated by the canister then turned the fire into a blowtorch.

The inherent danger of the oxygen canisters is well known and a host of safeguards and checks protect against fire. The MD-80 maintenance manual provides a six-step procedure for removing expired canisters. The work card that requires sign-off when maintenance is performed on the canisters delineates a seven-step process. Step 2 in both procedures states, “If generator has not been expended, install safety cap over firing pin.” Both the manual and the work card also contain warnings about the hazards of the generators. And the newer canisters include a label stating in all upper-case letters, “Warning. This unit gets hot! When removing unit install safety cap over primer. Do not pull lanyard. If activated place on surface that won’t burn.”

The work card for the generators was signed off by a mechanic who had just come onto his shift. He removed 10 generators; another 30 or so had been removed in a previous shift, so he signed for those as well. He then looked around for the safety caps, but was informed by his supervisor that there weren’t any. Both the mechanic and the supervisor were focused on the airworthiness of the planes, not on the disposal of old equipment, so the mechanic placed the generators on their side in the cardboard box, piling one on top of the other, and got back to the other work at hand. The work card required three other signatures, ending with the maintenance supervisor, who signed the “Final Inspection” line, and, aware that the generators needed safety caps, alerted the lead mechanic on the floor.

Before anyone applied the safety caps, the boxes made their way to Shipping and Receiving, where they sat unlabeled. On May 8 a stock clerk with some time on his hands asked the director of logistics, “How about if I close up these boxes and prepare them for shipment to Atlanta?” “Okay, that sounds good to me,” the director replied. The stock clerk repacked the generators, laying them side to side with some bubble wrap on the top of each box, and taped the boxes shut. He then labeled the boxes “aircraft parts” and the next morning sent them to the receiving clerk. The stock clerk told him to write, “Oxygen Canisters—Empty” on each box for parts identification and attach a label with the ValuJet Atlanta headquarters address. The receiving clerk put the word “empty” in quotes, probably because

he suspected that a 50-pound box held something more than empty canisters. They sat in the shipping area two more days because the driver was busy, but on May 11 he loaded the items on his truck and drove them to the ValuJet ramp area, where he was told to put them on a baggage cart. With the shipping ticket signed by a ValuJet employee, the driver left.

Though potentially hazardous, the oxygen generators are a critical part of mandated airplane safety equipment, and regulation also mandates their occasional replacement. The hazards were clearly identified in both the instruction manuals and the work cards. The work card was signed off in all four blocks. In the process of their removal and preparation for disposal, the oxygen canisters passed under the inspection of at least two floor mechanics, a senior mechanic, a supervisor, an inspector, a program manager, several technical representatives, and the director of logistics. With the possible exception of the last of these, all were aware of the risks posed by the oxygen generators, and all understood the need for the safety caps.

Flying an airplane is by nature fraught with tightly coupled processes. You can't pause the flight in midair to do some reengineering if something goes wrong. The tightly coupled process that caused the ValuJet accident was sparked by the physics of the oxygen generators and the fire that flashed through the airplane. The complexity of the process came from the human factor, layers of checks that became too burdensome to execute. This illustrates a critical aspect of normal accidents: Normal accidents are borne of complexity, so adding safety checks to try to overcome these accidents can be counterproductive, because they add to this complexity.

ACCIDENTS WAITING TO HAPPEN: INTERACTIVE COMPLEXITY AND TIGHT COUPLING

Interactive complexity is a measure of the way the components of a system connect and relate. An interactively complex system is one whose components can interact in unexpected or varied ways, where there is feedback that can lead the components to differ in their state or their relationship with the rest of the system from one moment to the next, where the possible stages and interactions are not clearly apparent or cannot be readily anticipated.

Does it sound chaotic? Does it sound like the way a nuclear power plant should behave? Surprise—a nuclear power plant is a textbook example of an interactively complex system. Why? Because its operation is not easily viewed, many important components are monitored only by instrumentation, and even then the instrumentation is inferring a function indirectly by measuring side effects and by-products—for instance, heat. Parts of the system are shrouded in container units, pipes, and valves. The process has nonlinear interactions and feedback loops, the best-known and most serious being the chain reaction as fissionable material reaches criticality. And the physics of the system remains less than fully understood—despite years of development and refinement of the nuclear power plants—especially regarding the events and optimal palliative measures near critical levels.³

Systems with high levels of interactive complexity are subject to failures that seem to come out of nowhere or that appear unfathomably improbable. With many moving parts and points of interaction, and with interactions that can bound out of control in any direction, the unavoidable missteps that in a normal system would be addressed in the course of each day's work may take a freakish leap into disaster. The start of the crisis can be something as trivial as a tag covering a warning light in the case of Three Mile Island or a failure to properly label an oxygen generator in the case of ValuJet. And these two cases illustrate the shortcomings in attempting to deal with interactive complexity by adding backups and safety features; in the cases of both Three Mile Island and ValuJet Flight 592, the safety features themselves contributed to the disaster.

Interactive complexity does not automatically arise from a large-scale system or from a system with many components. The postal service is an example of a very large system with many nodes and subunits, but it is not interactively complex. Letter delivery is subject to many small failures, yet rarely if ever do these culminate in a catastrophic result. The mail gets through. Short of the military, there are few organizations with as many moving parts, in terms of employees, branches, and items to be processed. The reason occasional failure does not translate into crisis for the post office is that it is a linear process. The path for any one letter is, from a processing standpoint, a straight line. Letters move from one node to the next as they make their way from point of pickup to delivery. If one letter is misdirected, it affects just that letter; the error is not propagated—it does

not feed back into the system to cause yet more letters to run off course. At worst, a letter is lost. Furthermore, the individual failures can be readily seen and corrected. Letters pass from bin to bin, from handler to handler, where the misrouting or errors have a chance to be seen and redressed.

Given the daily volume of transactions, financial markets mimic the post office in many ways, but with a critical difference: Rather than the volume, or even simply the diversity of financial instruments, it is the structure of some kinds of instruments that is the source of complexity. Options and other derivatives have payoffs that are not linearly related to the prices of their underlying securities. Observing their day-to-day movements gives no inkling of what may be in store if the market moves dramatically.

For swaps and other instruments that facilitate highly leveraged positions, the complexity arises because that leverage can link the market unexpectedly to events that are distant and economically unrelated. A market can spiral out of control simply because there is some group of overextended investors who happen to have positions that for one reason or another they are forced liquidate. These interrelationships cannot be anticipated in advance and will shift with the fortunes and market interests of the investors and speculators.

The implications of complexity in financial markets are more pointed than for most other industries: In the financial markets, some participants have a self-interest in gaming the system. Traders do not act as uninvolved parties. They are ready to take advantage of increased complexity in the products and the organizations to serve their own bottom line, making it all the more likely that the unanticipated crisis will appear. If rules strictly limit interest rate exposure but allow for interest rate neutral yield curve trades, a trader can still make an outright interest rate bet on Fed policy while appearing to be rate neutral. The end run is constructed by doing a yield curve trade of the one-year against the long bond. The system will see this as having no net interest rate exposure, because it will see the 30-year bond as having interest rate exposure that counteracts that of the one-year. But in fact the 30-year will have little reaction to a Fed policy shift while the one-year will. For the more nefarious trader, the complexity of the system provides plenty of cover. So just as the market is made more complex because of the feedback of the traders and other market participants, the complexity of the trading firm increases because of the actions of traders who take advantage of it to shield their actions. In short, complexity helps the malfessant.

Interactive complexity is going to produce unexpected, even bizarre incidents, although not necessarily turn those incidents into disasters. A system might be very complex but still allow time and opportunity to cope with problems without a cascade into an ever-deepening crisis. Complexity can be managed if there is time to observe and investigate and to take steps to intervene in a process before it runs out of control.

TIGHT COUPLING AND INTERACTIVE COMPLEXITY: AN X-RATED BEHAVIOR

The greatest dangers arise when there is both interactive complexity and a tightly coupled system that does not provide the time to intervene. Tightly coupled systems are not limited to the sphere of rocket launches and industrial processes. A task as simple as making bread is tightly coupled; as the ingredients are mixed and the yeast is added, the timing and steps must follow in a fairly precise and controlled manner.

Whereas a loosely coupled system provides time to improvise and come up with solutions, a tightly coupled system must be run and managed by the book. Simply put, a tightly coupled system provides no slack, in terms of either the time between steps, the ability to make on-the-fly alterations, or the opportunity to intervene. The by-the-book mode of operation that is required with tightly coupled systems is precisely the way to get into trouble with interactively complex systems. The book can never cover all possible contingencies. Interactively complex systems require a decentralized approach that provides for creativity at the operator level in dealing with unanticipated failures. If a system is both interactively complex and tightly coupled, management is faced with a dilemma; neither centralized, codified management nor decentralized, adaptive management will work.

A university is an example of an organization that is interactively complex but loosely coupled. It has many departments, each with a curriculum and a faculty that are only loosely controlled by the central administration. Students straddle across these departments, devising their own course schedules and activities. To graduate, a student navigates courses in the various departments based on a set of requirements dictated by the university. Things go wrong: Class schedules conflict, lectures are canceled, teachers come and go, and students fail to make it to class or attend the exams.

In spite of all this, students regularly make it through; few, if any, find a missed class unleashing a chain of events that derails their college careers. The inevitable problems do not translate into crises because the university system, while undeniably complex, is not tightly coupled. There is plenty of time for a review of the requirements, for appeals for cases where a prerequisite is not offered, and for consideration of alternative classes or even new majors. There are makeshift approaches for dealing with missed classes or exams. Put another way, there is plenty of slack in the system (not to mention many slackers).

Another example of a system that is complex but sheltered from catastrophic failure thanks to loose coupling is the hub-and-spoke system for airlines. This is a system that in theory adds efficiency, but in practice has become a questionable approach because of the feedback that propagates unavoidable failures. (I am referring to airline scheduling, not airline safety.) But there are plenty of options available and time to consider them—though that often leaves you cooling your heels in Atlanta. Flights can be delayed or canceled; in the worst case, aircraft en route can be rerouted. It may seem preposterous that a system has been developed in which a storm near Chicago can ground a flight going from a cloudless Dallas to a balmy San Diego, but at least it occurs without further incident.

These examples hint at how, just as loose coupling can diminish the effects of complexity, the effects of tight coupling can be diminished by simple, linear systems. If a system is linear in the sense that actions move in just one direction without any feedback, magnification, or potential propagation of effects, failures can be addressed without blossoming into crises. The simplest example of this sort of tightly coupled but linear system is an assembly line. The steps in an assembly line, be it for automobile manufacturing or soft drink bottling, are designed to be tightly coupled. The completion of each action leads immediately to the initiation of the next step in the process. Efficiency dictates that the tighter the linkage the better. Taking everything in at once, the operation of an assembly line can seem chaotic, but in fact each operation stands in isolation, and generally involves one simple, repetitive part of the process. The Charlie Chaplin *Mechanical Age* example notwithstanding, a failure at one point in the process ordinarily leads to a temporary line halt. This may be costly in terms of lost production time, but the failure does not propagate. In the worst case, the failure isn't immediately detected and some of the produc-

tion is flawed and must be discarded or recalled, or a machine failure leads to a backlog that must be cleared away.

The natural reaction to accidents and failures is to add safety measures. These might include systems of safety checks and regulations, such as the preflight checklists for aircraft or the sign-off procedures banks use for loans, and they might include redundant or fail-safe systems that provide an emergency backup in the event of a primary system failure.

The problem with adding safety features to systems that are complex and tightly coupled—the very systems that by nature have the greatest risk of catastrophic failure—is that they can actually increase the likelihood of these failures because they contribute to the source of the problem: interactive complexity. By adding that many more wires, switches, meters, and items for human oversight, safety systems make the operation more opaque. The wires, switches, and sensors also interact with the system; for every step in improving safety or monitoring potential failures, they introduce their own sources of failure. If the safety mechanisms are automatic—as they almost have to be in a tightly coupled system—they serve as one more variable that can add to an unexpected result and a nonlinearity in effect.

NORMAL ACCIDENTS AND ORGANIZATIONS

The core platform for normal accidents is engineering systems, but the issues also exist in organizations. By their very nature, many of the processes in financial organizations, from investment banking deals to trading desk operations, are tightly coupled. So the only ingredient that needs to be added to generate normal accidents is complexity, and as we saw in Chapter 7, the Citigroups of the world can provide that in spades.

If Three Mile Island and ValuJet are poster children for normal accidents in engineering, NASA presents the case study for how to set up organizations for normal accidents.

NASA knew there were erosion problems on the O-rings that sealed the solid rocket boosters for the space shuttle, yet continued to fly with these problems unabated until the *Challenger* disaster. NASA knew there were problems with foam debris striking the spacecraft in the flights before the *Columbia* flight, yet continued to fly until the debris damage led to catastrophic results. In the case of the O-rings, cold weather led to discussions to assess the risk the evening before the *Challenger* launch. (The

low temperature degraded the O-rings.) Initially, engineers argued to scrub the launch until the weather warmed, but then, after they were exhorted to “put on their management hats,” they demurred and assented to the launch going ahead as scheduled. Seven astronauts died in the subsequent explosion. In the case of the foam debris, evidence of much larger than normal debris led to postlaunch discussion in which the views of the engineers were muzzled. Indeed, even their request for a Department of Defense imaging of the shuttle to look for possible wing damage was countervailed. *Columbia* broke up on reentry, killing another crew of seven.

The root problem with both of the shuttle disasters was not operational or engineering related. While the shuttle launch obviously is a tightly coupled process, the analysis of the O-ring and debris risk did not have to occur in real time during the launch sequence. There was time to evaluate the effects before the launches, and in the case of *Columbia*, the possibility, however remote, of remedial action in the days after the launch. In addition, the O-ring and foam debris were not complex or nonlinear processes. They did not occur in surprising and unanticipated ways. The weather was just colder than it had been previously, leading the O-ring to erode a little more than it had before. The chunk of foam struck the shuttle as had happened many times before; it was just a larger chunk this time around. And these events did not suddenly lead to a cascade of events that spun out of control. From an engineering perspective, their impact on the fate of the shuttle was a straightforward, linear case of cause and effect.

So the shuttle disasters did not have the markings of a normal accident—one born of the connected technologies—at least as far as the O-rings and the foam debris were concerned. Rather, the complexity and the point of failure occurred in the organizational realm, which was multifarious, with specialized groups from both NASA and contractors across the country. And of course the process of a shuttle launch is also tightly coupled, meaning the go/no-go decision process left little time for reworking the decision points of the organizational hierarchy. As was stated by the *Columbia* Accident Investigation Board, “The organizational structure and hierarchy blocked effective communication of technical problems. Signals were overlooked, people were silenced, and useful information and dissenting views on technical issues did not surface at higher levels.”

For example, the shuttle program ranked the safety-of-flight concerns, and the Flight Readiness Review had already made an assessment that the O-ring (in 1986) and the foam debris (in 2003) were within safe bounds. The belief among managers that the O-ring and foam debris problems constituted low risks was reinforced by years of successful flights. For the engineers who argued against the launch of the ill-fated missions, the safety rankings of the Flight Readiness Review became more than just being a frame of reference for discussion; they formed a barrier of casebook precedent that was difficult to scale. The engineers were put in the position of having to demonstrate in short order that these problems were now unsafe. With each successful launch, the risks, which were clearly outside of the bounds of the initial shuttle expectations (the design did not envision ring erosion or foam debris damaging tiles during launch), had become further codified as acceptable; they were innocent until proven guilty.

YOU CAN'T PLAY IT SAFE

Layer one safety system on top of another and you will finally doze off into a world of unjustified complacency. You never end up as safe as you think because there are inevitably points of interaction that can hide failures. And the more layers there are, the more obscured those points become. When two redundant devices are added to a critical component, a 1 percent failure rate will drop to a rate of one in a million if the risk of failure for each of the three devices is independent. But this will never be the case, because they will employ similar design and engineering, use the same components, and probably tie into the same power or instrumentation systems. Worse still, they may end up interfering with one another's operation.

For example, at Detroit Edison's Fermi-1 experimental breeder reactor a zirconium plate was installed to protect steel cones in the reactor. This was a safety precaution the plant supervisor felt was a "super-cautious" step. It turned out that the super-cautious step almost resulted in a disastrous meltdown because the plate loosened over time, blocking the flow of coolant. And since this step was taken after the design as a last-minute safety measure, the plate was not recorded on the plant specifications or blueprints, making it all the more difficult to uncover the source

of the coolant blockage. Fortunately, the reactor was safely shut down after only a small fraction of the fuel melted.

Detroit stayed on the map, but things were not so lucky for another safety-related disaster. Like all nuclear plants, Chernobyl, located just 80 miles north of Kiev, had emergency backup generators in the event of a failure of the primary power source and a further safety precaution had another redundant power source that pulled power from the slowing power plant turbines to provide short-term juice until the diesel generators could kick in. Reactor Four was to be shut down for routine maintenance, an ideal time for operators to run a test on this new power backup. The test was to check whether, in the event of a shutdown, this scheme could generate enough electricity to operate the emergency equipment and core cooling pumps until the diesel power supply came online. To test this new backup system, the plant operators disconnected a number of the shutdown devices so that they could run on the backup source without interruption. As part of the test, the emergency core cooling system was switched off and the reactor carried on at half power. But during the test the reactor core started to heat up. With the automatic processes deliberately disabled, the operators attempted to intervene. Because they lacked either time or training, they were unable to effect a shutdown. The reactor started to become unstable.

The reason for the instability was rooted in the reactor design, which created a complex interaction among steam production, fuel temperature, and the generation of neutrons. With either an increase in power or a reduction in the flow of water, steam production increased. More steam increases the fission in the fuel, because steam, being less dense than water, does not absorb neutrons as well as water. But there is a second, countervailing effect: An increase in power also increases temperature, and high temperature reduces the production of neutrons. The net effect of the steam, on the one hand, and the temperature, on the other, varies with the power level. At full power the high temperature reduces the neutron production by more than the steam increases it. But at low power output the trade-off goes the other way: The temperature is not reducing the neutron output enough to counteract the lower absorption of the steam. Neutron output grows.

During the power test the power fell to the point where the lack of absorption from the steam dominated. The hapless operators were then in a

balancing act of trying to turn the power up to move back to the point where the temperature was again outweighing the effect of the steam. They tried to raise the power by switching off the automatic regulators and freeing all the control rods manually. Meanwhile the pumps that were powered by the slowing turbines were providing less cooling water to the reactor. This exacerbated the unstable condition of the reactor by increasing steam production even further. The increase in heat compromised the control rods, and hot fuel particles reacted with the water to cause a steam explosion, which destroyed the reactor core. Two minutes later a second explosion blew up the container unit and precipitated a runaway nuclear reaction.

Thirty people died in the accident or within a few months; over the next few years, according to the understandably conservative Ukraine Radiological Institute, more than 2,500 deaths were linked to the accident. Western sources put that number—mostly cancer-related—at well over 100,000.

What lessons do these disasters hold for the financial markets? In hindsight, it's difficult to understand why these accidents were allowed to spiral out of control. The root problem seems to be a combination of organizational complexity, particularly in regard to risk controls; an unanticipated event; and human error. The result is a cascading and ever-worsening series of events that overwhelm the organizational systems and processes.

One lesson can be inferred from the role regulation played in the Asian Crisis in 1997, described earlier in this chapter. A first line of defense for regulators against bank failure is to limit the risk the banks can take relative to their capital bases. The idea is to create a self-dampening reaction to crises: If a bank loses capital because of a bad investment or if a bank finds increased risk in its loans, it will have to reduce risk by selling off assets or calling back loans. Result: reduced exposure to additional losses. The problem is that selling assets or calling back loans does not occur in a vacuum. If the bank's actions take place in a market under pressure, they add fuel to the fire because the forced sales will drop prices further. This will reduce the capital bases of all of the banks, which in turn will require yet further liquidation by them. At the same time, the increased selling will raise the volatility of the market, implying greater risk for the bank's existing positions and signaling even greater demand to liquidate. The result is a liquidity crisis cycle that emerges in the attempt to control risk.

A DEMON OF OUR OWN DESIGN

The point is simply this: Risk controls, putting on layers of regulation and organizational oversight, cannot always fix the problems that arise from the complexity and tight coupling we have designed into the markets. Indeed, it might just make matters worse. This is not to say we should throw all regulation out the window. But a better approach for regulation is to reduce the complexity in the first place, rather than try to control it after the fact.

CHAPTER 9

THE BRAVE NEW WORLD OF HEDGE FUNDS

When I moved to Moore Capital, it was one of a handful of multi-billion-dollar hedge funds. It owed its success and preeminent position in the investment world to the trading talents of its founder, Louis Bacon, who started the firm in 1989 after a short stint at Lehman Brothers. He had benefited from the tutelage of Julian Robertson, another luminary, whose Tiger Management fund amassed \$20 billion of assets at its peak. Others in their ranks included Paul Tudor Jones, the founder of Tudor Investments, who, like Bacon and Robertson, has Southern roots, and George Soros, a Hungarian Jewish émigré. The list of high-powered, multibillion-dollar hedge funds expanded in the 1990s with a new generation that relied on computer power and analytical models, such as Long-Term Capital Management, D.E. Shaw, and Jim Simon's Renaissance Technologies, and has continued to balloon to this day.

It would seem that any discussion of hedge funds should include a taxonomy describing all the types of strategies and instruments, putting everything into a neat set of boxes. I believe that doing so is not particularly informative, for reasons that I will spell out in Chapter 11. But as background, it is worthwhile to take a look at basic types of hedge funds and the strategies they pursue to understand what sorts of things they do, how they try to make money, and how they sometimes run into trouble.

FUN WITH DATA

You need an edge. That's the basic rule for information-based trading, whether in stocks, bonds, or foreign exchange. You either know things others don't or interpret what others do know in a better way. In today's information-centric world, knowing things others don't know is difficult, so the advantage increasingly derives from how common sources of information are pieced together and analyzed.

The classic global hedge fund is top-down. You develop a long-term, fundamental view of the global economy, identifying the economies that are weak or strong, recovering or slumping. The high-level, global view provides the framework for the remaining analysis. The next ingredient is to establish a political view. For example, you really can't look at the yen without both a long-term view and a sense of the objectives and strategies of the Bank of Japan and the Ministry of Finance, or without a sense of how pending elections might change the government's approach to the market and the economy. Obviously, political insight is helpful in developing a political view, so having a friend in government doesn't hurt.

Then you've got to start crunching—analyzing the fundamental data. And not just number crunching, but context crunching: You've got to determine if newly released economic data are consistent with an established view or if they imply that you must change your hypothesis. It's also important to weigh each new release against other information and economic trends. Although economic data often have an immediate impact when announced, the impact may be short-lived, partly because the information has not been absorbed by the market into its collective long-term political and economic outlook.

You can't do much to protect against the immediate market impact from new information, because it comes as a surprise and everyone gets it

at the same time. Even if you can gain an edge by being quicker to react, the amount of capital that can be put into play is small. But frequently the most interesting aspect of new information is not what it is but how the market reacts to it. For example, is the follow-through on the news reaction shorter than usual? Suppose the market sold off prior to the data release, but then, when bad information actually hit, the price did not drop further. To many traders, that will be of more interest than if the price did drop. So after data are disseminated you've got to get into the weeds and develop a sense of market sentiment and market technicals, gaining an insight into underlying trends and potential turning points.

Finally, there's money, and, more specifically, where it's going. The flow of money into the investment markets is critically important, because the flow can give a sense of what types of investors are long and short, and how many might be overstretched in terms of leverage. Money flowing in or out of mutual funds and pension funds can move markets because of the sheer volume involved and the funds' need for liquidity. Large funds lumber along without much regard for the dynamics of the market because they are too big to take advantage of short-term market swings and because they are constrained by the demands of their investors or the liabilities they must fund. Furthermore, because they have large holdings, changes in positions take a long time to work through the market.

Related to money flows, and far more predictable, are seasonal effects. Examples include the tax-related January effect in the U.S. market and the tax-related effects in March and September in Japan, window dressing by investment managers at quarter end, and higher volatility during earnings season. These effects are predictable, because like the money flows of pension funds and mutual funds, there are investors who must act despite the implications their actions may have on prices. Seasonal effects may not imply clear trading opportunities, but can define when to add or reduce positions.

The top-down approach served well for a long period of time. But top-down met its match during the sky's-the-limit Internet bubble, for all the wrong reasons. What is fascinating is that the elite investors did not miss the point. They saw the Internet bubble for what it was, an irrational mania. But resisting its momentum proved fatal. That such experienced and luminary figures could fail in spite of understanding the market's dynamics is a lesson in how treacherous the hedge fund business can be.

BUBBLE BATHS

There are any number of stories that illustrate the extremes of the Internet frenzy, but one of my favorites is StockGeneration.com, a Web site that allowed its users to buy virtual stocks—for real money. On this “exchange,” set up as a gaming site in the laissez-faire Commonwealth of Dominica, the user could purchase any of 11 virtual stocks. The exchange set the prices of the shares and even provided a schedule of future appreciation for the stocks. For example, one of the stocks was guaranteed to appreciate by 10 percent a month. (As is stated on the Web site, “A guaranteed return of 10 percent per month without any risk involved isn’t so bad, is it?”) Others were more erratic, but with trends that were out of this world. One stock increased from 35 cents to \$14,000 in the space of two years. The list of winners who traded in this virtual market included some who saw their fortunes grow from a few thousand dollars to \$20 million—at least in their StockGeneration.com account. For its part, the exchange extracted 1.5 percent (again, in real money) of all buying and selling of the virtual stocks. The exchange was an unabashed pyramid scheme, where any attempt to cash out could be financed only by the inflow from other players. Without any real product or source of revenue, the total assets of this enterprise could not be higher than the total dollars received from the “investors,” yet all investors saw their fortunes rise. The site addressed this point more or less head-on: “As long as the mountain keeps on growing . . . there is no reason for anxiety. . . . As you can see we have succeeded in involving just 3.5 million people. . . . There are already more than 100 million people on the Internet, and this number almost doubles each year. See it for yourself: no ending in sight!”

In an age in which people are willing to invest money in virtual stocks, where by definition there are no prospects of earnings and where price appreciation is obtained through nothing short of an unsustainable bubble, it is not too hard to see how a real dot-com, with real prospects, no matter how dim, could attract investors.

Market bubbles have been explained by the tendency of investors to follow trends and by the dynamics of crowd psychology—the need for people to be part of a successful herd. But neither trend-following strategies nor irrational crowd behavior is necessary to create market bubbles. Even if we assume as a starting point that the stock market is a random walk and

is governed by rational behavior, and even if we assert at the outset that all trades reflect the full consideration of the most up-to-date information, merely the fact that there are winners and losers will lead to booms and busts that have little to do with the rational application of information.¹

The simplest market cycle is based on two psychological characteristics of investors. First, their risk tolerance increases as their market winnings pile up. If you are making money, you will be willing to take proportionally more risk in the market.² This is often termed the house effect, because it is akin to successful gamblers who raise their stakes because they are playing with the casino's money. Second, the more people win, the smarter they think they are. If investors make money on a market view, they adhere more strongly to that view, even if the real reason behind their success has nothing to do with these beliefs.

If by nothing more than luck, every now and then some investors will find themselves on the right side of a big market move. For example, the forces of nature might tip the scales in the direction of the optimists: A flurry of positive news comes out, a wave of fresh cash is invested in the market, or the optimists are just more successful in convincing others of their viewpoint. Having done well, the optimists will tend to invest more; emboldened by their success, they will be willing to take more risk because of their increased wealth. They will view the positive events as confirmation of their views and will hold them with all the more fervor. Now that they have more wealth, they will be more comfortable with their market positions, and possibly be willing to add to them. Closet optimists who earlier did not have enough conviction to trade on their views will join the bandwagon.

Conversely, those who lose will reduce their participation in the market, both because they have less wealth to invest and because they will be less willing to put what wealth remains into risky—and from their recent experience, unprofitable—markets. So the lucky ones will put in more money while the losers, who might have otherwise been willing to put shorts into the market against them, will pull out. The initial move will grow in magnitude, sustained by the actions of the winners and losers. As the market continues to go the way of the winners, there is a coalescence of opinion.

The weighted range of market opinions becomes narrower, with more and more capital focused on the investment rationale of the winners. This

occurs because whereas at the outset there were bulls and bears, soon all the bears are out of the fight, licking their wounds. Those who could pull the price down through their selling have sold; those who bet against the market through short positions have been stopped out. If the consensus opinion in the market initially was distributed close to 50–50 between the bulls and the bears, it is now vastly different. Significantly, the shift in opinion will have nothing to do with new information coming into the market or with market participants becoming more informed or intelligent; the shift will be directly attributable to the effect that winning and losing has on the participants. Those who win will think they are smarter and will know they are richer; those who lose will be less willing to keep their dwindling wealth in the market. Of course, having trend followers doesn't hurt, either, but it just adds fuel to an already self-sustaining fire.

Count on the popular and financial press to provide all kinds of explanations for why the market really should be going up, even if the reasoning defies logic or past market behavior. If it is Japan in the late 1980s, the journalistic analysis is that the traditional relationship of prices to earnings does not apply because that country's accounting methods are different. If it is Internet stocks in the late 1990s, the pundits point out that the information age is based on a new paradigm of value that is not well captured by traditional methods of accounting and the related modes of fundamental analysis.

What has really changed is not the basic information—a P/E ratio is a P/E ratio—but the implications derived from it. Taken in its most extreme form, when there is a story that can afford unbridled optimism and when the optimism is fueled by levered exposures, a market bubble is born.

In the case of the Internet bubble, the cycle had an accomplice in the form of a restricted supply of stock, or float. The scarcity of Internet shares was such that the market impact of each buyer contributed more than usual toward inflating the bubble. The float of a stock is the number of shares actually in the market and available for trading. Frequently, when a stock first comes into the market, some of the shares outstanding are restricted from sale, either because they are held by insiders or because they have been issued to major backers with restrictions on their sale. The float as a percentage of shares outstanding for new Internet issues was very low by historical standards, often 10 percent or less of the total shares outstanding. A small float means it doesn't take much market interest to move

the stock price. For example, if a stock has a float of five million shares and the average investor holds around a thousand shares, the market value is determined by the trading of only a few thousand investors.

The most immediate effect of a small float is to focus demand on the subset of investors who are the most optimistic. If a company has a hundred million shares and puts all of them into the market, the company has cast a pretty wide net to find investors willing to buy the stock. The marginal investor is likely to be more reserved about the company's potential than the most optimistic investors who were first in line to snap up shares. However, if only five million of those shares are put into the market, then they will be bid up by the most optimistic fringe. The stock obviously will be priced higher than if more of the outstanding shares were available in the market. And it will not take much investor interest to create an explosive change in price.

The issuance of small floats repeatedly led to rocketing price rises in Internet companies. For example, after the World Trade Organization (WTO) accord with China in November 1999, China.com, a small Internet company that served the Chinese audience, nearly doubled in price in one day. The volume was nearly twice the float—that is, every available share traded two times. On the basis of the trading on the float, which comprised just 5 percent of the shares outstanding, the company attained a market value of several billion dollars. (The price rise might have had to do with little more than the company having the name “China” and a dot-com in its name. For example, China Pacific, an iron and steel company in Sichuan province, had been such a poorly performing stock that it was delisted by the NASDAQ earlier in the year. However, in the euphoria brought on by the WTO accord, the stock surged in price from pennies to more than \$2 a share just on the merit of its name.) Other investors saw the price increase in China.com as a confirmation of the value of the company and they bid the price even higher. Multiplying the total shares outstanding by the price on the 5 percent of the shares that were trading, the total market capitalization of the company rose to many billions of dollars. This exercise instantly put the company on the map and gave it an air of stability.

The impact of small float on prices was further accentuated during the bubble by the wide dispersion in views about the Internet itself. It's hard to attract much difference of opinion about the mundane old-economy

industries that just pump out products with known costs and uses. The views of the most optimistic investors may differ in only the second decimal place from those of the average or pessimistic investors. But the firms that are on the precipice of the digital, virtual new economy or that are opening the door to the information-centric new era—though devoid of visible means of support—leave the most optimistic investors free to imagine prospects that are orders of magnitude beyond those of the average investors who remain stuck in the linear-thinking mode of revenues and earnings. The optimism at the extreme was fed by the same sort of carefully orchestrated marketing that turns the Super Bowl from a football game into an event, or pulls a hundred thousand people into a stadium to hear the Three Tenors sing over loudspeakers. The Internet fervor was fed by a self-reinforcing and self-financing hype machine. Billions of dollars went from investors into the coffers of the Internet firms and then was spent on advertisements touting the firms, often in publications that breathlessly described their brilliance. And which experts did the publications quote? A set of Internet analysts, ranging from Mary Meeker at Morgan Stanley to Henry Blodget at Merrill Lynch, shills who pushed virtually every company they covered in hopes of snagging their investment banking business. During this period, buy recommendations among Wall Street analysts outnumbered sell recommendations by a ratio of nearly a hundred to one.

The virtual nature of Internet businesses meant that when the bubble burst everything disappeared—all the money was gone; there was nothing left behind. E-commerce companies, from eToys to CDNow, surpassed the market capitalization of their bricks-and-mortar counterparts precisely because they did not have fixed plants or an earnings-based business model, yet, in a short period of time, they disappeared without a trace. There were no rusting plants, no fire sales of equipment, no tumbleweed rolling past abandoned warehouses; this was a bubble with no soapy residue.

Besides isolating the most optimistic investors and setting a market price that reflects their views, the small float sets secondary forces in motion that accentuate the price appreciation. With a small float, investor demand for shares ends up being accommodated in a costly manner: through short selling and through high turnover. Short selling accommodates demand by creating virtual float. This requires a concession in price. Shorting a stock, especially a stock that has seen such price appreciation

and that is subject to widely differing views of valuation, requires sophistication, capital, and a risk appetite beyond that required from the buyers. Each rise in prices reduces the ranks of short sellers as some of those with losses pull out and others perceive greater risks in continuing to participate, so ever-increasing prices are required to pull in new short sellers to expand the virtual float.

High turnover satisfies demand for shares by continually recycling the existing shares in much the same way that an increase in the velocity of money increases money supply for a fixed monetary base. The float can be thought of as the stock equivalent of the money supply and the volume per share as the stock equivalent of the velocity of money. Double the velocity of shares and the virtual float will double. Absent unusual liquidity needs or information shocks, this turnover can be generated only by price action. (And it doesn't hurt to have a bevy of day traders to make this happen.) It will take more of a price move to pull 10 percent of the float into the market than to pull in 1 percent. The result will be higher prices and price volatility. The resulting price appreciation and high volatility required to meet market demand through these two mechanisms will increase the perception of value, or at least pique further investment interest.

As the price moves up, the company and its primary investors will inject more shares into the market, extracting cash. The float expands, and in the best of cases the price drifts back to earth like a hot air balloon cooling off. If the timing is right (right, that is, as far as the insiders are concerned), the influx of shares hits the market while the stock price is still riding high. The result is a net loss for the investor public. If the shares in the market were constant during the ride up and back down, at least on net there would be no loss among the investors. But with more shares hitting the market near the peak, it is no longer a zero-sum game, with one public investor's profit equaling another's loss. When supply catches up with demand, pricing power declines. The stocks of many Internet companies behaved like a roller-coaster ride over the year or two of the bubble, with a price at the end not far from where it was at the start, but the investors discovered that sometime during the course of the triple loops their wallets had dropped out of their pockets.

If expanding the floats of the existing stocks doesn't meet demand, then the market will happily comply by creating more initial public offerings (IPOs). In the Internet boom, new companies sprouted like dandelions in springtime. The quality and prospects of the businesses were not

at issue. They were within the paradigms of the Internet demand, so they were purchased as equivalent commodities. The reason so many businesses with no apparent prospects for profitability found takers was that the demand they were fulfilling was speculative, not economic. This dynamic is not new. It has been in play in the markets as far back as the classic tulip bubble in seventeenth-century Holland.

WHY TULIP MANIA WASN'T CRAZY

Many of the stories related to the Dutch tulip mania of the 1630s are apocryphal, drawn for sermons on the bitter fruits of avarice, the fodder for object lessons on prudence and frugality. These stories often tried to bring perspective to the phenomenon by introducing hapless, but arguably more rational, outsiders into the middle of the frenzy. One concerns a sailor returning to Amsterdam from sea, unaware of the tulip boom. At the offices of the shipping company's warehouse, he sees what appears to be an onion, sitting unattended, and helps himself to it. The onion is actually a very rare Rosen tulip bulb, which the warehouse owner had just purchased and put down for a moment. Searching far and wide for the missing bulb, the distraught owner finally comes across the sailor, sitting alone on the edge of the dock eating the last bits of the precious bulb, one worth more than the sailor's lifetime pay. A second story tells of an Englishman, similarly unaware of the great value that the Dutch ascribed to this flower, who out of curiosity dissects a bulb belonging to a Dutch friend. Again, it is an exceptionally rare specimen, an *Admiraal van der Eijck*, whose price could buy a house complete with furnishings.

The stories, fanciful though they may be, only add color to facts that are dramatic enough to stand unadorned: In the fall of 1636, the prices of tulip bulbs, even those of the most common varieties, were bid up manyfold. Some of these common bulbs were so richly priced that they could have been exchanged for a house or could have supported a tradesman's family for a decade. Contemporary accounts estimate that in just one Dutch town more than 7 million guilders of tulips were traded, an amount that exceeded the total capitalization of the Dutch East India Company, the largest trading company in Europe at the time. Then, within a few short days in mid-February of the next year, the market simply disappeared.

THE BRAVE NEW WORLD OF HEDGE FUNDS

The explanation for the tulip mania usually centers on the irrationality of the market—the incredible excess in paying unbelievable prices for mere flowers. This misses the point. Botany and horticulture were an avocation for many wealthy Europeans in the seventeenth century. And tulips in Holland were among the most sought-after and valued. The rarest, most highly prized tulips were those with irregular colors, or “flames.” They hardly ever reproduced, and when they did, their offspring often did not share the same flaming characteristics as the mother bulb.³ Paying exorbitant sums for these rare flowers was no more unusual than for the very wealthy today to pay fortunes for a Rembrandt (who, by the way, was painting the rich in Amsterdam in 1637). But even before tulip mania seized Holland—and for decades after the bubble burst—such rare bulbs fetched huge sums. In fact, just one month after the common bulbs that fueled the mania could no longer find a buyer, a quantity of rare bulbs was sold in Haarlem, a center of the tulip mania, for more than 10,000 guilders.

Tulip mania came neither from these rare bulbs nor from their collectors. The fervor of the collectors may have been a trigger for the bubble, but it was distinct from the event. Rare bulbs were the exclusive province of the wealthy, and in any case were in such short supply that they could not possibly have supported a broad market upsurge. No, the objects of the mania were ordinary bulbs, which were in abundant supply—at least to meet the usual demand. They were purchased by the pound, sight unseen, and were traded for the simple reason that, unlike the rare bulbs, they were available. True collectors would never have touched this merchandise.

The demand that fueled the mania instead came from merchants and the working class—bricklayers, plumbers, and other tradesmen; clergymen and lawyers—who had only secondhand knowledge of rare bulbs, and who conflated the values of the ordinary and the extraordinary. What is more, most of the demand was not based on the bulbs themselves, but on a trading frenzy that adopted them as its currency. The florists and traders who bid the prices up to such absurd levels had little or no interest in owning or cultivating the tulips.

FUTURES SHOCK, 1635

The speculative fervor was fueled by the development of a forward market for tulips in 1635. The flowers could not be disturbed and delivered until

the spring, so bulbs were traded for future delivery. The forward contracts were first developed to facilitate trading among florists and traders in the “colleges” (basically, rooms in the back of popular inns and taverns). The trading season for tulips traditionally went from June until September, when the buyer could see the flowers in bloom and they could be safely returned to their flowerbeds. So without a forward market, the traders would have had part-time jobs. The forward market turned tulip trading into a year-round speculative enterprise for the masses. (Indeed, the connoisseurs never did follow this practice; their tulips were purchased on delivery.) What changed hands was not the physical commodity, but a certificate of ownership. The market had already substituted common bulbs for the rare ones, so it didn’t take much more of a leap to simply trade the pieces of paper that represented future delivery of these bulbs. The commodity itself became an abstraction, a convenience for feeding speculation.

In the early stages of the tulip mania, it was common to pay for bulbs in kind, bartering them for land, tools, and farm animals. But with the advent of forward markets, where bulbs that were not owned were sold and bulbs were purchased for future delivery, a natural outgrowth was that the claim to ownership did not require any capital outlay. Money changed hands only once the bulb was delivered. Most traders did not have even the remotest intention of actually holding onto their contract to delivery, so they could buy the bulbs forward without the capital necessary to effect their physical purchase. The colleges were informal markets, limited in infrastructure to a secretary who made note of the transactions; there were no credit departments or systems for the posting of collateral. Indeed, while Holland had repeatedly legislated restrictions for trading on margin, the informality of the tulip market seems to have allowed it to remain outside the regulatory purview. Thus, a trader could sell bulbs he did not own, and had no connections to secure, and buy bulbs with no capital to purchase. He operated on the assumption that he would sell off his forward commitments long before they came due, and do so at a profit.

With the objects of adoration impossibly rare and beyond the reach of all but the wealthy, alternatives were provided to accommodate the growing demand. The alternatives bore a remarkable similarity to those employed nearly four centuries later during the Internet bubble, when dot-com stocks became nearly interchangeable in the minds of investors,

regardless of the prospects for the underlying business. So, just as anything with a “.com” appellation found a clamoring of buyers during the Internet bubble, during the tulip mania if the masses could not get the real objects of desire, then any tulip would do. A measure of value—rarity or even beauty—did not matter; all that mattered was that other speculators were willing to agree that the lower-quality bulbs would count. At the same time, the forward market in tulips facilitated an increase in the “velocity” of the existing bulbs—that is, the number of times each bulb was traded. The forward contracts turned over many times during the course of the winter. In fact, there are accounts of bulbs, or certificates for future delivery of the bulbs, changing hands many times in the same day.

The willingness of the market to accept the common bulbs as an object of speculative excess assured the end to the bubble. When it burst, the inescapable fact was that all the trading was for a commodity of little value, that even at the going prices could be produced in enough supply to accommodate any level of demand. The end came quickly, in February of 1637. At an auction in Haarlem, several lots of tulip bulbs found no buyers. The implication for other tulips and forward contracts was clear, and those who had obligations, who had purchased the bulbs in the expectation of selling them off for a profit before delivery came due, had to find real money and real demand. For the common bulb, there was none. There was the speculative market based on forward contracts and margin, and then there was a real market that lay a world below it. By the time the speculators had found that market, their potential fortunes would be in ruins.

The seeds of the tulip mania may have been the unattainable lure of fashionable and rare tulips, combined with the newly accepted practice of substituting more common bulbs to meet that appetite. But the mania reached full bloom only with the innovation of forward contracts and the leverage these contracts afforded, which allowed traders to buy and sell commodities they did not own, had no intention of owning, and indeed did not even have the money to purchase outright.

“THIS CRAP IS GOING TO BE WORTH ZERO”

In the bucket shop era of the early 1900s, stock operators manipulated prices and churned positions to attract hungry investors who were looking to make a killing. In the Internet era, the mechanism of a very narrow

float was used to inflate stock prices and churn legitimate investors. In both cases, a high volume of trading was essential to move prices higher. In contrast, value-based market appreciation does not need much trading to elicit an increase in the market price. In fact, in the extreme case, envisioned by the efficient market paradigm, new information that creates a shift in value will lead all investors to revalue their positions, so that the price will change to a new equilibrium level without any trading.

The change in wealth with market appreciation can itself propel market prices, independent of any new information or any change in the view of a company's fundamental value. When people have more money, they are willing to put more of it at risk, especially when it came from speculative investment in the first place. There is a wealth cycle that can drive price appreciation, as investor demand leads prices to be bid up and as investors put more and more of their wealth into the market. The smaller the quantity of stock available to meet this demand, the higher the price will go.

At the time of the first major run-up in Internet stocks in 1998, the float of the major Internet stocks such as Amazon, Excite, and Yahoo! was around 10 percent of the total shares outstanding. Daily volume averaged over 10 percent of float, compared with a volume/float ratio of 1 percent for the similar but more established companies like Intel, Microsoft, and even AOL. In contrast, in a value-based market the float will not matter, because investors will not be willing to buy the stock at higher than fair value, and fair value will be determined by the total shares of the stock outstanding, not by the stock that happens to be in the market.

The warning signs for this cycle therefore were large price changes coupled with large flows, high volume-to-float ratios, high short interest relative to float, and—a natural outgrowth of these factors—high volatility. There were also disparities in correlation. The places where the winners were piling on their investments became hot spots—sectors with high internal correlation but low correlation with the rest of the market.

The end of the Internet bubble arrived when the period of restrictions on trading IPO holdings passed and the float expanded beyond the point of the demands of the extreme Internet optimists. Up to October 1999, the amount of IPO issuance that had become unlocked amounted to less than \$40 billion. Then, in the last quarter of 1999, the IPO unlock exploded, with more than \$50 billion worth released in December alone.

In January 2000, another \$65 billion worth of IPO shares were unlocked and came into the market, and nearly \$100 billion worth was released in the following three months. There was now more stock looking for buyers than there were super-optimists ready to buy. The marginal share had to find a buyer who was not all that hyped up about Internet prospects, and the downward spiral began.⁴

What is surprising in all of this is that the flood from unlocked IPO shares and the implications this would have on the minuscule float were public information. Anyone could know how much would be freed up and when it would come to market. While this became a subject of research interest, the implications it might have for the Internet bubble led to little action at the time. This was not for lack of notice.

For a moment during the Internet bubble, Warren Buffett, a world-class investor, looked like Warren Befuddled. So did hedge fund masters such as Julian Robertson, whose Tiger Management fund was beaten to the ground by his doggedly rational addiction to value. George Soros, who broke the Bank of England, was broken the same way—fighting the equity crowds clamoring to buy more stock. His book *Irrational Exuberance* (Princeton University Press, 2000) earned Robert Shiller notoriety for his call that we were in the middle of a bubble, but in fact many investment professionals recognized it for what it was, the chorus getting louder as the bubble swelled. However, it is one thing to say the market has run amuck; it is another thing to trade against it. One of the most surprising fallouts of the Internet bubble was the closing, in rapid succession, of both Robertson's and Soros's long-admired funds. Each took a different path in trading against the "insanity," and each failed.

Robertson's Tiger investment funds enjoyed a remarkable run, returning an average of more than 30 percent each year in its first 18 years. Even after stumbling badly in its last two years, with a drawdown of nearly 50 percent, it still returned 25 percent annually over the course of its existence. Robertson persisted in a strategy that made sense before the Internet boom and after the bubble burst, but was poison in between: value investing. He elected to ignore the bubble, during which he felt "earnings and price considerations take a backseat to mouse clicks and momentum," and stay the course as his stocks dropped in attraction and in value. Unfortunately, he had little choice but to do so because he held positions that were so large as to be illiquid. The most prominent of these was a 25 percent stake in

US Airways; the value of Tiger's position dropped by two-thirds from the time Robertson acquired his stake, and there was no way to sell the position off. Indeed, since his position had become so well known, even rumors that he might be thinking of selling his holdings would sink the stock price. (Two months after Robertson closed shop in March 2000, vindication for his keystone position came in for a landing. US Airways agreed to be bought by United Airlines, and its price nearly doubled overnight.)

Bad as it appeared, Tiger's last year and a half was perhaps even worse than the 50 percent drawdown suggests. The fund may well have lost more money in the last 18 months of its existence than it had earned in all of the previous 18 years. It started in 1980 with less than \$10 million in capital; by 1998, it had in excess of \$22 billion of capital, the vast majority of this money coming from new investments. And it is from this height that it began its fall into oblivion, generating losses of some \$10 billion from peak to closing. This is almost certainly more than all the dollars it earned for its investors previously, when the fund posted spectacular returns, but on a much smaller dollar base. This means that when all the chips are counted, one of the most stellar traders of all time may well have spent 20 years and acquired a personal fortune while on net losing money for his investors.

The same might also be said of the other luminary of the hedge fund universe. George Soros's investment performance eclipsed even that of Robertson; his Quantum Fund generated average returns exceeding 30 percent over more than three decades. Soros made his largest gains in the macro markets of foreign currencies, where he was best known for a huge bet against the British pound that netted him nearly \$2 billion. Like Robertson's funds, at their peak Soros's funds commanded more than \$20 billion of capital.

While Robertson stuck to his value-investing guns, Stanley Druckenmiller, the manager for Soros's Quantum Fund, didn't fight the tape. The Internet bet pulled Druckenmiller out of a hole in 1999, moving the fund from negative territory to plus 35 percent for the year. But with large positions, it was impossible for the fund to navigate the shallow float of the Internet stocks, and when the market turned in March 2000, Druckenmiller could not climb out of the positions without sizable losses. In March the fund lost 12 percent; through April another 20 percent.

Robertson blamed his demise on "irrational markets." Soros similarly professed bewilderment, saying that "historical measures of value no

longer apply.” Robertson justified his direction, stating, “We are not smart enough to know when they will begin to perform, but we do know that they are value. Buying value has always been our strategy, and it’s our strategy now. And we do know that value will win out eventually.” Alas, too late for him.

In contrast to these two funds, yet another sage of the markets remains open for business: Warren Buffett. Like Robertson, he did not understand where the value was coming from for the Internet stocks, and like Robertson he elected to sidestep the party. But he differed from Robertson in the concentration and leverage that he put into play. Although his holding company, Berkshire Hathaway, dropped nearly 50 percent in value by early 2000, it could ride out the frenzy in part because it owned actual businesses—furniture, newspapers, jewelry—that had actual sales and profits.

If value is your religion, the Internet bubble left you with damnable choices. You could try to ride the bubble up and be quick enough to bail out before it burst, as Soros’s Quantum Fund attempted. Unfortunately, the success of that approach depended critically on the ability to move out of the positions when the market turned. The Quantum Fund was an elephant trying to ride a balloon. It couldn’t work. Quantum was too heavily invested in relatively thinly traded stocks—at least stocks that became thinly traded when the market turned. A second approach was to keep the faith in value, ignore the bubble worshippers, and stay focused on the stocks that had underlying value. Characteristically, Robertson and Buffett chose that path, although there was a critical difference between the two. Buffett had no problems with redemptions, because his investors owned shares that they could sell, and did, at whatever the market would bear. Robinson needed to worry about redemptions, as some of his investors defected to Internet stocks. But his fatal mistake was to lever and take on illiquid positions; the combination was like a vice that squeezed his ability to adjust his positions and thus the fund’s staying power. Buffett, while having his reputation tarnished in the short term, remained intact.

Early on during the Internet mania I discussed the inflated valuations of many of these companies with Stan Shopkorn, a great equity trader who was vice chairman of Salomon in the early 1990s. He held positions in a number of them, and I asked him how he could be in these markets. “Well,” he said, “that’s what makes me a trader and you a risk manager.”

While Buffett kept away from a market he could not understand, Robertson traded the value stocks against it and Soros realized after the fact that “we went into the new economy but we overstayed our welcome.” In the midst of the bubble, one manager who had watched his investments in Internet IPOs grow twenty-fold, ballooning to upwards of \$500 million, asked me how he could escape with his profits. He was restricted and could not sell. He knew a day of reckoning would come, and that “in a year or less this crap is going to be worth zero.” Well, there wasn’t much he could do; before he finally could get out, the bubble burst.

When the cycle breaks, a large price decrease will accompany an imbalance to sell. The trick is knowing when that imbalance will occur. For an unpredictable period of time, the wealth feedback gives reconfirming signals that endure too long for many people to resist. They can’t remain on the sidelines watching everyone else get rich, or hold on to a short position going south. It is all in the timing. Even the tulip mania in seventeenth-century Holland lasted for more than two years, but when it ended, you didn’t want to be the one holding the tulips.

PAIRING OFF: THE EMERGENCE OF STATISTICAL ARBITRAGE

In the perfect market paradigm, assets can be bought and sold instantaneously with no transaction costs. For many financial markets, such as listed stocks and futures contracts, the reality of the market comes close to this ideal—at least most of the time. The commission for most stock transactions by an institutional trader is just a few cents a share, and the bid/offer spread is between one and five cents. Also implicit in the perfect market paradigm is a level of liquidity where the act of buying or selling does not affect the price. The market is composed of participants who are so small relative to the market that they can execute their trades, extracting liquidity from the market as they demand, without moving the price.

That’s where the perfect market vision starts to break down. Not only does the demand for liquidity move prices, but it also is the primary driver of the day-by-day movement in prices—and the primary driver of crashes and price bubbles as well. The relationship between liquidity and the prices of related stocks also became the primary driver of one of the most powerful trading models in the past 20 years—statistical arbitrage.

If you spend any time at all on a trading floor, it becomes obvious that something more than information moves prices. Throughout the day, the 10-year bond trader gets orders from the derivatives desk to hedge a swap position, from the mortgage desk to hedge mortgage exposure, from insurance clients who need to sell bonds to meet liabilities, and from bond mutual funds that need to invest the proceeds of new accounts. None of these orders has anything to do with information; each one has everything to do with a need for liquidity. The resulting price changes give the market no signal concerning information; the price changes are only the result of the need for liquidity. And the party on the other side of the trade who provides this liquidity will on average make money for doing so. For the liquidity demander, time is more important than price; he is willing to make a price concession to get his need fulfilled.

Liquidity needs will be manifest in the bond traders' own activities. If their inventory grows too large and they feel overexposed, they will aggressively hedge or liquidate a portion of the position. And they will do so in a way that respects the liquidity constraints of the market. A trader who needs to sell 2,000 bond futures to reduce exposure 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 dumps 2,000 contracts into the market, 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 selling was not based on an informational edge.

So the principal reason for intraday price movement is the demand for liquidity. This view of the market—a liquidity view rather than an informational view—replaces the conventional academic perspective of the role of the market, in which the market is efficient and exists solely for conveying information. Why the change in roles? For one thing, it's harder to get an information advantage, what with the globalization of markets and the widespread dissemination of real-time information. At the same time, the growth in the number of market participants means there are more incidents of liquidity demand. They want it, and they want it now.

Investors or traders who are uncomfortable with their level of exposure will be willing to pay up to get someone to take the position. The more uncomfortable the traders are, the more they will pay. And well they

should, 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. Thus the demand for liquidity not only is the source of most price movement; it is at the root of most trading strategies. It is this liquidity-oriented, tectonic market shift that has made statistical arbitrage so powerful.

Statistical arbitrage originated in the 1980s from the hedging demand of Morgan Stanley's equity block-trading desk, which at the time was the center of risk taking on the equity trading floor. Like other broker-dealers, Morgan Stanley continually faced the problem of how to execute large block trades efficiently without suffering a price penalty. Often, major institutions discover they can clear a large block trade only at a large discount to the posted price. The reason is simple: Other traders will not know if there is more stock to follow, and the large size will leave them uncertain about the reason for the trade. It could be that someone knows something they don't and they will end up on the wrong side of the trade once the news hits the street. The institution can break the block into a number of smaller trades and put them into the market one at a time. Though that's a step in the right direction, after a while it will become clear that there is persistent demand on one side of the market, and other traders, uncertain who it is and how long it will continue, will hesitate.

The solution to this problem is to execute the trade through a broker-dealer's block-trading desk. The block-trading desk gives the institution a price for the entire trade, and then acts as an intermediary in executing the trade on the exchange floor. Because the block traders know the client, they have a pretty good idea if the trade is a stand-alone trade or the first trickle of a larger flow. For example, if the institution is a pension fund, it is likely it does not have any special information, but it simply needs to sell the stock to meet some liability or to buy stock to invest a new inflow of funds. The desk adjusts the spread it demands to execute the block accordingly. The block desk has many transactions from many clients, so it is in a good position to mask the trade within its normal business flow. And it also might have clients who would be interested in taking the other side of the transaction.

The block desk could end up having to sit on the stock because there is simply no demand and because throwing the entire position onto the

floor will cause prices to run against it. Or some news could suddenly break, causing the market to move against the position held by the desk. Or, in yet a third scenario, another big position could hit the exchange floor that moves prices away from the desk's position and completely fills existing demand. A strategy evolved at some block desks to reduce this risk by hedging the block with a position in another stock. For example, if the desk received an order to buy 100,000 shares of General Motors, it might immediately go out and buy 10,000 or 20,000 shares of Ford Motor Company against that position. If news moved the stock price prior to the GM block being acquired, Ford would also likely be similarly affected. So if GM rose, making it more expensive to fill the customer's order, a position in Ford would also likely rise, partially offsetting this increase in cost.

This was the case at Morgan Stanley, where we maintained a list of pairs of stocks—stocks that were closely related, especially in the short term, with other stocks—in order to have at the ready a solution for partially hedging positions. By reducing risk, the pairs trade also gave the desk more time to work out of the trade. This helped to lessen the liquidity-related movement of a stock price during a big block trade. As a result, this strategy increased the profit for the desk.

The pairs increased profits. Somehow that lightbulb didn't go on in the world of equity trading, which was largely devoid of principal transactions and systematic risk taking. Instead, the block traders epitomized the image of cigar-chewing gamblers, playing market poker with millions of dollars of capital at a clip while working the phones from one deal to the next, riding in a cloud of trading mayhem. They were too busy to exploit the fact, or it never occurred to them, that the pairs hedging they routinely used held the secret to a revolutionary trading strategy that would dwarf their desk's operations and make a fortune for a generation of less flamboyant, more analytical traders. Used on a different scale and applied for profit making rather than hedging, their pairwise hedges became the genesis of statistical arbitrage trading.

Although IT is a support function, during the mid-1980s Morgan Stanley's technology department was operated as a fiefdom by Bill Cooke, who parceled out programming services like a sovereign handing out favors. A young programmer in this group, Gerry Bamberger, was assigned to work on the equity trading floor to improve the block desk's ticket entry process so that the traders could track their trades and their intraday

P&L more efficiently. Part of his task was to identify and marry blocks with their paired hedges so the two could be monitored as a single entity.

As he ran through the block trading positions, Bamberger noticed a common behavior between the two stocks in paired hedges. The price of the stock that was the object of the block trade would run up or down in price depending on whether there was a buy or sell order. Once the block trade was completed, more often than not the stock price would then settle back down to where it started, while its hedging partner would usually not move much at all. These were common characteristics of the pairs stocks in block trades. Moving a big block of stock demanded some sort of price concession. The art of block trading was to minimize this concession.

What was interesting, but expected, was the effect of the occasional sudden news event on the pair. If a surprise event hit the market, both sides of the pair moved together. Both stocks were in the same industry, so any information in the market that was not specific to just one of the two stocks would tend to affect them both in the same way. What would affect the pair's price—and what the pairing of the stocks helps to isolate out of the market cacophony—was the demand for stock-specific liquidity.

Bamberger started to think of the pairs not as a block to be executed and its hedge, but as two sides of a trading strategy. If all the pairs the block desk maintained were monitored and one started to rise unaccountably, one of two things was possible: On the one hand, the stock that was rising could be the object of a block trade or related liquidity demand. If so, its rise would be limited. Once the demand had been filled it would revert to its original spread relative to its partner. On the other hand, the rise could be the result of some information that was specific to that stock. Because information is by definition random, the future course of the stock would also be random. It might continue to rise, or it might turn the other way. Each stock was paired with another stock, so only company-specific information would affect the relative pricing of the pair. Any broader information would make both stocks move, and the relative value of the pair would remain unchanged. The company-specific effects could be diversified away by holding many pairs since they would be independent from one company to another. In this way Bamberger's trading strategy gave obeisance to efficient markets. He assumed information would move prices as a random walk, but he constructed the portfolio in a way that negated the impact of that information, at least in a statistical sense. That

left the portfolio solely exposed to liquidity demand. His strategy of selling the better-performing of a pair while buying its underperforming partner should, on average, make money.

Bamberger took some historical data for a set of the pairs to test his hypothesis. The effect was subtle, because for any individual pair the relative price movement was drowned out by a much larger amplitude of noise coming from the overall market. But using statistical analysis he could show in the aggregate a clear tendency toward mean reversion—the stocks would get back in line with their history. There was money to be made, but the key was to hold many, many pairs to average out the market effects.

The pairwise stock trades that form the elements of statistical arbitrage trading in the equity market are just one more flavor of spread trades. On an individual basis, they're not very good spread trades. It is the diversification that comes from holding many pairs that makes this strategy a success. But even then, although its name suggests otherwise, statistical arbitrage is a spread trade, not a true arbitrage trade.

Bamberger pitched his strategy, and surprisingly, given the politics of the firm, the equity division was willing to let him give it a try. He got a desk just to the side of the futures traders, mounted with monitors that tracked every pair: Ford and GM, American Airlines and United Airlines, International Paper and Georgia-Pacific, and so on. People were dumbfounded by the results: In his first year of trading he made \$6 million.

Then, in the classic Morgan Stanley style of the time, office politics reverted to the mean, too. Bill Cooke, the head of the systems area and Bamberger's boss's boss, gave one of his lieutenants, Nunzio Tartaglia, oversight for the operation at the latter's request. With money lying on the table and a scrappy young programmer as the only impediment in his path, it was inevitable that Tartaglia would wrest the pairs program from Bamberger. The only unusual thing about this power play was that it did not follow the regular pattern of glomming onto business by way of old school ties.

In Morgan Stanley's clubby world, Tartaglia was a strange breed of nonconformist. He was an engineer, not a financier, which he made clear by wearing white short-sleeved shirts and a pocket protector. He was pint-sized, muscular, and tough. He had studied briefly in a Jesuit seminary,

and in a different age he might have been on the leading edge of the Inquisition. He trusted no one. To protect his authority, he did not allow those reporting to him to meet or to talk with one another. It was clear to Bamberger that he would last about as long as it took Tartaglia to understand the computer code behind the trading system. They had one conversation and that was it. While many who were at Morgan Stanley at the time can retell the story of Bamberger's discovery of stat arb, its epilogue remains a blank page. Bamberger headed off to Princeton Newport, where he continued to pursue the strategy with the legendary Ed Thorp for a year, made a small fortune, and then disappeared from the world of trading.

This move was a natural. Thorp was the first great analytically oriented trader. An MIT mathematician, he was the author of *Beat the Dealer* (Random House, 1966), which introduced the concept of card counting to blackjack. He wrote the book to spread his techniques after he had been banned by most casinos and could no longer pursue the strategy for profit. The next year he wrote *Beat the Market* (Random House, 1967), which formed the basis of convertible bond arbitrage, a strategy that remains the mainstay of many hedge funds. While I was still a professor in 1983 (the year before I left academia for Wall Street), I got a call from Ed to interview at his firm in Newport Beach, California. The visit was a culture shock, a little like Woody Allen's trip to Los Angeles in *Annie Hall*. I was met at the airport by one of Thorp's lieutenants, who drove a Porsche, wore his sweater tossed over his back with the arms draped over his shoulders, and referred to his girlfriend as his "lady." When we got to the Princeton Newport offices, he parked next to two other Porsches, owned by other members of the firm. Ed met me in tennis whites and as we walked down the corridor to his office we stopped by a map where he showed me the various jogging trails he used. At dinner that night he picked me up in *his* Porsche, the prized 935 model. At the time his house was the most expensive in the city, a \$7 million hilltop estate with an observatory. I got an offer but in the end it didn't work out, and fortunately so: The Princeton side of his firm and the source of his funding, Oakley Sutton Management, became embroiled in tax litigation, and the firm eventually dissolved. Still, the visit warmed me to the idea of working outside of academia, so when I got calls from investment banks the next year, I was ready to listen.

ARRIVEDERCI TARTAGLIA

Tartaglia matched Bamberger's revenue of \$6 million the year after he took over the strategy. He started a new department at Morgan Stanley christened Analytical Proprietary Trading (APT). He automated Bamberger's techniques, linked them to the SuperDOT network that had been developed for program trading and index arbitrage, and applied them to an array of thousands of stocks, often holding a portfolio containing more than 600 names at a time. In 1986, with his new scale of operation, he brought in \$40 million.

As the money rolled in, the department's size and accoutrements swelled. APT grew to 40 professionals and an endless array of high-tech toys. They bought Silicon Graphics machines that used 3-D glasses to try to discern patterns in stock data. They outfitted a corner conference room with a huge doughnut-shaped table and state-of-the-art electronics. Tartaglia brought in mathematical talent ranging from computer scientists to chaos theorists. He nurtured pet projects, including the application of complex differential equations to model stock market movements. But for all the growth in body count and intellectual horsepower, the only strategy that made money was the one that Bamberger had been forced to hand over to Tartaglia—a strategy that could have been run with two or three people. In 1988 the stat arb strategy took in only \$6 million, far from enough to pay the bills for the bloated group. Tartaglia, who had squandered his political capital by keeping the entire firm at arm's length, was out of the firm and APT was closed down.

Analytical Proprietary Trading died a quick death but spawned a thousand children. One was David Shaw, a Tartaglia recruit. A faculty member at Columbia with a PhD in computer science from Stanford, he joined APT in 1986. Two years later, he left to start his own stat arb fund following a row with Tartaglia, who felt he was trying to steal his thunder. Almost from the start Shaw's fund pulled in \$100 million plus a year from the strategy. Although he has since limited his day-to-day involvement, pursuing research in biotech instead, his firm, D.E. Shaw, continues to this day. (One of his Princeton recruits, Jeff Bezos, left his ranks in 1994 to start his own enterprise, Amazon.com.)

Thanks to Gerry Bamberger, who started as a programmer on Morgan's equity desk, the way trading was done and the function it performed

had changed. As a result of his work, the computational power for statistical analysis was unleashed on the markets and—using the newfound execution capabilities of the equity market—a machine was created to harvest opportunities to provide liquidity. Bamberger had moved at least one segment of the market from that of hunter-gatherer to farming.

A remnant of APT remained in the form of a small trading group run by a latecomer to APT, Gene Flood. The operation was struggling and I was asked to join the group and help out. In May of 1990, I met Gene for the first time in the group's office, a large converted conference room in an obscure corner of an administrative floor in Morgan Stanley's midtown offices. Gene came to Morgan Stanley with a Harvard undergraduate degree, a PhD from MIT, and a warm, disarming demeanor that rested on the foundation of strong Christian beliefs. He was a refreshing figure in the tangles of Morgan Stanley's trading operations.

My first look under the hood was somewhat of a letdown from what I had expected from a group run by someone with such strong academic training. Rather than carrying on the scientific lineage of APT, the group used standard technical trading methods such as the relative strength index (RSI), moving averages, and breakouts. In place of the battery of computers that fed APT's mean-reversion analysis, the group's strategies were displayed on two walls lined by whiteboards, one board for each strategy. Black magic markers designated long positions; red markers, short positions.

One of Flood's immediate objectives was to raise the group's profile by acquiring better real estate. He wanted to move to the fixed income trading floor, even though our fledgling trading group's strategies had nothing to do with the business of fixed income customer flows. But, like staffers vying for a cubbyhole of an office in the West Wing, many computer programmers and research staff gladly gave up quiet offices to lay claim to six square feet of desk space on the trading floor—anywhere on the trading floor, even some out-of-the-way nook where a risk taker had never set foot. In Flood's case, maybe the idea was that what was lacking in P&L could be jumpstarted by being closer to the action and to the source of trading ideas. In any case, we found ourselves at a row of desks next to some of the back-office clerks who cleared trades. We worked while listening to the drone of CNN's audio on the Gulf War, piped in all day long through one of the clerks' speakerphones.

The change in venue did not change the outcome, and the group died less than a year after I joined. I then moved to the Treasury trading desk, where I continued my proprietary trading. Each night I commandeered six of the firm's most powerful computers, IBM RISC workstations. Considered powerful at the time, they likely wouldn't even be found in a scrapheap today. Five of them had 256 MB of memory. The sixth was souped up with 512 MB with a 2 gigabyte hard drive. That one cost more than \$1 million. The workstations did the bidding of the mortgage and derivative groups during the day, but lay idle at night. I got my runs ready to go and launched them before leaving work at 7 P.M., then gave them a bedtime check from home to verify they were humming along. The next day I would get to work early, clear my runs from the computers, and spend the day going through the results and programming the next batch of analysis.

All the horsepower was needed because I was trading based on tick-by-tick data. The objective, like that of statistical arb pairs trading, was to divine, through the pattern of tick-by-tick trading, instances where the price move was based on liquidity demand, and thus would likely revert over time. The market maker at a broker-dealer often has information that provides an advantage in differentiating liquidity-based from information-based trading. For example, market makers know who is coming to them to make a market, and often also know if that person tends to trade based on information or liquidity. But a trader with nothing other than the market on the other side can never be sure that a price discrepancy is the result of liquidity demand. The discrepancy could be caused by information or even manipulation. This possibility leads to uncertainty about the quality of the security, will lead prices to move further than they would otherwise, and will also make the market less liquid.

Both stat arb traders and market makers share a problem called adverse selection. They don't know what they don't know, which is easily described in the context of used cars. When you're buying a new car it is difficult to know its relative quality, but once you've driven it for a while you can assess if it is a good car or a lemon. You now have information that no one else does. A customer who buys a lemon will likely try to get rid of the car by selling it in the used car market. No one in the market will be able to tell easily whether it is good or bad, so the owner will be able to sell it for more than it is really worth. In contrast, anyone who is fortunate

enough to buy a good new car will be more likely to hold on to it: If the owner sells it to buy another car, he or she once again goes into the good car/bad car lottery and may not be as lucky the next time around.⁵

Just as a car buyer can never be sure whether information is being withheld by the seller, in the financial markets a buyer can never be sure whether there is something going on with a stock that is beyond his purview. The person on the other side of the trade might have insider information on the company, or he might know that there is a much larger overhang of potential selling, the demand the buyer sees being a first trickle in what will emerge as a flood of selling. Even beyond the problem of adverse selection, if somebody waves a white flag and tries to overcome the adverse selection bias by announcing who they are—confirming to everyone's satisfaction that they are trading strictly because of a liquidity need and 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 there is no guarantee the price will not fall further between the time the trader takes on the position and the time he seeks to resell it. There may be many other liquidity-driven sellers behind this sale, or there may be a surprise economic announcement that affects the market.

The adverse selection problem is especially troublesome for market makers, and particularly for market makers in specialized arenas, such as corporate bonds, mortgage securities, and emerging markets. Their long-term business depends on standing ready to meet transaction demand. A bank's emerging market trading desk might like nothing better than to duck for cover when a crisis occurs, but it has to be ready to make trades regardless of the situation. The desk's longer-term livelihood depends on keeping things orderly, and that means pulling in its clients' inventory and watching its own inventory grow even though it knows things will likely get worse before getting better. Of course, the desk can try to move prices down to stem market demand to sell. In theory, lower prices will lead fewer to sell and will finally get some fence-sitters to buy. But it doesn't work in a crisis, as we saw during the 1987 crash.

The liquidity supplier should expect to make money on the trade because of the risk faced in accommodating the seller. The more volatile the market, the higher the probability that prices will run away from the liquidity suppliers. The less liquid the market, the longer the position must

be held, and so the greater the risk. The liquidity supplier should also expect to make money because of the opportunity cost in holding cash free for the speculative positions and the time spent monitoring the market and making trading decisions.

Though Morgan Stanley management killed the stat arb golden goose, it revived the approach in 1992 when the equity division hired a young Princeton graduate, Peter Muller, to build a stat arb system. He started from scratch, plugging away on the SAS statistics package to build his approach. Peter's first months of trading at Morgan Stanley were rocky; he was one month from being shut down, but then he hit his stride, beginning a streak of years of profits with nary a month of down performance. The computerized trading pulled in hundreds of millions a year for the firm and tens of millions for Muller, and the system was sufficiently automated that he could leave the day-to-day administration to his staff while he indulged in his other interests—performing his songs at New York clubs and kayaking in Ecuador—before finally retiring from the firm in 2001.

Besides the use of pure mean reversion between pairs, Muller and other stat arb traders exploited another market characteristic: leads and lags in the price changes of stocks within the same sector. When information or liquidity hit a market sector, the first stocks to respond were the larger, more visible names. For example, Wal-Mart and Costco would be among the first names investors would think about if they wanted to change their holdings of discount stores. Their prices would move first, and then the effect would trickle down to the less traded names like PriceSmart and Fred's. Rather than mating pairs, stocks in each sector were divided into groups of tightly aligned leaders and laggards. If the prices of the leaders ran up while the laggards remained unaffected, the laggards would be bought and the leaders would be shorted to hedge against them. The traders used the price moves of the leaders as a signal to take positions in the laggards, so they would be ready to meet the wave of liquidity demand when it finally hit those stocks.

Operationally, the lead-lag strategy looked much the same as the original Bamberger mean-reversion pairs trading. Stocks that have already moved up are sold against a similar group of stocks that, relatively speaking, have underperformed. But the dynamics are more refined and the selection criteria are cleaner. With the original pairs trading, one of the

stocks could be going up for reasons that have nothing to do with the overall sector. In those cases, there is no expectation it will revert back to the level of the other stock in the pair or that the other stock will later follow suit and start to rise. The simpler pairs trade is based on the idea that frequently a relative change in a stock price will be short-lived because it is due to a transitory liquidity need. When that is not the case, the effect of the future course of the stock's price on the overall trading portfolio will be diversified away—its price effect will be drowned out—by the large number of pairs that are being held.

With the lead-lag strategy, a whole group of leading stocks is monitored. The odds that a group of stocks will move together based on random effects are much lower than the odds that one stock of a pair will start to deviate. If this group goes up as a whole, that is an indication the move is sector driven, and if they go up while the laggards do not, that is an indication there will be a lead-lag effect that can be exploited. The subsequent profit can come not just from this leading effect dissipating before the laggards are affected, with the usual stat arb result of a mean reversion in price; it can also come from the demand reaching the laggards and having them run up in price as well.

Statistical arbitrage is now past its prime. In mid-2002 the performance of stat arb strategies began to wane, and the standard methods have not recovered. This is not surprising, given the simplicity of the strategies, the ease of entry, and the proliferation of computer power. My son David had the bad luck to get started in this sort of strategy just as the window of opportunity was closing. The strategy had performed admirably in years of back-tests and in the first months of operation, but then sputtered along doing next to nothing. He closed it down from active trading after six months and then ran it on paper for another year, with no better results.

The stat arb concept remains, but in place of the stat arb strategies of the late 1980s and the 1990s is an incarnation called high frequency trading. It performs the same liquidity function, but by monitoring aberrations in supply and demand based on real-time information. It is an arms race fought as much in the IT battlefield as in finance, where the edge from eliminating one node in the order entry route—shaving a few milliseconds off of an execution—can determine the success of a trade.

LONG-TERM CAPITAL MANAGEMENT'S SCANDALOUS BIRTH

I was in the office of Salomon Brothers chairman Bob Denham for a casual meeting when he turned to answer the phone. The call was brief, with hardly a word spoken. He then turned to me and said, "They gave him nothing."

The "him" was John Gutfreund, formerly the King of Wall Street, later the vanquished chairman of Salomon. Gutfreund had joined Salomon in 1953 after college because his father was a golfing buddy of Billy Salomon, the son of one of the founding brothers and a decade later the first managing partner of the firm. At that time the Wall Street world, especially the lower rungs Salomon inhabited, was not a destination address for aspiring graduates. For Gutfreund, it was "just a job." He quickly moved up in the ranks and in the 1960s used Salomon's distribution network to push the firm into bond underwriting. It was Salomon's clients, after all, who were the institutional investors buying the bonds. By the mid-1970s Gutfreund was the heir apparent, finally taking the helm from Billy Salomon in 1978.

The "they" that Denham referred to was the arbitration board that heard Gutfreund's demand for compensation following his resignation from Salomon in the aftermath of the 1991 U.S. Treasury securities scandal. In that episode, the head of the government trading desk, Paul Mozer, submitted false bids in Treasury auctions. Mozer was trying to circumvent a 35 percent limitation on the maximum bid for the supply that any bidder could submit in an auction of Treasury securities, a limitation that came about as a direct result of Salomon's earlier activities in submitting bids in amounts far in excess of the amount of securities being auctioned.

These bids, submitted at the direction of Mozer, were just one more manifestation of the disdain he had for the Treasury regulations that got between him and profits. He also expressed his disagreement with the decision to adopt the new rule in discussions with Treasury officials, in news articles, and through immature so-called protest bids placed far above the market rate.

The specific trigger for the scandal was the February 1991 five-year note auction, in which Salomon submitted a bid in its own name for 35 percent of the auction amount and two additional bids in the names of two Salomon customers, Quantum Fund and Mercury Asset Management.

The last two were false bids made without the knowledge of either customer. After the auction, Mozer had the clerk write trade tickets “selling” the auction allocations from these two unauthorized bids to customer accounts in the names of Quantum Fund and Mercury Asset Management at the auction price. At the same time he had the clerk write tickets “selling” the same amounts from those accounts back to Salomon at the same price. The idea was to make it look like the customers had received the securities awarded in response to the unauthorized bids and had then sold them to Salomon. Mozer then circumvented the firm’s procedure by preventing confirmations from being sent to the customers, who understandably would have been puzzled if they heard about the transactions.

Unfortunately for Mozer, it happened that an entity related to Mercury Asset Management also participated in the auction, and between its bid and the fictitious Salomon bid, the 35 percent limitation was exceeded. This resulted in a letter from the Treasury Department to Mercury Asset Management with a copy to Mozer. The cat was out of the bag, so taking matters into his own hands, Mozer brought the letter to his supervisor, John Meriwether. When Meriwether finished reading it, Mozer told him that the Mercury Asset Management bid referred to in the letter was in fact a bid for Salomon and had not been authorized by Mercury. Meriwether understood the seriousness of the matter and passed the letter on to Salomon’s general counsel, Donald Feuerstein, and to its president, Thomas Strauss. Gutfreund was out of town, so the men convened when he returned. When the meeting finally took place near the end of April, Feuerstein told the group that he believed Mozer’s actions were criminal and had to be reported. Incredibly, no report was made until mid-August.

In reconstructing the reason for the delay, the government investigation suggested that each of the four had a different understanding of the way the revelation would be handled. Meriwether thought Strauss would reveal the matter to Gerald Corrigan, the president of the Federal Reserve Bank of New York. Feuerstein thought Gutfreund wanted to think about how the bid should be reported. Strauss thought he and Gutfreund would report to Corrigan, but only after Gutfreund spent more time thinking it through. Gutfreund thought the decision of who would do the reporting had already been made for him and he did not have to speak with Corrigan about the matter. A similar passing of the buck led to a failure to investigate Mozer’s conduct. Meriwether thought that once he had taken

the matter to Strauss he had passed on the responsibility to take action as well. Strauss thought Meriwether, as Mozer's direct supervisor, would take whatever steps were necessary or required as a result of Mozer's disclosure. Feuerstein thought once a report to the government was made, the government would instruct Salomon about how to investigate the matter. Gutfreund thought the other executives would take whatever steps were necessary to properly handle the matter.

In any case, the February auction was neither the beginning nor the end of Mozer's illegal activity. Even after his disclosure of the bid, he submitted two other unauthorized bids in auctions of U.S. Treasury securities, using yet other unwitting Salomon clients as the conduits. The most notable of these bid scams occurred during the May 1991 two-year note auction. Salomon and two customers received 86 percent of the two-year notes. Reports appeared in the press concerning rumors of a possible squeeze in the May two-year issue by Salomon. Gutfreund met with Treasury officials to discuss Salomon's role in the auction, although he didn't bother to mention that the same person who was responsible for the problem with the May auction also had submitted false bids in the February auction.

In early August an internal investigation into Mozer's auction activities concluded that besides the false bids in the February and May auctions, Mozer had also made false bids in the December 1990 four-year Treasury note auction and the February 30-year bond auction. The results of the investigation were reported to Feuerstein on August 6 and to other members of senior management of Salomon, including Gutfreund, Strauss, and Meriwether, the next day.

At that point the issue could no longer be forestalled. Mozer was suspended and Salomon finally pressed forward with notifying regulators. They issued a press release, but it was sketchy on details and bypassed the most egregious aspect—that senior management at Salomon had known of the problems for months without taking any action. This point only came out over the course of the next few days. After aggressive questioning by Charlie Munger, Warren Buffett's close associate, Feuerstein admitted that "one part of the problem has been known since late April." As he was pressed on this point, the chronology of inaction came out. It was this that Buffett later said was "inexplicable and inexcusable." A later press release addressed the delay in shrouded terms, explaining that it was "due to

a lack of sufficient attention to the matter.” This was a little better than the farfetched first draft of the release, which stated that it was due to the “press of other business.”

Though this alone might have been enough to do in Gutfreund, his fate was sealed when he failed to inform the board of yet another letter he had received from the Federal Reserve, a letter that asked for a full reporting within 10 days of the “irregularities, violations, and oversights,” suggesting that these issues, if not satisfactorily resolved, might end the business relationship between Salomon and the Fed. That is, it might end Salomon’s appointment as a primary dealer, which would cut the legs out from under the firm’s fixed income operations.

Buffett learned later that Corrigan expected the letter to be promptly given to Salomon’s directors, whom he believed would just as promptly recognize that top management had to be changed. When the directors didn’t act, Corrigan thought they were being defiant. The truth, of course, was that they were simply in the dark. Buffett did not hear about any Fed letter until later in the week, when he spoke to Corrigan, and even then he assumed the Fed had only sent a request for information. He did not actually see the letter until more than a month later, after he heard Corrigan refer to it in congressional hearings.

In Buffett’s opinion, the Fed’s belief that its letter had been ignored stoked the fury with which the regulators came down on Salomon a few days later. There was no shortage, Buffett said, of “vital matters” that Gutfreund, Strauss, and Feuerstein kept from the directors in the previous months, all the while acting as if things were perfectly normal. But not conveying the Fed letter to the board was, in Buffett’s thinking, “the atom bomb.” Or, as he described in a more earthy manner: “Understandably, the Fed felt at this point that the directors had joined with management in spitting in its face.”⁶

The next day Salomon issued a press release that stated it had “uncovered irregularities and rule violations in connection with its submission of bids in certain auctions of Treasury securities.” The release described several of the violations and stated that Salomon had suspended two managing directors on the government trading desk and two other employees.

That same day Gutfreund and Strauss disclosed for the first time that the firm had known of a false bid in a U.S. Treasury auction since late April. Five days later, Salomon issued a second press release, which pub-

licly disclosed for the first time that Gutfreund, Strauss, and Meriwether had been “informed in late April by one of the suspended managing directors that a single unauthorized bid had been submitted in the February 1991 auction of five-year notes.”

That Sunday, at a special meeting of the board of directors, Gutfreund, Strauss, and Meriwether resigned. Feuerstein resigned his position as chief legal officer later that week.

Gutfreund had spurned a separation offer of \$8 million and went to arbitration asking for \$55 million, hoping that a split-the-difference negotiation would get him what he really hoped for, something around the \$15 million that was his due for his options and pension. Negotiations went nowhere, and the arbitration board apparently decided (its decisions are not public) that his mismanagement, which had almost cost Salomon its existence, did not merit any further compensation. The firm had argued as much, and went even further by saying that he had deceived the Salomon board through his inaction, and thus might have justifiably been fired for cause. With disgruntled shareholders and lawsuits still circling the firm after the Treasury scandal, this decision was the end of the line—there was no way Salomon would be able to pay Gutfreund even if so inclined.

Gutfreund never seemed to appreciate the gravity of his actions and appeared mystified by Salomon’s refusal to settle his claim. At the time, he called the firm’s actions to avoid settling his claims “relentless, baseless, and bizarre.” Ultimately, Salomon paid \$290 million in fines to settle charges related to the scandal. As for Gutfreund, in a few years he had gone from King of Wall Street (as declared by *BusinessWeek*) to an ignominious resignation and having the bulk of his wealth wiped away. Salomon would regain its footing, although in a tamer world; the firm and Wall Street did not miss Gutfreund.

The most significant impact from the scandal ultimately was Meriwether’s resignation, which ushered in the hedge fund boom. Meriwether recruited his star players from the Salomon arbitrage desk to form the Greenwich, Connecticut, hedge fund Long-Term Capital Management (LTCM), nicknamed Salomon North. It wasn’t long before a trend was established, as the move of many of the highest-paid and most talented traders on Wall Street into the hedge fund world started others thinking along the same lines. It’s easy to see why Meriwether’s band of traders caused a gravitational shift on Wall Street. They were incredibly smart and

successful. They had proprietary chic. Even sitting in the middle of the Salomon trading floor, they maintained an aura of celebrity, a clique with its own culture of inside jokes and secret nomenclature to describe the yield curve and the strategies they employed. When they moved off, driving to work in Ferraris, working in shirtsleeves by a harbor off Long Island Sound, and making huge levels of compensation—in the first few years their fund had returns in the 30 to 40 percent range—the sell-side traders started to feel like workaday drones slavishly making markets. No matter that they were pulling down six-figure bonuses. Some of the best followed the path to the more glamorous and more lucrative hedge fund world. And as that occurred, the message to those who remained was that they were also-rans in the trading game.

THE MARCH OF THE LONG/SHORT HEDGE FUNDS

Much of my academic work and early focus on Wall Street was research in the mathematically elegant realm of option theory and related derivatives, including yield curve and mortgage models. My first foray into proprietary trading involved crunching gigabytes of tick data for various financial instruments on an array of workstations that was close to having a supercomputer at my fingertips. Then I moved into risk management, which at the time centered on the transformation of a few financial and economic state variables such as the yield curve of interest rates into signals for relative value and macro trades. But over time I was drawn to the most standard of investments, buying and selling stocks.

In moving to Moore, I moved from the control-oriented risk management of the large sell-side firms to a tactical use of risk: risk as a tool to make money instead of trying not to lose it. On the sell side the challenge was to corral risk within bounds set by the board—and even before that, determining how to define the risk and then measure it across hundreds of trades in a myriad of financial markets. These problems generally did not exist in the hedge fund world, because the board and the traders were one and the same. And they focused on a relatively small set of markets that they knew inside and out. Instead, the issue for hedge funds was in sizing trades and modulating risk to make the most of market opportunities while maintaining the discipline to stay on the sidelines if the opportunity wasn't there. While the main focus at Moore was macro trading, I also

worked with a group of long/short equity managers to develop a suite of tools to help them evaluate their trades, so that they could understand where they derived their edge and where they were wasting time and effort. I realized this product could be of broad value to hedge funds, so, with the go-ahead from Moore, I created a company called Scribe Reports to sell the analysis to other long/short equity hedge funds.

Scribe Reports provides what I call skill assessment to hedge funds. By looking at the history of managers' positions, what stocks were bought and sold, and what was going on in the market at that time, managers can get a sense of whether they are good at shorter-term or longer-term trading, whether the low liquidity trades they put on are worth the risk, and whether they are better at stock selection or market timing. In terms of risk management, rather than simply measuring the portfolio's current risk level—a trivial exercise for equity portfolios—the analysis tries to determine the link between risk and sizing: that is, how good the manager is at modulating risk—pushing leverage and position size up and down—in response to market conditions and trading success.

Scribe Reports also proved valuable for helping make asset allocation and hire/fire decisions for investors, funds of funds, and those within a hedge fund who were overseeing the stable of managers. The problem with hedge funds is that they are deliberately opaque. They don't want anyone to know their positions, lest others trade against them or emulate their strategies. So an investor's information about a hedge fund is typically limited to monthly performance results and occasional PowerPoint presentations from the manager or investor relations types. This obviously is not a lot of information, and it is sometimes funny to see how it gets multiplied in an attempt to get more out of it. For example, I know of one firm that pulls in the monthly performance numbers for a hedge fund and pumps out nearly a hundred different statistics in charts and graphs that look like the result of a workshop on "Fun with Excel." Most hedge funds have less than five years of monthly data, so this turns the notion of statistics on its head. Rather than having statistics capture data and summarize their most important points in a few numbers, you end up with 50 or 60 monthly return numbers being sliced and diced in more ways than the number of raw data points.

In any case, Scribe Reports was a business that was not well suited to marketing while I was at Moore, because Moore was a large and strong

competitor for investor dollars for most of my potential clients. I went from Moore to Ziff Brothers Investments, and while there I continued to pursue this venture. ZBI traded only with internal funds—those of the Ziff family, of publishing fame—in what is termed in the investment world a family office; it was never chasing after other hedge funds' investor dollars. Not long after I joined ZBI I moved from risk management to portfolio management. I redirected the efforts of a small group of PhDs who had been providing quantitative analysis for the traditional portfolio managers toward running an internal portfolio based on quantitative trading models.

While the trading center for Moore was macro strategies, at ZBI the center was equities, and the portfolio I managed was an equity portfolio. So between this and Scribe Reports, I had moved solidly into the world of equity hedge funds, and I found equities to be a very attractive market. There are many state variables that underlie the price of a stock, and with years of data on thousands of stocks there is a statistical soup of observations where relationships can be coaxed out in many dimensions (although this bounty is not necessarily an advantage—for those without sufficient discipline or statistical knowledge, there is enough data to find just about any relationship you want). Relationships might emerge based on common industries, factor exposures to energy, interest rates and other economic variables, a value versus growth bias, or a host of other factors. There are quarterly accounting fundamentals, tick-by-tick trade and limit order book data, analyst views, and raw economic data. And while rich in statistics and information, the relationships can generally be treated as being linear because the noise emanating from so many data sources obscures any higher-order effects.

Contrary to the academic notions I cut my teeth on, any number of surprisingly simple ideas seem to bear fruit with equities. One reason for this is that there are many large asset management firms and pension funds that are slow to make adjustments because of limitations in liquidity or because they have a decision-making bureaucracy; others are busy maintaining appearances, which leads them to skirt out-of-favor stocks and make predictable portfolio “window dressing” adjustments. Then there's human nature: Some market participants have objectives other than creating return on investment. For example, for most security analysts the critical factor is being more right than other analysts: If you think earnings will

be \$2 a share, you just have to be closer to \$2 a share than the other analysts. If everyone else is estimating \$1 a share, you might move your estimate up to \$1.10. There is no need to stick your neck out and risk looking like a fool. So analysts tend to be behind the curve when a company has unexpected trends in earnings. I call this the *Price Is Right* effect, after the game show where all that matters is to be closer to the correct price than the other contestants.

And as I have already discussed, investors on occasion seem to stray from rational decision making. This is true of even experienced professionals. They make decisions differently based on whether they already hold the position. They also make decisions based on arbitrary benchmarks, the best known being what is called the disposition effect, a tendency to hold positions that are underwater in hopes they can eke out some sort of a profit down the road. The other side of the coin from the disposition effect is the tendency to quickly get out of positions that are making money. I believe the reason many investors get out of winning trades quickly, and, for that matter, the reason they hold onto losing trades, is that they measure their success at least in part by their win/loss percentage. That is, they are trading to win rather than to maximize cumulative return. If you are trading to win, the amount of the profit or loss is less important than the number of winners versus losers. The rule to follow if you are trying to maximize your winning percentage is to get out of a profitable trade quickly in order to chalk up one more in the win column and to hold onto the losers in hopes they turn around. After all, if a trade fails to turn around, it is still just one more loss, whether it is down 2 percent or 20 percent.

Investors also harbor irrational behavior that I believe is central to the empirically demonstrated success of momentum strategies. If new information comes into the market that suggests the current market price is too low, prices should adjust rapidly up to the new level. But prices often move up slowly, giving rise to success in buying into a trend. One reason for the slow pace of adjustment is that many investors must move into new positions gradually, either because they have to work through an approval process or because they have a large amount to invest and do not want to disrupt the market. But another reason is that investors have arbitrary time frames in assessing market moves. If a security price goes up during the day, some investors will hold back because they feel that they are late

to the party. But the next day or two, it is as if the benchmark price has reset. The previous run is largely forgotten, and if the prices are stable in the interim these same investors will then go in, further ratcheting up the price. As a result, the price will go up in fits and starts.

After two years I left ZBI and pursued a long-term interest of mine, running a hedge fund as a family enterprise. My two oldest sons were out of college, with degrees in math and computer science, so I had the raw material. I looked at two possible homes for doing this. One—you'd think I would learn—was back in the mines at Citigroup, where a colleague, Tanya Beder, was heading up the hedge fund unit, Tribeca Investments. The other, and my eventual choice, was FrontPoint Partners, an organization established by three émigrés from Morgan Stanley and Tiger. FrontPoint was founded with an appealing and strategically farsighted structure. It provides the hedge fund managers with all the back-office, execution, legal, and marketing infrastructure in return for part of the management fee. But the portfolio managers also can have a partnership stake in the mother ship, so that the managers have their interests aligned. This structure was set up not only to overcome the hired gun mentality of most hedge funds that have separate portfolios, but also to overcome one of the biggest drawbacks of hedge funds and investment firms in general: the difficulty of monetizing their performance.

Although going public or selling out is becoming of increasing interest, even hedge funds with great performance can find it difficult to go into the market and sell their businesses the way other businesses can. A business built to generate revenues of \$100 million might be able to create a market value of a few billion dollars. But if a hedge fund pulls in \$100 million in fees, it will be hard-pressed to find someone who will be willing to buy it for \$1 billion or more. (From an economic standpoint, this may justify hedge fund managers' high fee income. They cannot monetize future earnings as can be done in other businesses, so all the benefits have to be accrued through the yearly draws of income.) As the saying goes, a trader is only as good as his last trade. So a hedge fund will have a low price-earnings multiple, because there is low confidence the earnings are sustainable.

And there is also the risk all the assets—the talent of the portfolio managers—will either walk out the door or lose their incentive to trade if the firm is sold. FrontPoint's business approach sought to overcome these

problems by making the managers de facto shareholders from the outset and by creating a large enough pool of managers, focusing on several dozen markets, to reduce the risks that are inherent in a one-trader shop.

This monetization became a reality far sooner than the founders or portfolio managers at FrontPoint had initially anticipated. In late 2006 FrontPoint Partners was sold to Morgan Stanley, becoming part of the new MSAM, now called Morgan Stanley Investment Management (MSIM). And so I closed the loop, working again in the same firm where my Wall Street career started. I had moved from the growing institutionalization of the investment banking firms into the freewheeling buy-side hedge fund world, only to be reeled back in. But the absorption of hedge funds' talent, talent that largely came from the investment banking world in the first place, is a mark of success. Hedge funds have moved into the mainstream. They are of increasing interest to pension funds and other major institutional investors, and these professional investors demand a scale and degree of risk control that the larger institutions are well positioned to provide. And thus the Wild West of hedge funds will be tamed. With the exception of the very largest hedge funds, which already have sufficient scale to satisfy the institutional investors' demands, hedge funds will increasingly be consolidated and pulled into the larger institutional structures.

CHAPTER 10

COCKROACHES AND HEDGE FUNDS

When I attended graduate school at MIT in the late 1970s, the perfect market paradigm was fast emerging as a framework to analyze the financial markets. It had already revolutionized the study of economics, and for academics the perfect market paradigm promised, for the first time, to provide a rigorous mathematical approach to understanding and interpreting the financial markets. I became enthralled with this promise after taking a course from Bob Merton, one of those rare men who is both a brilliant researcher and a great teacher. I had embarked on my graduate work with the intention of changing the underdeveloped world through developmental economics, but ultimately the elegance of financial economics was more alluring. My dissertation turned to a subject area of mutual interest to finance and economics: the transmission of information through the markets, with Merton as one of my advisers.

The MIT campus at that time contained the preeminent economics department in the country. The cornerstone was Paul Samuelson, one of the first Nobel Prize winners in economics and the man responsible for bringing mathematical rigor and scientific analysis to the discipline. As an undergraduate at the University of Chicago in the early 1930s, with the country mired in the Depression, Samuelson moved from the study of physics to economics. The switch was propitious, as he quickly realized that the tools of calculus and the principles of maximization he learned in his studies of physics could be applied in a similar manner to economics. During his graduate school years at Harvard, he developed a new, mathematically oriented approach to economics. His doctoral dissertation describing his revolutionary approach won the David A. Wells Prize in 1941 as the best doctoral dissertation in economics. His career was off and running.

Samuelson expanded his doctoral dissertation into a work entitled *Foundations of Economic Analysis*, which was finally published after being delayed by World War II, in 1945. He stayed on as a junior fellow at Harvard, but the 1940s were not a hospitable time for Jews in the Ivy League. He failed to receive a permanent appointment at Harvard, so he headed down the Charles River to MIT. There, joined by other Jewish academics, he became the leading light in an economics department that would have an enormous impact on academia, governmental economic policy, and the financial markets. Two who joined on shortly after his appointment were Robert Solow, a Brooklyn-born economist who did his undergraduate work at Harvard, and Franco Modigliani, who fled fascist Italy in 1939. Over the next 40 years, Samuelson, Solow, and Modigliani worked on some of the most important economic issues of their time. Solow's most significant contribution involved the application of ordinary differential equations to understand the determinants of economic growth. Modigliani's work spanned areas of corporate finance and macroeconomics. Like Samuelson, both Solow and Modigliani were awarded the Nobel Prize in economics.

Samuelson and Solow brought a mathematical bent to economics that became its sine qua non and, over time, an end in itself. The math was fun, looked impressive, and provided proof positive of modern economics' scientific heritage. Indeed, the more high-powered the mathematical technique, the greater its allure. Gerard Debreu, a French-born mathematician-turned-economist (and also a Nobel laureate), stated that no other academic field

was as quick as economics to embrace and actively apply the frontier developments of mathematics. As a result, economics became a safe haven for mathematicians and physicists who could not find tenure-track positions in their own increasingly competitive fields.

Despite its benefits, the mathematical focus created limits for economics that have endured for several generations. In classical physics, any number of real-world effects such as friction or air resistance are assumed away to make mathematical analysis more tractable. Perfect vacuums and ideal gases provide a set of simplifying assumptions that allowed for the development of theories of the physical world. Similarly, in the study of economics it is necessary to assume a construct of frictionless markets to build a market theory out of the tools of mathematics.

This assumption of frictionless markets included instantaneous and costless transactions devoid of real-world constraints. Buyers and sellers bought or sold at posted prices, with no associated fees, and their actions had no impact on the market—in the nomenclature of economics, the market participants were atomistic. Moreover, to permit sophisticated spanning arguments and the application of fixed-point theorems from topology, it was assumed there were securities available for every possible contingency; every risk and possible event or state of nature not only was identified, but was also represented by a market security. Economics could be successful only if it predicted the behavior of people, but if real people were brought into the equation—people who don't think like a computer, don't perform mathematical optimization before every decision, and don't (or can't) process all available information before they take action—it would get in the way of the mathematics that filled the pages of the academic journals. So in order to apply the tools of mathematical optimization, the idealized financial world also assumed perfect rationality on the part of the market participants, all of whom make investment decisions with complete information.¹

The resulting structure is the perfect market paradigm. Its key features are readily available information that reaches all market participants simultaneously, rational investors who can instantly and correctly interpret the implications of that information, and liquid markets that can immediately and without cost accept trades based on these implications. The perfect market contains a complete set of financial instruments to allow investors to trade on every possible contingency. Capital is not a constraint to action. There is infinite leverage—investors can borrow without limit.

IMPERFECTIONS IN THE PERFECT PARADIGM

The perfect market paradigm assumes that markets are efficient—that is, that all information is imbedded in the market price. In an efficient market, no trader can make money trading on news; by the time a trader gets the news all the other market participants will have the news as well, and the price will have adjusted to the correct level given that information before any trade can be made.

Take a trip to a Wall Street trading floor and it is easy to see how one might end up making the efficient market assumption. Traders are news junkies. The typical trading floor is peppered with screens showing CNN and CNBC; electronic tickers streaming around the room; a half-dozen or so screens surrounding every trader, flashing red or green for each downtick and uptick; and other news screens displaying only a headline for breaking stories and slightly more information for stories that are a few minutes old. There is a phone bank with direct lines to the brokerage houses; no waiting even for a speed dial—a press of the button and the market maker is on the other end of the line.

To keep everyone in the information flow, the large broker-dealers try to cram all the traders on the same cavernous floor. When important news breaks, a flurry of activity erupts around a trading turret, and a crescendo of noise emanates from the intercoms and squawk boxes that connect the traders on the floor. The sound level—which normally is like background noise in a crowded restaurant—multiplies in intensity. Besides keeping plugged in to all the information washing across the trading floor, traders receive e-mail and instant messages that stream into their in-boxes almost as quickly as news stories roll by on their screens, along with research reports from brokers, as well as commentary and recommendations from specialized research firms. While each e-mail contains information or analysis, many are little more than marketing pitches intended to entice traders to do business. In spite of this apparent frenzy, for any one trader it is a workaday world—answering the phone to execute orders, keeping track of trades on a blotter, and making sure that his net position is within the prescribed risk parameters. But to an observer, the buzz and activity on the floor can take on the feel of a chaotic street scene in downtown Calcutta.

In the perfect market paradigm, all of this is for naught. Markets are efficient, which is to say that they react immediately and appropriately to

all relevant information; when the news comes out, they adjust instantly. Since information coming into a market is by definition unknown and random, and since the market reacts fully and immediately to this new information, market prices move about randomly. From this comes the assertion that the market is a random walk.

Full adherence to the efficient markets hypothesis leaves much of the financial industry in a paradoxical position. It is precisely the activity of the many people trying to track down information to make profitable trades that leads the markets to be efficient. But in the aggregate this leaves investors and traders unable to extract profits for all their trouble. On the margin, investors and traders provide no value, but in the aggregate they are the bastions of the markets.

This leads to the question: What are all these people doing who make a living managing money and trading? By luck, some will outperform the market and some will underperform, but no one can predict who will be a winner and who will be a loser. At best, these professionals are economic deadweight; at worst, they are a net negative because they are a drain on economic resources, as all of their trading does nothing but rack up transaction costs, unnecessarily driving down returns. They are needed to make the markets efficient, but because they do so, they cannot profit from it.

The blind adherence of academics to the efficient markets hypothesis is caricatured by the professor who sees a \$20 bill lying on the sidewalk as he walks down the street with a student. He passes it with barely a second glance, and the student asks, "Why aren't you picking the money up?" The professor smiles, and with a knowing shrug of his shoulders explains, "Well, if there really were a \$20 bill on the sidewalk, someone would have picked it up already." The efficient markets hypothesis states in effect that there are no spare twenties for the taking.

The efficient markets hypothesis also leads to other questions, the first being: Why do we have such voluminous markets with continuous trading and second-by-second reporting of prices? The standard answer of the information-centric perfect market paradigm is that they are there to provide prices "for planning purposes," giving manufacturers and others involved in the real economy up-to-the-minute information so they can make better decisions in resource allocation. Because information moves prices, the logical inference from this view is that information also is the

culprit causing turmoil in the markets. News leads investors to revise their view of the world, and they readjust their holdings accordingly.

IT'S THE LIQUIDITY, STUPID

This focus on information flow was the basis of economic analysis into financial markets. This happened as much because of information's mathematical convenience as it did its correctness. But the information-driven efficient market view, and indeed economists' love affair with mathematical analysis, seems to ignore the facts of how markets actually operate. If information drives market prices, and if the sole role of trading is to provide better resource allocation in economic production, then why do intraday prices bounce around as much as they do? The price of a futures contract or a stock moves around much more than one would rationally expect from new information. What type of information could possibly cause the price of a stock to jump by one tick, then down by one tick, then up three ticks, up another two ticks, second by second, throughout the trading day? And if the objective of markets is to provide information for the production sector, how do we justify the enormous overhead of a continuous market with real-time information? And just how much planning can you do when prices are jumping all over the place, anyway?

These questions bothered me throughout my MIT catechism. Something was clearly missing in the academic view of the world, something that is immediately apparent once you step into the real world of Wall Street. In one news segment on *Saturday Night Live*, the newscaster announces, "And today on the New York Stock Exchange, no shares changed hands. Everyone finally has what they want." The punch line raises a legitimate issue: Why is it that shares are always trading? Why are we never satisfied with what we have?

The principal reason that prices vary, especially in the short term, is liquidity demand. That is, far more than acting as a conveyor of information, the objective of markets is to provide liquidity. Market liquidity is essential to allow assets to be bought and sold quickly and with low transaction costs. And, most importantly, it is in the froth of liquidity that profits are made and that the market demons are spawned.

Implicit in the perfect market paradigm is a degree of liquidity where not only are securities readily bought and sold, but also the individual's act

of buying or selling does not affect the price. The market is composed of so many small players relative to its size that they can execute their trades, extracting liquidity from the market as they demand, without moving the price. But that is where the perfect market vision starts to break down: Not only does the demand for liquidity move prices, but it—and not information—is also the primary driver of prices. As we will see, it is the primary driver of crashes and bubbles as well.

LIQUIDITY IN THREE EASY LESSONS

Liquidity demanders must move the market to meet their needs. Liquidity suppliers try to profit from the moves of the demanders; they seek to meet the liquidity demand—for a price; they will buy or sell on demand, so long as the price is right. They have a view of the market and take a position in the market when the price deviates from where they think the fair price should be. If the deviation is because others in the market need liquidity, then the liquidity providers will supply that liquidity and extract a profit once the price returns to fair value. If, though, the change in price is not transitory, they will face a loss on the trade. To liquidity demanders time is more important than price, while to liquidity suppliers price matters more than time.

The first line of liquidity providers is the market makers and specialists on the exchange floor. The market maker is the intermediary for transactions, moving the price based on the liquidity demander's needs in order to attract the appropriate amount of supply. Market makers trade with a very short horizon, meeting each bid and offer. They get to buy at the bid price and sell at the offer price, so on average they make the bid/offer spread on each trade. But they face the risk that they will be flooded by buys or sells without any countervailing trades, forcing them to either hold inventory as the market moves against them or drop their price and sell at a discount. In fact, the changes in prices are the result of the price adjustments of the market makers when their inventory becomes too lopsided. In some markets where there is no central exchange, such as corporate bonds and mortgage-backed securities, the broker-dealers take on the role of price setters. The trader may already have the bond available in inventory, but more often than not will have to take the other side of the client's position and wait until someone else comes along who wants to take the

bond off the trader's hands. These fixed income instruments do not trade as frequently as exchange-traded securities, so the trader may enlist the services of the salespeople on the trading floor to call around and see if any clients are interested in the position.

The next line of liquidity providers includes hedge funds and other speculators, and then, working on a longer time frame, investors. The statistical arbitrage trader falls into this group. Traders and investors come to the market with a wide range of strategies. Some are value oriented, looking for deviations between market prices and their perception of the fundamental worth of the stock. Others rely on price signals, other so-called technical information, or just noise. These players attempt to profit by predicting the future course of prices. They do not think of themselves as being liquidity providers. But without realizing it, they often respond to the signals of liquidity demanders. For example, if a pension fund has a persistent need to sell a particular stock, the stock's price will drop as it tries to entice a buyer, and the stock will end up deviating sufficiently from perceived value to attract some value investors. It does not really matter if the investors have the right view of value. Fundamental investors, like all other investors, are all over the map, and there will be some out there for whom a given price drop will move things below their threshold. In terms of their responses to market prices, value investors are just the flip side of the trend followers. Whatever the process by which they measure value, they will sell into a rising price, because it is moving up relative to their value target, and buy when the price declines. The trend follower exacerbates liquidity demand. The liquidity demander moves the price up to elicit a seller, and the trend follower, rather than selling, adds to the fire by buying even more. From a liquidity perspective, the strategy of trend following makes sense if there is a good chance the initial surge in prices will be followed by yet more liquidity demand.

Liquidity providers serve a valuable economic function. Their business is to keep capital readily available for investment and to apply their expertise in risk management and market judgment. They look for instances of a differential between price and value, and as they trade to exploit that differential, they provide liquidity to the market. In short, they take risk, use their talents, and absorb the opportunity cost of maintaining ready capital. For this, they receive an economic return.

The argument that hedge funds provide liquidity to the market and that providing liquidity is a service that merits compensation may seem es-

oteric and a bit of a stretch. But we do see return differentials between securities that are highly liquid and securities that are similar in almost all attributes except liquidity. One of the simplest examples is the liquidity premium between on-the-run and off-the-run bonds that is exploited by relative value traders.

Liquidity also has a social value: It provides economic freedom. It allows wealth to be accessed and used to take new opportunities. To appreciate its importance, we can look at a world where liquidity of wealth was absent: medieval England.

Primogeniture and the Role of Liquidity

The economic way of life in medieval England was framed by two economic realities: Virtually all wealth was in the form of land, and the land could not be sold. This is the very definition of an illiquid market.

Wealth was held in land—even up to the seventeenth century it was the universal outlet for savings in England—primarily because there were not many alternative investments. As late as the sixteenth century, more than 80 percent of production was based on agriculture. What limited industry there was bore little resemblance to the Industrial Age that would follow centuries later. Some towns had specialized industry—brewing, salt making, iron working, paper mills—but this was still not characteristic of the economy as a whole. And because land was the preponderant store of wealth, it was also the source of social stature and political power. A large landed estate gave its owner great local influence in controlling elections and sharing in patronage and opened the door for him to join the gentry.

As early as the time of the Norman Conquest in 1066, land could not be sold or even used as collateral for a loan because feudal lords exchanged it for a knight's military service. A knight could no more transfer his land than he could pass on his military obligation. This became one of the most characteristic features of feudal England, though initially it was a new and revolutionary concept.

By the late thirteenth century, this limitation on the right to transfer land carried the weight of law because of another import into England: primogeniture, which restricted land transfers. Under the Statute De Donis in 1285, all landholders had a right to their land only for the course of their lives, after which the deed was transferred according to the rules of

primogeniture, which meant it generally passed to the oldest son. A will or testament could distribute personal goods, but it meant nothing when it came to the disposition of the land. Even in the case of treason, the land would be forfeited only for the generation in question and then would be reinstated to the next, based on primogeniture.

The objective of primogeniture was to prevent the dilution of wealth, in order "to see the holding that he [the landholder] has maintained against the world for a lifetime remain essentially intact to provide a material basis for the perpetuation of the family line." Support for primogeniture was further buttressed by England's disdain for the hodgepodge of land ownership in France, where dividing up the land was the norm. In one French community some 2,000 acres were divided into more than 5,000 parcels belonging to 170 different owners, and a particular walnut tree was asserted to have 32 owners. By contrast, whatever its downside, England at least had a stable base; primogeniture and its related perpetuities assured that the English family estate "neither waxes nor wanes in the course of generations."

The implications of primogeniture were most evident in the social organization of the countryside. At the top was the dominant resident landowner, the squire. Following him was the clergyman, usually appointed by the squire and often his relative. The peasants who worked the land were tenants-at-will and often were also the squire's debtors. The tradespeople and shopkeepers in the village rented from the squire, and the laborers lived in his cottages. They would work on his farm in the summer and do other jobs for him in the winter. The villagers looked to the squire for advice in worldly matters just as they did to the clergyman in spiritual matters, and his word carried the weight of law. He was also the benefactor for the school and the local charities. All in all, his land ownership vested the squire with authority that was just short of sovereignty.

The farmers, tradespeople, and laborers also transferred their positions from one generation to the next based on the rules of primogeniture. This came about initially out of convenience for the squire, because it seemed to be the easiest way of assuring a continuing supply of labor and skills from one generation to the next. Over time, convenience turned into a custom, then finally carried the force of law.

Although England would later serve as the cradle for the Industrial Revolution, for centuries the illiquidity of wealth perpetuated the stagna-

tion—or at least the lack of development—of its society. With wealth locked in the land, the medieval man, whether squire or serf, had no freedom. Both nobles and peasants were bound to their land or their trade. The medieval man had no sense of progress; indeed, the world remained unchanged from one lifespan to the next. Just as the squire had inherited his land, so the farmer inherited his tenancy and the villager his trade. His son and later his grandson would live in the same house and work in the same field or shop. The same picture formed throughout medieval Europe. Other than the changing of the seasons and the marking of religious holidays, the course of time had little meaning. Life passed with intergenerational anonymity. There was no path for venturing out and taking the risks of new opportunities, and there were no opportunities in the village structure where wealth could be applied even if it were emancipated. As a result, few would ever find cause to travel beyond the confines of their own villages. Indeed, in many areas of Europe, to travel any distance—as might be required in the event of war—was to risk never finding the way home.²

From the middle of the fifteenth century, objections started to increase against the massing of so much land in so few hands. Not surprisingly, primogeniture created discord within families, but it also stifled economic incentive, caused turmoil among tenants, and became the source of fraud and worse. This was chronicled in a famous passage from William Blackstone's *Commentaries on the Laws of England* (Book II, Chapter 7):

“Children grew disobedient when they knew they could not be set aside; farmers were ousted of their lands made by tenants-in-tail; . . . creditors were defrauded of their debts; . . . innumerable latent entails [contracts prohibiting transfer] were produced to deprive purchasers of the lands they had fairly bought; . . . treasons were encouraged, as estates-tail were not liable to forfeiture longer than for the tenant's life.” Other writers declared that perpetuities “fight against God, by effecting a stability which human providence can never attain to” and that “nothing is more unwise than to attempt to bind posterity with parchment.”³

The cost of illiquidity became increasingly apparent on a practical level as new avenues for wealth and investment opened up. The development of overseas commerce and the increasing involvement of leading merchants in the lucrative business of lending money to the government expanded investment opportunities. Finally, there were prospects for building fortunes apart from land ownership. The landed gentry had an

incentive to extract the wealth from their land to pursue other opportunities. Although the sixteenth century continued to see the rise of an aristocracy and a portion of the landed gentry still obsessed with creating an entrenched position in society, the tide had largely turned. Their interest in the status quo was far outweighed by others who demanded liquidity for their land-locked wealth. Various artifices became employed to circumvent the rules of primogeniture.⁴ Demographics also weakened the bonds of primogeniture because many family lines ended with no male heir.⁵

The newfound liquidity in the land progressed during the Tudor and early Stuart reigns, resulting in the rapid growth and independence of the English gentry and their servants. But like the deprivations that followed the economic freedom in post-communist Europe, this freedom did not come without a cost. The ability to sell land led to an increase in poverty. No longer could the medieval villager count on being looked after by the local landowner as an asset and responsibility, as had been the case for generations. As the squire-villager relationship eroded in Tudor England and as tenures of land were freed up, the tenant was left relatively unprotected to face a rapidly changing world. The result was an increase in the prospect for economic failure and poverty, not just for the tenant but for the landowner as well.

Now that land could be used as collateral, it opened up new possibilities for borrowing and lending. Not surprisingly, landowners began to incur debt with increasing frequency as the sixteenth century wore on. An embryonic capital market developed in London, and by the seventeenth century in other cities as well. The interest rates were substantial—they could rise to 10 percent or more—and the loans were for short periods. The results were perhaps predictable: Because land was used as collateral, numerous estates were put on the market to fulfill the loan obligations. Thus, with the end of perpetuities, a single break in the fortunes of the landholder could scatter the accumulation of wealth that had been trapped for generations. This was a liability to some but a benefit for the populace and economy overall as the wealth slowly diffused into a broader population and, more critically, as the wealth found its way into financing the emerging opportunities of the Industrial Revolution. Indeed, I would go so far as to say that this emergence of liquid wealth lies at the genesis of the Industrial Revolution.

The lesson of medieval England is clear: Liquidity allows capital to migrate to greater opportunities. Extending the point, liquidity suppliers are

providing an economic benefit and, as such, should expect to be compensated for their services. In the world of modern trading, the liquidity supplier rightly should be paid not just because of the risk entailed in accommodating the seller, but also because of the opportunity cost involved in making cash available for speculative positions and in taking time to monitor the market and make trading decisions. The compensation should also be a function of the volatility in the market: The more volatile, the higher the probability in any time period that prices will run away from the liquidity suppliers. In addition, compensation should be a function of the liquidity of the market; the less liquid, the longer the position must be held, and the greater the risk.

Looked at in this manner, hedge funds and other speculative traders provide an economic service similar to that of retailers who try to anticipate market demand and stockpile accordingly. Consider a clothing retailer who is approaching the summer season and must decide which types of swimsuits to order. There are indications of which styles will be of interest, the clothing equivalent of the fundamental information that tells investors what earnings and economic numbers will look like. But as the season begins, the retailer can gauge more directly the demand for different styles and colors, and keep abreast of emerging fads simply on the basis of customer demand. The retailer is a liquidity supplier, and the customers who must have the “in” swimwear are the liquidity demanders. The retailer will profit from her astuteness by being able to sell more suits at full price while other store owners, having missed the sweet spot, will either mark down their less popular inventory (thus becoming liquidity demanders themselves) or pay a higher price to get the in-demand swimwear from wholesalers at the last minute. In getting ahead of the curve, anticipating demand and putting in an early order to stockpile the swimsuits, the retailer is not just making a profit; she has made the economy more efficient. The desires of the customers will be filled. The customers will not be standing in line at six in the morning or paying above-list price to try to get the swimwear they want; their demand will be filled efficiently and at a price that is fair.

Though vilified and demonized by many, hedge funds and other speculative traders are not gamblers or financial parasites. In the aggregate, by supplying capital to hold risky securities, traders and hedge funds serve to reduce market volatility and improve prices for both buyers and sellers.

For the investor who must liquidate because of a margin call, or the bank that finds itself hitting risk limits that force it to sell in a declining market, the cumulative impact of the hedge funds and speculators can be the difference between staying in business and facing default. And, as can be appreciated against the backdrop of the crises discussed earlier, this impact can extend to the overall market: The liquidity supply afforded by hedge funds and speculators will often make the difference between a market closing down only a few percentage points on the day and the market sliding abruptly into a crisis.

TOO MUCH INFORMATION

Markets react nearly instantaneously to up-to-the-second news about world and economic events. At the same time, every tick in price and the size of every trade is broadcast throughout the world. The companies we invest in are subject to accounting systems and disclosure rules to ensure that investors are provided with accurate information to make their decisions. Despite this constant free flow of information, however, we are caught unawares as one crisis after another devastates the financial markets.

The temptation is to think that information and openness will enable us to control risk and eliminate market crisis. To that end, many argue that we need even more disclosure, and they offer a new solution: position transparency. The argument is that if the market—or perhaps only a subset of the market such as the regulator or the banks that help finance positions—knew the nature and the size of positions that traders are holding, we would be better able to recognize emerging problems and prevent them from escalating into market crises.

As plausible as this position is on the surface, I believe it is wrong. I suspect our problem is not too little information; rather, we may be at the point where we have too much.

In the rigorous world of mathematics or the path of deterministic physical systems, we find that, if we dig deep enough, every analysis rests on assumptions. We cannot necessarily know or verify these assumptions, but they must be accepted as true in order to establish a foundation for our inquiry.

In other words, the progress of knowledge is framed not only by what we know, but also by gaining a better understanding of what we cannot

know. Such is also the case in the mundane world of financial markets. There are some things we cannot know because the very process of learning those things will generate feedback and change the markets in unpredictable ways. Practically speaking, those who seek to make markets transparent might not fully appreciate that traders will know that their positions are being scrutinized and may alter their behavior accordingly. On one level, more information may be ferreted out of the market, but the net effect may be less understanding of the market dynamics because little may be known about the implications of that newfound information.

Even worse, the whole exercise may end up making the markets less stable. If others can readily discern a trader's positions or trading interests, that trader will be less willing to provide liquidity, especially in times of crisis. The trader will be worried about being, in the lingo of the market, "picked off." Other traders will sense his distress and need to liquidate, so they will trade against him; consequently, prices will have to move further to entice traders to take on positions. Our knowledge of the speculative trader's position very well may increase his risk; as a result, we will not only fail in the quest to understand the market in more detail, but the process of trying to increase our knowledge may actually increase the risk of crisis.

The Limits of Knowledge

In 1901 Bertrand Russell, the British logician and philosopher, began an intense 10-year study that would culminate with the completion of a three-volume, 1,800-page tome of nearly impenetrable mathematics called *Principia Mathematica*. This study, for which he enlisted mathematician Alfred North Whitehead as a co-laborer, aimed to demonstrate that "all pure mathematics follows from purely logical premises and uses only concepts defined in logical terms." Its goal was to provide a formalized logic for all mathematics, to develop the full structure of mathematics where every premise could be proved from a clear set of initial axioms.

Russell observed of the dense and demanding work, "I used to know of only six people who had read the later parts of the book. Three of those were Poles, subsequently (I believe) liquidated by Hitler. The other three were Texans, subsequently successfully assimilated." The complex mathematical symbols of the manuscript required it to be written by hand, and its sheer size—when it was finally ready for the publisher, Russell had to hire

a panel truck to send it off—made it impossible to copy. Russell recounted that “every time that I went out for a walk I used to be afraid that the house would catch fire and the manuscript get burnt up.” For the 20 man-years of work—including, for the last three years, 12-hours-a-day bouts of writing once the mathematics had been worked out—the Cambridge University Press estimated it would face a loss publishing the book. (Today the three volumes cost a hefty \$600.) After the Royal Society chipped in a grant to cover some of the cost, Russell and Whitehead were left shelling out £50 each, thus earning –£5 a year apiece for each of their 10 years’ work.⁶

Momentous though it was, the greatest achievement of *Principia Mathematica* was realized two decades after its completion when it provided the fodder for the metamathematical enterprises of an Austrian, Kurt Godel. Although Godel did face the risk of being liquidated by Hitler (therefore fleeing to the Institute of Advanced Studies at Princeton), he was neither a Pole nor a Texan. In 1931, he wrote a treatise entitled “On Formally Undecidable Propositions of *Principia Mathematica* and Related Systems,” which demonstrated that the goal Russell and Whitehead had so single-mindedly pursued was unattainable. (This work, one of the keystones of mathematics and logic, was, incredibly enough, written as part of Godel’s qualifying dissertation for entrance into the teaching profession.)

The flavor of Godel’s basic argument can be captured in the contradictions contained in a schoolboy’s brainteaser. A sheet of paper has the words “The statement on the other side of this paper is true” written on one side and “The statement on the other side of this paper is false” on the reverse. The conflict isn’t resolvable. Or, even more trivially, a statement like; “This statement is unprovable.” You cannot prove the statement is true, because doing so would contradict it. If you prove the statement is false, then that means its converse is true—it is provable—which again is a contradiction.

The key point of contradiction for these two examples is that they are self-referential. This same sort of self-referentiality is the keystone of Godel’s proof, where he uses statements that imbed other statements within them. This problem did not totally escape Russell and Whitehead. By the end of 1901, Russell had completed the first round of writing *Principia Mathematica* and thought he was in the homestretch, but was increasingly beset by these sorts of apparently simple-minded contradictions

falling in the path of his goal. He wrote that “it seemed unworthy of a grown man to spend his time on such trivialities, but . . . trivial or not, the matter was a challenge.” Attempts to address the challenge extended the development of *Principia Mathematica* by nearly a decade.

Yet Russell and Whitehead had, after all that effort, missed the central point. Like granite outcroppings piercing through a bed of moss, these apparently trivial contradictions were rooted in the core of mathematics and logic, and were only the most readily manifest examples of a limit to our ability to structure formal mathematical systems.

Just four years before Godel had defined the limits of our ability to conquer the intellectual world of mathematics and logic with the publication of his Undecidability Theorem, the German physicist Werner Heisenberg’s celebrated Uncertainty Principle had delineated the limits of inquiry into the physical world, thereby undoing the efforts of another celebrated intellect, the great mathematician Pierre-Simon Laplace. In the early 1800s Laplace had worked extensively to demonstrate the purely mechanical and predictable nature of planetary motion. He later extended this theory to the interaction of molecules. In the Laplacean view, molecules are just as subject to the laws of physical mechanics as the planets are. In theory, if we knew the position and velocity of each molecule, we could trace its path as it interacted with other molecules, and trace the course of the physical universe at the most fundamental level. Laplace envisioned a world of ever more precise prediction, where the laws of physical mechanics would be able to forecast nature in increasing detail and ever further into the future, a world where “the phenomena of nature can be reduced in the last analysis to actions at a distance between molecule and molecule.”

What Godel did to the work of Russell and Whitehead, Heisenberg did to Laplace’s concept of causality. The Uncertainty Principle, though broadly applied and draped in metaphysical context, is a well-defined and elegantly simple statement of physical reality—namely, the combined accuracy of a measurement of an electron’s location and its momentum cannot vary far from a fixed value. The more precise the measure of the electron’s location, the less accurate the measure of its momentum; the more precisely one measures its momentum, the less exact will be the measurement of its location. The reason for this, viewed from the standpoint of classical physics, is that accurately measuring the position of an electron requires illuminating the electron with light of a very short wavelength. But the

shorter the wavelength the greater the amount of energy that hits the electron, and the greater the energy hitting the electron the greater the impact on its velocity.

What is true in the subatomic sphere ends up being true—though with rapidly diminishing significance—for the macroscopic. Nothing can be measured with complete precision as to both location and velocity because the act of measuring alters the physical properties. The idea that if we know the present we can calculate the future was proven invalid—not because of a shortcoming in our knowledge of mechanics, but because the premise that we can perfectly know the present was proven wrong. These limits to measurement imply limits to prediction. After all, if we cannot know even the present with complete certainty, we cannot unfailingly predict the future. It was with this in mind that Heisenberg, ecstatic about his yet-to-be-published paper, exclaimed, “I think I have refuted the law of causality.”

The epistemological extrapolation of Heisenberg’s work was that the root of the problem was man—or, more precisely, man’s examination of nature, which inevitably impacts the natural phenomena under examination so that the phenomena cannot be objectively understood. Heisenberg’s principle was not something that was inherent in nature; it came from man’s examination of nature, from man becoming part of the experiment. (So in a way the Uncertainty Principle, like Godel’s Undecidability Proposition, rested on self-referentiality.) While it did not directly refute Albert Einstein’s assertion against the statistical nature of the predictions of quantum mechanics that “God does not play dice with the universe,” it did show that if there were a law of causality in nature, no one but God would ever be able to apply it. The implications of Heisenberg’s Uncertainty Principle were recognized immediately, making him famous. And it became a simple metaphor reaching beyond quantum mechanics to the broader world.

This metaphor extends neatly into the world of financial markets. In the purely mechanistic universe of classical physics, we could apply Newtonian laws to project the future course of nature, if only we knew the location and velocity of every particle. In the world of finance, the elementary particles are the financial assets. In a purely mechanistic financial world, if we knew the position each investor has in each asset and the ability and willingness of liquidity providers to take on those assets in the event of a forced liquidation, we would be able to understand the market’s vulnerability. We

would have an early-warning system for crises. We would know which firms are subject to a liquidity cycle, and which events might trigger that cycle. We would know which markets are being overrun by speculative traders, and thereby anticipate tactical correlations and shifts in the financial habitat. The randomness of nature and economic cycles might remain beyond our grasp, but the primary cause of market crisis, and the part of market crisis that is of our own making, would be firmly in hand.

As stated earlier, the first step toward the Laplacean goal of complete knowledge is the advocacy by certain financial market regulators to increase the transparency of positions. Politically, that would be a difficult sell—as would any kind of increase in regulatory control. Practically, it wouldn't work. Just as the atomic world turned out to be more complex than Laplace conceived, the financial world may be similarly complex and not reducible to a simple causality.

The problems with position disclosure are many. Some financial instruments are complex and difficult to price, so it is impossible to measure precisely the risk exposure. Similarly, in hedge positions a slight error in the transmission of one part, or asynchronous pricing of the various legs of the strategy, will grossly misstate the total exposure. Indeed, the problems and inaccuracies in using position information to assess risk are exemplified by the fact that major investment banking firms choose to use summary statistics rather than position-by-position analysis for their firmwide risk management despite having enormous resources and computational power at their disposal.

Perhaps more importantly, position transparency also has implications for the efficient functioning of the financial markets beyond the practical problems involved in its implementation. The problems in the examination of elementary particles in the financial world are the same as in the physical world: Beyond the inherent randomness and complexity of the systems, there are simply limits to what we can know. To say that we do not know something is as much a challenge as it is a statement of the state of our knowledge. If we do not know something, that presumes that either it is not worth knowing or it is something that will be studied and eventually revealed. It is the hubris of man that all things are discoverable. But for all the progress that has been made, perhaps even more exciting than the rolling back of the boundaries of our knowledge is the identification of realms that can never be explored. A sign in Einstein's Princeton office

read, “Not everything that counts can be counted, and not everything that can be counted counts.”

The behavioral analogue to the Uncertainty Principle is obvious. There are many psychological inhibitions that lead people to behave differently when they are observed than when they are not. For traders it is a simple matter of dollars and cents that will lead them to behave differently when their trades are open to scrutiny. Beneficial though it may be for the liquidity demander and the investor, for the liquidity supplier transparency is bad. The liquidity supplier does not intend to hold the position for a long time, like the typical liquidity demander might. Like a market maker, the liquidity supplier will come back to the market to sell off the position—ideally when there is another investor who needs liquidity on the other side of the market. If other traders know the liquidity supplier’s positions, they will logically infer that there is a good likelihood these positions shortly will be put into the market. The other traders will be loath to be the first ones on the other side of these trades, or will demand more of a price concession if they do trade, knowing the overhang that remains in the market.

This means that increased transparency will reduce the amount of liquidity provided for any given change in prices. This is by no means a hypothetical argument. Frequently, even in the most liquid markets, broker-dealer market makers (liquidity providers) use brokers to enter their market bids rather than entering the market directly in order to preserve their anonymity.

The more information we extract to divine the behavior of traders and the resulting implications for the markets, the more the traders will alter their behavior. The paradox is that to understand and anticipate market crises, we must know positions, but knowing and acting on positions will itself generate a feedback into the market. This feedback often will reduce liquidity, making our observations less valuable and possibly contributing to a market crisis. Or, in rare instances, the observer/feedback loop could be manipulated to amass fortunes.

One might argue that the physical limits of knowledge asserted by Heisenberg’s Uncertainty Principle are critical for subatomic physics, but perhaps they are really just a curiosity for those dwelling in the macroscopic realm of the financial markets. We cannot measure an electron precisely, but certainly we still can “kind of know” the present, and if so, then

we should be able to “pretty much” predict the future. Causality might be approximate, but if we can get it right to within a few wavelengths of light, that still ought to do the trick. The mathematical system may be demonstrably incomplete, and the world might not be pinned down on the fringes, but for all practical purposes the world can be known.

Unfortunately, while “almost” might work for horseshoes and hand grenades, 30 years after Godel and Heisenberg yet a third limitation of our knowledge was in the wings, a limitation that would close the door on any attempt to block out the implications of microscopic uncertainty on predictability in our macroscopic world. Based on observations made by Edward Lorenz in the early 1960s and popularized by the so-called butterfly effect—the fanciful notion that the beating wings of a butterfly could change the predictions of an otherwise perfect weather forecasting system—this limitation arises because in some important cases immeasurably small errors can compound over time to limit prediction in the larger scale. Half a century after the limits of measurement and thus of physical knowledge were demonstrated by Heisenberg in the world of quantum mechanics, Lorenz piled on a result that showed how microscopic errors could propagate to have a stultifying impact in nonlinear dynamic systems. This limitation could come into the forefront only with the dawning of the computer age, because it is manifested in the subtle errors of computational accuracy.

The essence of the butterfly effect is that small perturbations can have large repercussions in massive, random forces such as weather. Edward Lorenz was a professor of meteorology at MIT, and in 1961 he was testing and tweaking a model of weather dynamics on a rudimentary vacuum-tube computer. The program was based on a small system of simultaneous equations, but seemed to provide an inkling into the variability of weather patterns. At one point in his work, Lorenz decided to examine in more detail one of the solutions he had generated. To save time, rather than starting the run over from the beginning, he picked some intermediate conditions that had been printed out by the computer and used those as the new starting point. The values he typed in were the same as the values held in the original simulation at that point, so the results the simulation generated from that point forward should have been the same as in the original; after all, the computer was doing exactly the same operations. What he found was that as the simulated weather pattern progressed, the

results of the new run diverged, first very slightly and then more and more markedly, from those of the first run. After a point, the new path followed a course that appeared totally unrelated to the original one, even though they had started at the same place.

Lorenz at first thought there was a computer glitch, but as he investigated further, he discovered the basis of a limit to knowledge that rivaled that of Heisenberg and Godel. The problem was that the numbers he had used to restart the simulation had been reentered based on his printout from the earlier run, and the printout rounded the values to three decimal places while the computer carried the values to six decimal places. This rounding, clearly insignificant at first, promulgated a slight error in the next-round results, and this error grew with each new iteration of the program as it moved the simulation of the weather forward in time. The error doubled every four simulated days, so that after a few months the solutions were going their own separate ways. The slightest of changes in the initial conditions had traced out a wholly different pattern of weather.

Intrigued by his chance observation, Lorenz wrote an article entitled "Deterministic Nonperiodic Flow," which stated that "nonperiodic solutions are ordinarily unstable with respect to small modifications, so that slightly differing initial states can evolve into considerably different states."⁷ Translation: Long-range weather forecasting is worthless. For his application in the narrow scientific discipline of weather prediction, this meant that no matter how precise the starting measurements of weather conditions, there was a limit after which the residual imprecision would lead to unpredictable results, so that "long-range forecasting of specific weather conditions would be impossible." And since this occurred in a very simple laboratory model of weather dynamics, it could only be worse in the more complex equations that would be needed to properly reflect the weather. Lorenz discovered the principle that would emerge over time into the field of chaos theory, where a deterministic system generated with simple nonlinear dynamics unravels into an unrepeated and apparently random path.

The simplicity of the dynamic system Lorenz had used suggests a far-reaching result: Because we cannot measure without some error (harking back to Heisenberg), for many dynamic systems our forecast errors will grow to the point that even an approximation will be out of our hands. We can run a purely mechanistic system that is designed with well-defined and

apparently well-behaved equations, and it will move over time in ways that cannot be predicted and, indeed, that appear to be random.

Lorenz's observation has led to a new approach to modeling physical processes. In 1993, shortly after I had taken on the role of risk manager at Morgan Stanley, I traveled to New Mexico to attend a conference on economics and finance at the Santa Fe Institute, a vibrant multidisciplinary think tank, an "institute without walls" with no permanent faculty, which encourages collaboration between scholars and researchers in many fields. The principal conceptual thread running through the Institute's research asks how apparently simple systems, like that discovered by Lorenz, can produce rich and complex results. Its method of analysis in some respects runs in the opposite direction of the usual path of scientific inquiry. Rather than taking the complexity of the world and distilling simplifying truths from it, the Santa Fe Institute builds a virtual world governed by simple equations that when unleashed explode into results that generate unexpected levels of complexity.

The main exercise of the conference I attended was to create artificial markets with traders and investors who followed simple and reasonable rules of behavior and to see what would happen. Some of the traders built into the model were trend followers, others bought or sold based on the difference between the market price and perceived value, and yet others traded at random times in response to liquidity needs. The simulations then printed out the paths of prices for the various market instruments. Qualitatively, these paths displayed all the richness and variation we observe in actual markets, replete with occasional bubbles and crashes. The exercises did not produce positive results for predicting or explaining market behavior, but they did illustrate that it is not hard to create a market that looks on the surface an awful lot like a real one, and to do so with actors who are following very simple rules. The mantra is that simple systems can give rise to complex, even unpredictable dynamics, an interesting converse to the point that much of the complexity of our world can—with suitable assumptions—be made to appear simple, summarized with concise physical laws and equations.

The systems explored by Lorenz were deterministic. They were governed definitively and exclusively by a set of equations where the value in every period could be unambiguously and precisely determined based on the values of the previous period. And the systems were not very complex.

By contrast, whatever the set of equations are that might be divined to govern the financial world, they are not simple and, furthermore, they are not deterministic. There are random shocks from political and economic events and from the shifting preferences and attitudes of the participants. If we cannot hope to know the course of the deterministic systems like fluid mechanics, then no level of detail will allow us to forecast the long-term course of the financial world, buffeted as it is by the vagaries of the economy and the whims of psychology.

PRIMAL RISK AND THE LIMITS OF KNOWLEDGE

I run downstairs to grab some breakfast. There is white and rye bread and some bagels (sesame seed, poppy seed, and plain) in a cupboard, butter and a couple of jars of jam in the fridge. Or I can grab one of four boxes of cereal. When I am done eating, I head back upstairs to figure out what to wear. Then it is off to work. I can take the bus, the subway, or a taxi, or in nice weather just walk. I live on 85th Street and Riverside Drive on the west side of Manhattan and my office is a mile and a half to the southeast on 53rd and Sixth, so I can stroll along a number of scenic paths through Central Park or Riverside Park or be more efficient by taking advantage of the fact that for part of the way Broadway cuts toward my office on a southeast diagonal.

Everyone runs through the same sort of tasks every morning, with more or less the same set of choices. We don't really think about it much—we just go through the day employing coarse rules of thumb and heuristics—but if we were absolutely rational, we would leave no possible action unexplored. As a rational man—at least the economics version of the rational man—I should enumerate all combinations of breakfasts, all possible sets of clothes, and all paths to work, evaluate each one in turn given my current set of preferences, grade them relative to the others, and then finally arrive at the one that I find to be the best. In economics parlance, I should optimize my utility. I could do this only by looking at all the choices and ranking them one against another based on my comparative preferences.

This process could take a long time. For example, for my trip to the office, do I want to save money and get some exercise by walking, or do I take the subway? If I walk, do I take a path to maximize safety, to take in the scenery, to minimize distance, or to bypass congestion? The streets

from my apartment to work more or less comprise a 33 by 6 grid, and there are a lot of ways to work down through that maze of streets.

No matter how we make our decisions, it is not as rigorous or as complete as the mathematical world of utility optimization demands. We go about our lives leaving information by the wayside and possibilities unconsidered. We are satisfied to do the job only partway. And the examples here don't even take into account the real complication of our lives: uncertainty. For most choices we make, the outcome is measured in probabilities rather than certain results.

Though it forms the foundation of the economic theory of behavior, economists are often apologetic about the burden that they place on us by invoking the "rational man," arguing that this notion should be taken with a grain of salt. They might concede that we do not make decisions according to the rational man model, which involves a process of understanding the nature of the problem, defining clear preferences in comparing possible alternatives, and formulating complex optimizations in order to arrive at the final course of action. Rather than assert that we make decisions as would the rational man, they assert that we act *as if* we follow this procedure, even though we do not do so explicitly. In other words, our outcome, however derived, ends up being one that could have come about as the result of this formal "rational man" approach.

Even when I was a graduate student, the disconnect between how people behave and how they should behave left me perplexed. But that was in the heyday of efficient markets and the idea that investors are utterly rational, so it was difficult to present an argument for behavior outside that framework. My sense was that the culprit leading to our apparent nonresponsiveness to optimality was a lack of knowledge, a realization that the world could change in ways we could not anticipate or model. In the approach of the day, taking this route just meant adding a probability distribution into the equation, so that the ever-rational agent was now optimizing *expected* utility. But to me the point was to consider the effect of uncertainty that could not be labeled, much less represented by a probability of outcomes. It was hard to get anyone in the academic circles in economics to concede such a metaphysical level of uncertainty, so I moved the argument into the biological realm.

The assertion that there are some risks that cannot be known can be more readily grasped by thinking about unintelligent life forms. While

some may not be ready to admit it, there is a philosophical issue in discussing that which we are assuming we cannot even know to the point of discussing. In any event, we are more comfortable with the notion that lower life forms aren't as smart as we are. We are willing to concede that nature has surprises that are wholly unanticipated by our nonhuman cohabitants. A disease that destroys a once-abundant food source, the introduction of chemicals into a pristine environment, and the eruption of a volcano in a formerly stable geological setting are examples of events that could not be anticipated by lower life forms even in probabilistic terms and therefore could not be explicitly considered in rules of behavior. They are nature's equivalent to the unforeseeable risks that I argue also exist for us. Because of their foothold in the biological, I call these primal risks.

COCKROACHES AND THE BENEFITS OF COARSE BEHAVIOR

While the focus of the academic community continued to be on making economics look more and more like physics, I felt biology was a better place to look for a frame of reference. I developed a biological analogue for economic behavior, writing a paper in collaboration with a fellow MIT graduate student, Joe Langsam, entitled "On the Optimality of Coarse Behavior Rules," which was the lead article in a 1985 issue of the *Journal of Theoretical Biology*.⁸ (Langsam earned a PhD in economics from MIT, but in the course of his graduate studies discovered a love of mathematics and added a PhD in that discipline as well. Shortly after coming to Morgan Stanley I recruited him to join me in fixed income research, where his double PhD earned him the sobriquet of "Doctor-Doctor." Brilliant and with an unusually intuitive sense of both mathematics and finance, he now heads that department.)

The best measure of adaptation to unanticipated risks in the biological setting is the length of time a species has survived. One that has survived for hundreds of millions of years can be considered, de facto, to have a better strategy for dealing with unanticipated risks than one that has survived for a short time. In contrast, a species that is prolific and successful during a short time period but then dies out after an unanticipated event may be thought of as having a good mechanism for coping with the known risks of one environment but not for dealing with unforeseeable changes.

By this measure, the lowly cockroach is a prime case through which to study risk management. Because the cockroach has survived through many unforeseeable changes—jungles turning to deserts, flatland giving way to urban habitat, predators of all types coming and going over the course of the countless millennia—it can provide us with a clue for how to approach unanticipated risks in our world of financial markets.

What is remarkable about the cockroach is not only that it has survived for so long, but that it has done so with a singularly basic and seemingly suboptimal mechanism: Its defense mechanism is limited to moving away from slight puffs of air, puffs that might signal an approaching predator. This risk-management structure is extremely coarse; it ignores a wide set of information about the environment—visual and olfactory cues, for example—that one would think an optimal risk management system would take into account. The rule the cockroach obeys is so simple that it depends only on its giant fiber nervous system; it is a reaction that does not need to be filtered through its brain, but rather goes directly from the sensory hairs that detect the puff of air to the thoracic ganglia controlling its leg motion.

This same pattern of behavior—using coarse decision rules that ignore valuable information—appears in other species with good track records of survivability. The crayfish, another old-timer in the evolutionary tree, uses a winner-take-all escape mechanism where a stimulus triggers a set of neurons, each dictating a pattern of action, and one variant of behavior then suppresses the circuits controlling the alternative actions. That is, although a number of different stimuli are received and processed, all but one of them are ignored.

FATE FINISHES THE FURU

If the cockroach is an example of how to design a creature to survive in a world with unanticipated environments, the furu, a once-dominant fish in Lake Victoria in the middle of Africa, is a good example of how a specialized creature can be defeated by unanticipated environmental changes. Lake Victoria, the world's second largest freshwater lake, covering an area the size of Ireland, lies in east-central Africa, bordering Kenya, Tanzania, and Uganda, and is the chief headwater for the Nile River. For its huge size, it is relatively shallow, with an irregular coastline

of countless inlets and swampy bays. Straddling the equator, its evaporation creates morning clouds that define the weather for a wide region of the continent.

Living in protected isolation in this vast and varied habitat, the small perchlike furu specialized to a remarkable degree, diverging from a single species over the relatively brief 12,000-year life of the lake to at least 300 species, ranging in size from 4 to 12 inches. There are furu that survive as scavengers living off of the organic waste of the lake bottom; algae scrapers that feed off of shoreline rocks; snail crushers that have developed a second set of powerful teeth in their throats for that function; snail shellers that have developed long teeth to pull the snail out before it can fully retract into its protective shell; larvae eaters that sift insect larvae out through mouthfuls of mud; prawn eaters that inhabit the deeper waters where prawns live; and “child eaters” that eat the newly hatched eggs of other furu just after they are released from their mother’s mouth, or in some cases by first ramming the mother to dislodge the eggs from her mouth.⁹

For the biologist, the furu of Lake Victoria rival the finches that Darwin studied in the Galapagos Islands. But for the fisherman they are just small, bony trash fish. In the summer of 1954, the lure of the lake to the naturalist and the fisherman was changed forever by the actions of a Kenyan game fisheries officer with a bucketful of Nile perch.

Unlike the diminutive furu, the Nile perch is a marketable catch—an adult can weigh upwards of 100 pounds. In the mid-1950s the fish was introduced to other African lakes, including Lake Kyoga in Uganda, to the north of Lake Victoria, with spectacular results: Commercial fish production rose tenfold in just a few years. But these other lakes contained species of fish that had time to adapt to the Nile perch or had habitats where the Nile perch did not tend to go. Neither of these conditions turned out to be the case in Lake Victoria.

In the two decades following the initial stocking of Nile perch in Lake Victoria, naturalists who were following the furu found that they were increasingly pulling Nile perch out of their nets. Soon the only place they came across the furu was in the stomachs of the predatory Nile perch, “often still whole and wearing surprised expressions.”¹⁰ It seems the furu did not know what hit them. Defenseless and apparently clueless to the voracious predator that had been unleashed in their midst, they were rapidly becoming extinct. But their impending extinction was not the result of nat-

ural selection based on fitness in the usual sense; they were diverse and suited for almost every conceivable element of the Lake Victoria ecology. There was, however, one component of behavior where this was not so, a component that had not mattered at all in the thousands of years they had inhabited the lake but that made all the difference once the Nile perch was introduced. With the exception of a few of the insect- and snail-eating species, the furu at some point in their life cycle move out of the littoral areas and head for the open waters. Because it is such a large fish, the Nile perch tends to stay in deeper waters, so fish that stay near the shoreline, inlets, and rocks might go their whole lives without running into one. For Lake Victoria, this represents a lot of secure real estate. But the furu never had any evolutionary need to distinguish between the shallow coastline and the deeper waters. This did not represent a failure of fitness or an inability to adapt to its environment. Its path toward extinction was just a result of the dumb luck that someone had introduced an alien species into its waters. It was no different than if a nuclear holocaust blanketed the earth with radiation, killing off all mammals and allowing the more resistant insects to flourish. The relative immunity of the insects to radioactivity is not the result of any evolutionary plan. It has nothing to do with a strategy for survival; it is just a lucky byproduct of their evolutionary development.¹¹

PRIMAL RISK AND THE CASE FOR COARSE HUMANS

The cockroach's ganglions are programmed to address the unforeseeable, primal risk that has escaped our axiomatic approach to probability and risk measurement. It might have escaped us because we do not readily admit to being subjected to this type of risk, a risk that we might describe as "free-floating anxiety," or perhaps because by construction it is a risk that we can do little to address.¹²

The cockroach and the furu are just two of many examples I can cite in biology to illustrate the benefits of coarse behavior and the perils of fine-tuned behavior in reacting to the broad range of natural uncertainty. Many species seem to have gotten the message, displaying coarse behavior that appears to ignore information or fails to differentiate when a focused and finely differentiated behavior would appear to be optimal. For example, the great tit does not forage solely on the small set of plants that maximize its nutritional intake; it will forage on plants with a lower nutritional

value and fly afield in order to do so. The salamander does not fully differentiate between small and large flies in its diet. It will forage on smaller flies even though the ratio of effort to nutrition makes such a choice suboptimal. This behavior, although not totally responsive to the current environment, enhances survivability if the nature of the food source unexpectedly changes.

We also see animals increase the coarseness of their response when the environment changes in unforeseeable ways. For example, animals placed for the first time in a laboratory setting often show a less than fine-tuned response to stimuli and follow a less discriminating diet than they do in the wild. In fact, in some experiments, dogs placed in a totally unfamiliar experimental environment would curl up and ignore all stimuli, a condition called experimental neurosis. This problem led one experimental biologist to declare that “observing rats in mazes can tell you nothing other than how rats behave in mazes.”

The coarse response, although suboptimal for any one environment, is more than satisfactory for a wide range of unforeseeable ones. In contrast, an animal that has found a well-defined and unvarying niche may follow a specialized rule that depends critically on that animal’s narrow perception of its world. If the world continues on as the animal perceives it—with the same predators, food sources, and landscape—the animal will survive. If the world changes in ways beyond the animal’s experience, however, the animal will die off. Precision and focus in addressing the known comes at the cost of reduced ability to address the unknown.

What is the application of this concept to humans and markets? How can we use this idea to understand how people will behave? In our *Journal of Theoretical Biology* paper, Joe Langsam and I introduced a concept we call the omniscient planner, the purpose of which is to formulate mathematical tests to determine whether a particular coarse behavior is consistent with improving survival in the face of unanticipatable uncertainty. The idea of the omniscient planner rests on the following premises:

- You have an omniscient view of the future and you know all the types of risks that a species will face.
- You are required to program rules for that species to give it the best chance of survival, not in the current environment but across all of those possible future environments.

- You have one critical constraint: You cannot communicate through your rules any information regarding these unknown future states. (This might sound a little bit like the *Star Trek* Prime Directive.)

Generally, the survival rules set by this omniscient planner will not be the same as the optimal survival rules of any one environment. They will tend to differ by being more coarse; that is, they will tend to lead to an animal that appears to ignore information.

Now extend this idea from biology into the financial world. Imagine an omniscient planner who is creating rules of behavior for an investor. The omniscient planner will know all of the information that will be available for the investor to use, but also can see various crises that will emerge in the future, crises that cannot be anticipated by the investor given the history of the market to that point or the information the investor will have at his disposal. The constraint the omniscient planner has is the inability to communicate this knowledge to the investor. The rules can only dictate behavior based on the information the investor can observe from his nonomniscient viewpoint. In our paper we show that faced with these prospects, the omniscient planner will set up a trading rule that will differ from an optimization based solely on the information that is known. And the rule will appear to ignore some of this information and will be less than optimally reactive to the observed environment; in other words, the investor who follows the rules of the omniscient planner will display coarse behavior.

Perhaps we have been wired to behave in a coarse manner. We run our lives ignoring some information and reacting less than fully to our circumstances. And in doing this we are following a response that may be the best for the much broader world, one replete with unanticipatable events.¹³ The paradox is that this uncertainty leads us to take actions that are more predictable and less fine-tuned. We are wired to leave \$20 bills on the sidewalk.

OUR NOT-SO-EFFICIENT REALITY

The elegant mathematical framework of optimal behavior that underlies neoclassical economics cannot reconcile itself to actual human behavior because humans are dealing with a type of real-world risk that is not described by statistical distributions. Going one step further, this behavior

also leads to the refutation of one outcropping of neoclassical economics, the efficient market paradigm. Hedge funds and other investors make money in spite of a world with full information because people do not act based on full information. That market inefficiencies occur speaks to the fact that there is more to life than financial markets and financial risks. Our coarse behavior literally leaves money for the taking. It may well be that this is not a failing, but rather an insignificant cost for the genetic structure that allows us to better survive as a species.

As another example of the optimality of coarse behavior rules, imagine that the CEO of a corporation has been given an omniscient view of the world and can see all possible sources of risk. Armed with this knowledge, the CEO is allowed to design an approach to risk management—but with one constraint: Although aware of all the risks that the company has not identified, the CEO is not permitted to convey any information about these risks. The CEO can create a structure to address both the seen and the unseen risks, but cannot open up the unseen to view. The CEO can use full information to create the structure, but the resulting structure cannot convey information that would not otherwise be known. How will the risk-management structure the CEO designs differ from the structure he would design if he were never afforded this omniscient view?

If the CEO follows the lessons from the biological world, he will transfer resources away from managing the known risks and reconfigure the risk-management structure to better respond to the risks that remain unknown. Knowing that the structure cannot address all the risks, he will streamline and simplify formal processes and procedures because some of those procedures will obscure the unseen risks. The CEO will decide that it is better to spend less time focusing on detailed investigation of the known risks and more time thinking and reacting to the unknown risks. Similarly, he will simplify the risk-management models and analyses. Specialized analysis, although important in providing perspective for what is known, can only coincidentally do the same for the unknown.

Coarse measures will be more likely to indicate—although perhaps not fully elucidate—areas of unanticipated risk. And being more concise, these measures will be easier to discuss and analyze intuitively than specialized measures. Our omniscient CEO will reduce the organizational complexity and hierarchy of responsibility. This complexity, although perhaps effective in a specific environment, will obscure unidentified risks that fall

across organizational lines and slow the company's ability to respond to events that are beyond the design and function of the risk-management system. In short, this CEO will eschew pinpoint targeting of the observed risks in favor of lower-resolution, 360-degree radar that is more likely to capture the unobserved risks.

The unanticipated events that are catastrophic for a species are not difficult to understand in retrospect; we can see that the species was simply not designed to anticipate and react to certain events. As recent history has shown, the same is true of catastrophic risks in the financial world. Whether it is Amaranth, Kidder, Peabody & Company, or Barings Securities, the catastrophic risk that devastated these firms can easily be described in one or two sentences (and described without resorting to Greek letters or statistics). That the risks are obvious after the fact is not surprising; every loss can be explained *ex post*. It is tempting to think "they should have known" or "how could anyone run a firm like that?" But as simple as these problems appeared to be, and as simple as it might have been to prevent them, new events continue to surprise us.

The root of the problem is not the complexity of the unseen risks; it is the complexity of the organization, as we have seen from the airline and nuclear energy industries. Yet we continue to turn up the dial—witness the spate of mergers among financial institutions. The ever-increasing execution speed, the coupling of previously discrete processes, the reduction in the number of financial firms, and the refocus of these firms toward broad-based retail service have reduced their capacity to absorb or buffer trading risk.

All of this makes the lesson of the cockroach that much more germane. The finely tuned approach to risk—the approach that would seem optimal in any one world—may in the long run prove suboptimal. Given the complexity and fundamental unpredictability of nature, an approach that is coarse and less complex may be the best long-term risk-management strategy. The lesson for the corporation or investment firm is even more pointed: Unlike the biological world, in the business world the more intricate risk-management structures may actually make the situation worse, leading to greater complexity rather than simply a less robust response.

When providing for safeguards against market failures, it may be dangerous to assume that institutions behave rationally. In a rational institution, a company's objectives are in line with those of the shareholders: that

those in the organization seek to maximize the shareholders' value, that the regulatory institutions follow their mandates to acquire relevant information for monitoring the financial institutions, and that all the members of these organizations focus in unison to execute the public trust.

It is hard to recite this image of the rational institution without breaking into a smile. If individuals do not behave rationally—at least in the economic sense of the term—then it is hard to expect institutions, which are made up of individuals, to be rational. The rewards for those in most institutions are not aligned with those of management, which in turn may not be aligned with the interests of the owners. This has given rise to an entire field of analysis called agency theory, which looks at how the disconnect between the incentives of the managers and those of the owners leads to economic distortions.¹⁴

When searching for the causes of crises, lack of rationality is an easy target. Organizations are inefficient, and larger organizations tend to be less nimble in response to the unexpected and more prone to failure. While this view gets at part of the problem, it does not explain why, despite overlay after overlay of safety measures and regulation, things continue to go wrong. Perhaps the problem rests in the very structure and nature of organizations and institutions. The conflicting goals and mounting inefficiency of operating in a large, multilayered environment are part of this institutional structure, and will exist whether the individual members of the organization are rational or not.

Once a normal accident occurs, controls can address recurrences. But if a system is already at a level of complexity where normal accidents are common, adding one control after another will exacerbate complexity and obscurity. Indeed, the conventional response to the unanticipated events will produce a cycle where better and better controls of the identified risks will cause an upward spiral of surprise events. This has obvious and important implications for the limits of regulation. If risk management can fail in unanticipated ways, then adding more controls can't address the issue. This potential for failure makes a statement about the prevalence of system risks and should lead to a more coarse, not a more detailed, response. If we build finely tuned structures of rules and systems that only address the risks we can embrace and understand, we will continue to wander down the path of complexity, and pay its exorbitant toll.

THE DANGER TO THE SYSTEM IS THE SYSTEM

As experience has shown, the problems that lead to crisis emerge from within the market. In the case of the Internet market bubble, it was the coalescence of trading views, leading all those remaining in the market to bid against one another, encouraged at each turn by the increases in prices that were nothing other than their own doing. In the LTCM meltdown, it was the liquidation forced by the creditors, which led prices to drop, causing the successive liquidations that pushed creditors to take even more draconian actions. In the 1987 crash, it was the hedging actions that led to the price declines, which, through the nonlinearity of the strategies, led to even more aggressive hedging.

We can react to opportunities only based on the knowledge that we have. We can manage risks only when we can identify them and ponder their possible outcomes. We can manage market risk because we know securities prices are uncertain; credit risk because we know companies can default; operational risk because we know missteps are possible in settlement and clearing. But despite all those risks we can control, the greatest ones remain beyond our control. These are the risks we do not see, things beyond the veil. The challenge in risk management lies in our ability to deal with these unidentified risks. It is more than a challenge; it is a paradox: How can we manage a risk we do not know exists? The answer is that we cannot do so directly. But we can identify characteristics of risk management that will increase our ability to react to them.

CHAPTER 11

HEDGE FUND EXISTENTIAL

The hedge fund world got a lot more crowded in the half dozen years after I entered the industry. Elite firms such as of Moore, Tudor, Tiger, LTCM, and D.E. Shaw were joined by dozens more percolating up from the Wall Street establishment. Their clientele evolved, moving up the chain from wealthy individuals to large institutions. Hedge funds became an investment phenomenon, and as such they became subject to increased regulatory scrutiny because of their supposedly wild ways, which provided fodder for journalistic forays into greed and fraud. The reality is a bit more sober. For starters, it is hard to come up with a common feature or quality that unites the 5,000 plus hedge funds. With so much focus on them, it is worth asking what hedge funds really are. Although many outfits are branded as such, do hedge funds really exist as a definable entity? Is there such a thing as “hedgefundness”?

These are existential questions with practical import. We are in a period of mounting interest in hedge funds, interest that is backed by a growing demand for performance measurement, hedge fund indexes, and tracking portfolios, as well as for transparency and regulation. The unspoken assumption is that hedge funds are a homogeneous entity, for only in that case does it make sense to analyze, index, and risk manage hedge funds as a class.

I believe much of what is proposed for hedge fund oversight and analysis will turn out to be a fruitless exercise because the concept of hedge funds defies a meaningful definition. If we persist in trying to categorize them, we will run headlong into the entire universe of possible investment trading strategies applied to the universe of tradable securities and financial instruments. We will have a definition that provides no distinction.

In this sense, there is no such thing as a hedge fund. It is not part of a homogeneous class that can be analyzed in a consistent way. The hedge funds/alternative investments moniker is a description of what an investment fund is not, rather than what it is. The universe of alternative investments is just that: the universe. It encompasses all possible investment vehicles and all possible investment strategies minus the traditional investment funds and vehicles.

If you think in terms of leverage, alternative investments are the entire universe, with the exception of those funds that are restricted to leverage that is no greater than one to one. If you think in terms of positions, alternative investments are the entire universe, except those funds that are restricted to only being long. If you think in terms of securities, alternative investments are the entire universe, except those funds that are restricted to a somewhat arbitrary and generally evolving set of traditional assets: domestic stocks and bonds, or perhaps—a more recent extension of the definition of “traditional”—the stocks and bonds of the G-10 countries.

It is true that the vast majority of investment wealth is concentrated in traditional strategies, but for all their size, these strategies constitute just a small part of the overall investment universe, one galaxy in a system of billions. What we call alternative investments is really the wide world of investments minus that small slice known as traditional management. From the perspective of the traditional investment world, the burgeoning market of hedge funds must be like zooming in on a bustling island from a satellite, only to discover, when the camera is pulled back, that beyond the

island's shores lies an entire continent. Or perhaps it looks like Saul Steinberg's famous "View of the World from Ninth Avenue" *New Yorker* cover where the land west of the Hudson River, from New Jersey to Japan, is sparsely sketched out, an afterthought.

Defining hedge funds as a complement to traditional investment funds, as the universe minus a subset, rather than as a sideshow to traditional investment management, is more than an issue of semantics. It changes the way we must think about hedge funds. It changes the potential value of broad-scale studies of hedge funds, and it reorients our thinking on what it means to regulate them. If they do not exist as a well-defined class, we can never get our arms around them. And if we think we have succeeded in doing so, then what we have really done is embrace the whole of investments.

The problem I have sketched out becomes readily evident when we look at recent attempts to create a cogent classification of hedge funds. The most common approach to classifying hedge funds, one used by Hedge Fund Review, CSFB/Tremont, and Standard & Poor's, is to organize them based on trading styles. For example, the Standard & Poor's Hedge Fund Index has three styles: arbitrage, event-driven, and directional/tactical. Each of these styles has three strategic subsets. Arbitrage consists of equity market neutral, fixed income arbitrage, and convertible arbitrage; event-driven has merger arbitrage, distressed, and special situations; directional/tactical has long/short equity, managed futures, and macro.

The problem with this sort of classification, based as it is strictly on the trading style or strategy type, is that it has to be revised over time as new strategies emerge and existing ones fail.

An alternative classification matrix, which I developed in 2001, attempts to overcome this problem, but in so doing reveals the existential issue for hedge funds.¹ This approach classifies hedge funds by five characteristics:

1. *Asset class.* The broadest category, it defines the market in which the fund operates. These include fixed income, equities, currencies, and commodities. There can also be a "multiclass" to capture "none or some of the above," and this specifically includes global macro funds.

2. *Direction.* As the name implies, the direction of the manager's activity in the asset: long, short, long/short, and neutral. We might think of a long strategy as having an average correlation with the relevant index of over .5 (the figure indicates the strength of the relationship; a perfect correlation equals 1); long/short as being between $-.5$ and $.5$; and neutral as being between $-.2$ and $.2$. Another category for direction that is useful is event, which depicts strategies that usually have low correlation with the market, but on occasion the correlation can be very high (e.g., during a liquidity or credit crisis).
3. *Investment type.* This provides more information about the specifics of the investment process or strategy. For example, in the neutral classification there is relative value and statistical arbitrage; the event classification would include merger arbitrage, credit arbitrage, and distressed debt. Investment type is the one component of the analysis that will vary over time with the introduction of new investment strategies.
4. *Geographic region.* Where is the fund trading? Differentiation may be limited to G-10 and emerging markets, or the region can be broken out in more detail.
5. *Liquidity.* Some funds trade short-term and in instruments that can be traded easily. Others are less liquid because of either their strategies, the types of instruments they hold, or the size of their holdings. A second approach to classifying liquidity is turnover, which addresses the percentage of portfolio turnover on a monthly or annual basis.

With this categorization scheme, a typical technology fund might be "Stocks-Long/short-Relative value-U.S.-Highly liquid." A Japanese distressed debt fund would be "Bonds-Event-Distressed-Japan-Illiquid," while a U.S. corporate bond fund might be "Bonds-Long-Credit arbitrage-U.S.-Liquid."

This matrix provides a stable and robust framework for hedge funds. But it leads to a critical question for those who want to put all hedge funds into one little basket: What, then, is the classification scheme for non-hedge funds? What is the difference between the set of strategies embraced by these classifications and the universe of all strategies? If this is an effective categorization framework for hedge funds/alternative investments, what is the categorization framework for the alternative to this al-

ternative? The answer is, there is none. This categorization for hedge funds actually is a categorization for all investment strategies. After all, what investment strategy is not typified by some direction (especially since “neutral” is one choice), operating on some general asset class, and focused on some geographic region?

The same question arises when we consider hedge funds as a subject of study and research. I know of at least two institutes that are focused on the study of alternative investments. One is at the London School of Business, the other at the University of Massachusetts, Amherst. There are also several journals that focus on hedge funds and alternative investments.

If hedge funds are heterogeneous to the point of being the entire investment world less the small subset of traditional investment strategies, then hedge funds are a questionable topic for study. It would be like opening up a program to study all objects made of materials other than wood, or initiating research on contemporary history for every country but France. You could do so, but I don’t know how that study would be much different from simply having a study of all materials or of all modern history. In fact, the proper study of hedge funds cannot be differentiated from a general study of investments. Issues of risk, return, and liquidity apply to all hedge fund strategies, and indeed to the whole range of possible investments.

Consider the following scan of articles from various issues of the *Journal of Alternative Investments*, just one of a number of journals on hedge funds: “Currency Market Trading Performance”; “Timber Investment”; “Current Attitudes to Private Equity”; “Convertible Arbitrage: A Manager’s Perspective”; “Macro Trading and Investment Strategies”; “Commodity Trading Advisor Survey”; “Stock Selection in Eastern European Markets”; “Market Neutral versus Long/Short Equity”; “Merger Arbitrage: Evidence of Profitability”; “Analysis of Real Estate Investments in the U.S.”; “Benefits of International Small Cap Stocks.” What is the common ground, other than being related to investments? If this is a sample of articles specific to hedge funds, what would articles on the broader world of investments outside of hedge funds look like?

CAN WE REGULATE HEDGE FUNDS?

Hedge fund regulation is a topic of discussion that has reached a crescendo since the LTCM debacle, and subsequent wipeouts such as

Amaranth in the fall of 2006. But it is questionable whether regulating hedge funds as “hedge funds” can succeed, or whether it even makes sense to try. With so broad a classification, seeking a uniform approach would be like developing a single set of traffic rules to apply for all modes of transportation, from pedestrians to commercial jets. Or, actually, since alternative investments exclude traditional unlevered long-only investments, it would be like doing so for all modes of transportation except, say, passenger sedans.

It may be fruitful to impose regulations on leverage, add to the rules imposed on short sales, or discuss approaches to regulate offshore entities. But starting down the regulatory path with hedge funds as the objective is to fail before beginning, because this will be regulating an entity that cannot be well defined.

Let’s look at what happens when we do segment hedge funds in one of many possible dimensions.

On one extreme are the analytically driven funds. These would include options and volatility trading and some of the fixed income strategies like relative value trading and complex mortgage products. It could also include statistical arbitrage strategies, which are usually computer driven, are often short-term, and are tightly controlled to maintain market neutrality. On the other extreme are the strategies driven by market or economic events. These would include distressed debt and merger arbitrage, as well as opportunistic emerging market funds—funds that trade in countries that are on the precipice of crisis.

Related to the diverse sources of risk is the relationship between the availability of risk management tools and the ability to measure and manage risk. The more complex strategies—fixed income relative value and option arbitrage, to name two—usually have risk management as an integral part of their trading. It is difficult to be successful in trading options without good option pricing models, and these models provide the measures for volatility and gamma risk. In relative value trading, the margin of success comes from understanding convergence and model risk and in sizing trades properly for the opportunity. In these instances, the risk measurement is quantitative, and the tools to do that measurement are usually tied directly into the trading business.

By contrast, the event risk for distressed debt and merger arb is simply not objective and quantifiable. That is not to say there are no hedge funds

or investment houses that shove each position of these portfolios into a risk model and come out with statistics on the other side. But the results cannot be taken too seriously, because the main risks are idiosyncratic, one-time events, and those risks can be assessed only through experience and study of the particulars. Mathematical tools are of little value. It is hard to develop any sort of population for statistical analysis. Probabilities might be useful, but they come from the subjective feel for the likelihood and impact of one-off political events and business decisions.

It is clear that the ability to do risk management varies from one end of the spectrum to the other. Fortunately, those strategies with the most complex and quantitative risks also have the best tools in place to deal with them. Those strategies with more subjective risks do not generally have these tools, but then, those tools would not do them much good. That is, along this spectrum the capability to do risk management is in phase with the meaningfulness of the results. The level of detail to which risk management can reasonably be performed is in line with the degree to which that detail would have any value. There are those who can and already do and those who don't and probably can't.

The nature of the risks varies greatly from one extreme to the other. For option and volatility trading, gamma risk is a dominant issue; for relative value trading, control of leverage is most critical. For merger arb or distressed debt, the most important risks are one-off risks that are firm specific. I would hate to be the person assigned to try to develop one risk management system to deal with the entire spectrum. It is clear from a risk management perspective that hedge fund classification provides no unification, hardly even a common thread.

Not only do hedge funds cover the waterfront of investment strategies, of which the traditional asset management strategies are only a small part, but those executed by hedge funds are only one point on the spectrum of a surprisingly broad realm of business endeavor that might be called asset allocation businesses. We have already discussed the shorter-term role of the hedge fund in providing liquidity to the market; in that capacity the hedge fund manager is acting as a quasi market maker, taking over a function that has traditionally been the domain of the floor trader or the sell-side firm, but which, with the speed of communication and transparency of the market micro structure, can now be outsourced. In taking a longer-term view, as do funds that deal in the corporate arena—long/short funds,

private equity, vulture capital funds—there’s a shared objective for fund managers, entrepreneurs, CEOs, and corporate division managers. In each case, the objective is to run a set of assets to its best result. It is the focus and the problems in executing the objective—and thus the set of critical talents—that differ. For the entrepreneur, the problem may be mapping out a way to effectively employ seed capital; for the division manager, it may be providing leadership and judgment to best marshal the skills of the workforce; for the CEO it may be moving the company’s production and marketing assets toward the best opportunities within reach. The hedge fund manager is looking at the same issues these businessmen are: What is the vision for the start-up? What is the quality of the work force? Where are the resources of the firm best spent? The hedge fund portfolio manager sits somewhere along the continuum in trying to put resources to work efficiently and to identify opportunities in the market, but the allocation decisions are placed at a more macro level, deciding not between one division and another, but between one company and another.

THE HALF-LIFE OF HEDGE FUNDS

Karl Marx observed that the capitalist system requires an ever-expanding market in order to maintain profits. For capitalism to succeed, it must move forward by developing new products or opening up markets in backward colonies. Innovation creates a product that is ahead of the market, while foreign expansion pushes the current products on new markets. The Marxist view applies both literally and metaphorically to the penchant for innovation in financial products.

The traditional Treasury market, which comprised bare-knuckled traders whose gaze never moved past the price-yield relationships on their books, was fresh fodder for the complex world of fixed income options and swaps. In the same manner, investors in the conventional 30-year mortgage bonds were a Marxist “backward” market, easily “exploited” by trading desks that could use sophisticated models and specialized expertise to develop collateralized mortgage obligations (CMOs) and other mortgage derivatives. Each innovation added complexity to the market for the simple reason that complexity was what sold and what made the most money.

A large segment of the firms in the hedge fund/alternative investment world employ opportunistic strategies that feed off the inefficiencies

born of this complexity. For these firms, there is a problem even beyond the existential: the ever-changing reality of existence. The half-life of most innovation-driven inefficiencies is just three or four years. It takes a while for the decline in opportunities to be detected in fund performance, so opportunistic hedge funds may survive a bit longer than this, but whatever classification scheme or regulation structure is put in place will be operating against a morphing objective and, because future innovations cannot be predicted, an objective that will be changing in a way that is difficult to anticipate. The most successful strategies fade away and new ones emerge two or three years down the road, often based on securities that are new to the market.

To see this point, consider the history of opportunistic strategies. Although they were not executed within the traditional hedge fund structure, some of the early opportunistic strategies included basis trading on the cheapest-to-deliver bond shortly after the introduction of the Treasury bond futures, and cash-futures index arbitrage in the years following the introduction of the S&P and the Value Line futures. Both strategies peaked within a few years, and a decade later amounted to little more than background radiation in the trading firmament. O'Connor's Partnership was making hundreds of millions of dollars by applying the Black-Scholes formula to options in the nascent Chicago Board Options Exchange in the late 1970s and early 1980s, with a cadre of young traders grabbing their pricing sheets at the start of the day and taking their posts along the CBOE trading floor to apply delta hedges to mispriced options. By the mid-1980s, the writing was on the wall for margin contractions in the floor market-making business, and O'Connor's sold itself to Swiss Bank.

On the heels of the cash-futures and index arbitrage opportunities came statistical arbitrage, which was the first to emerge in a hedge fund structure. In 1985, the first statistical arbitrage strategy was developed at Morgan Stanley, by Gerry Bamberger, a young information technology (IT) person who had been assigned to work on some hedging issues on the equity trading floor. As we discussed earlier, Bamberger developed a pairs trading strategy that resulted in a burgeoning business for Morgan Stanley and spawned D.E. Shaw and a host of other stat arb firms. Although this strategy survived longer than some others, predictably the out-sized returns have dissipated over the past few years. Meanwhile, at Salomon Brothers, the focus was on fixed income relative value trading.

Starting with the “two-plus” yield curve model, a two-factor yield curve model that explicitly accounted for the effects of Federal Reserve action, Salomon’s legendary fixed income arbitrage group generated billions of dollars for the firm. This group then left Salomon to form LTCM. A few years later, the relative value opportunities collapsed, and, straining to maintain outsized returns in a less fertile market, LTCM ultimately became the most visible hedge fund supernova.

The birth and death cycle for opportunistic strategies will no doubt continue. Macro hedge fund strategies now face a world that has compressed into three major markets, with the central banks no longer ripe to be gamed. With concentrated demand for convertible bonds and default swaps, convertible bond and credit derivative arbitrage funds are spring-loaded for crisis.

Most of what remains are long/short equity funds. These funds, which include niche funds such as emerging markets, private equity, and event driven funds, make up more than half of all hedge funds, and are essentially traditional equity management in a different guise. Add an S&P futures position to the long/short equity fund and you have a fund that is hard to distinguish from a traditional long-only asset management fund. The only differences are that long/short equity funds tend to have more turnover in their positions and they tend to have more amplitude in their risk levels. But even here, once an index overlay is added, the end result will not look very different from the profile of an active asset management fund. The hedge fund simply concentrates the active component of the decision making. Once the appropriate adjustments are made, the remuneration to the long/short equity hedge also is similar to that of its long-only cousins. A 1-percent-and-20-percent fee structure leads to the same ballpark return for the hedge fund manager as a 100-basis-point fee will for the manager of a larger but unlevered long-only fund.

WILL HEDGE FUNDS TAKE OVER THE INVESTMENT WORLD?

Hedge funds are the unconstrained version of traditional investment funds in that they do not have restrictions on shorting, leveraging, or expanding to innovative asset classes. They can do everything a traditional manager can do and then some. Because of this, hedge funds should dom-

inate the traditional funds in generating returns. Looking at it another way, any traditional investment manager who finds himself passing up an opportunity to improve returns because he cannot get short exposure, cannot lever his exposure to a trade idea, or cannot take on a promising position because it lies outside of his allowable universe will be left behind by an equally talented counterpart who is following an identical investment method in a hedge fund.

This simple point, that an unconstrained investment process will dominate a constrained one, means that in the end hedge funds, whether in their present or some reformulated structure, will move from being the alternative to being the standard.

DO YOU BELIEVE?

A clergyman interviewing a parishioner asks, “Do you believe in baptism by immersion?” The parishioner unhesitatingly replies, “Do I believe in it? Why, I have actually seen it.” I have seen hedge funds and worked in hedge funds, but like the clergyman, the question is not about just having seen hedge funds but seeing what qualities exist beyond the name. From one hedge fund to another, we may observe a family resemblance, but there is not enough of a subject to grasp, study, risk manage, or regulate. To study hedge funds is to study the world of investment strategies. As for risk managing and regulating hedge funds, faced with so general a task, what more can one say than “be careful”? I believe that we will discover, as we continue in our attempts to grapple with hedge funds, that we have enveloped the entire world of investments under a name that sounds a lot more exotic than it is.

CONCLUSION

BUILT TO CRASH?

The question posed by this book, simply put, is: Why can't the financial markets seem to get their act together? Why, in spite of reduced risk in the underlying economy, in spite of the march of innovation and the contributions of financial engineering, do we not enjoy reductions in financial risk that we find in other areas of our lives? Why are markets actually becoming more crisis-prone?

One answer can be found in the effects of innovation. It is undeniable that innovation has had some positive effects on the markets. It has improved the markets by making them mechanically more efficient. The markets are more liquid and quicker to react to information. Information flows more freely and is distributed more widely, and prices are readily available to virtually all participants. Trades are executed nearly instantaneously worldwide at transaction costs that are a small fraction of what they were a few decades ago. And, whether developed with the intent of better meeting the demands of investors or, more cynically, to stave off commoditization and maintain profitability, we are awash in new and innovative instruments.

But the positive effects of innovation come at a price. Innovation increases complexity. Many innovative instruments are in the form of derivatives with conditional and nonlinear payoffs. When a market dislocation arises, it is difficult to know how the prices of these instruments will react. Innovation and mechanical efficiency have also increased complexity by pushing markets to become more interconnected. Thanks to globalization,

CONCLUSION

a problem in one market can affect another even when there is no economic relationship between the two simply because investor portfolios or bank credit lines have exposures to each. Innovation has also led the markets to become more tightly coupled. This tight coupling, and the resulting higher liquidity, makes it easier to take on levered positions, because more liquid and readily priced securities make for better collateral.

The combination of tight coupling and complexity is a formula for normal accidents—accidents that are all but inevitable as a result of the structure of the system. We have analyzed these sorts of accidents in airlines with the ValuJet crash and in nuclear power plants with Three Mile Island and Chernobyl. In all of these cases disaster was triggered by simple and apparently innocent actions that initiated a chain of compounding problems because of the complex nature of the system. The tight coupling from one link to the next precluded any kind of stop lever; no one could sit back and say, “Wait a minute, let’s shut things down and think about what is happening here.” What is all the more troubling is that attempts to add layers of safeguards or regulation to prevent these disasters may actually do the opposite by increasing the complexity. The catalyst for the ValuJet crash was a regulated safeguard. The problems at Chernobyl started with a safety test.

Financial risk is also higher because the markets increasingly assume a mathematically precise rationality, as opposed to the way we actually do, or indeed really should, behave. People do crazy things all the time, yet the efficient market paradigm assumes that investors take all information into account and react quickly and rationally. The world is not well described by this paradigm; we tend to be coarse in our responses and we leave information by the wayside. We do this because we conduct our lives with a sense of unanticipatable, primal risk that remains unconsidered in the market’s design. Recall the cockroach—scurrying along over millions of years, as jungles turn to deserts and cities rise and fall, ignoring most of the information the environment has to offer—versus the furu, optimized and specialized to take advantage of every nook and cranny of its niche in Lake Victoria. We are wired to deal with a type of uncertainty we cannot recognize, and this leads us to exhibit behavior that is inconsistent with the mathematical rationality that underlies the paradigm on which financial engineering is based. We fail to take the degree of fine-tuned actions that conventional optimization would dictate, and fortunately so, be-

CONCLUSION

cause, like the furu, conventional optimization pairs off only against the current world with the risks and uncertainties that can be identified within it.

The implications of this uncertainty are profound, extending beyond the financial markets. But because the markets feed on risk and are largely free of friction and institutional constraints, apparently suboptimal behavior may have its most significant and obvious effects there. For example, coarse behavior leads to a paradoxical corollary: Greater uncertainty leads to more predictable behavior. Within the limited world of finance, this predictability is a decided negative. It might be the reason traders and hedge funds can pick us off to make profits. But within the broader world in which we live, and the many possible worlds that might unexpectedly arise, this predictability and the coarse behavior from which it stems leaves us—and our biological compatriots—more capable of survival.

Market crises are not born from nature. They are not transmitted by economic or natural catastrophe. The machinery of the market itself can take a small event and distort it. The more closely we try to follow the ideal, thereby adding complexity and more tightly coupling the actions of the market, the more frequently crises will occur. Attempts at that point to add safety features, to layer on regulations and safeguards, will only add to the complexity of the system and make the accidents more frequent. And when blowups happen in the future I can guarantee that the focus will be directed improperly: not at the issues of market design but at hedge funds where the events are observed. They will be implicated for the simple reason that they are engulfing more and more of the risk-taking landscape. The perception of hedge funds being what it is, they will take the blame and become subject to increased regulation. But blaming hedge funds is a little bit like *The Simpsons* episode in which a meteorite hits Springfield and the townspeople gather, shouting, “Let’s burn down the observatory so this never happens again!” True, the hedge funds are the institutions that have the appetite for the risk; but there is nothing inherent in hedge funds, nothing that they represent as a unified set, that makes them the singular cause of anything.

So if we are subject to risks that we cannot even anticipate, if we have built a world of complexity and tight coupling where adding regulation only makes matters worse, is there any more we can say other than “get used to it”?

CONCLUSION

In the basement of a rundown office building on West 30th Street in New York's Garment District resides Renzo Gracie's Brazilian Jiu Jitsu Academy. It is not only a place where I train several times a week, but one that also offers a great lesson in demonstrating a method for dealing with the endogenous risk of the market.

How "Jiu Jitsu" and "Brazilian" came to be joined in one breath is an interesting story. In 1912 a large group of Japanese immigrated to northern Brazil. They were assisted by Mitsuyo Maeda, a noted Japanese jujitsu expert who had traveled throughout the Americas and Europe teaching the art before emigrating to Brazil. Gastao Gracie, a diplomat and businessman, had befriended Maeda and arranged for the group's immigration. To show his appreciation, Maeda taught jujitsu to Gastao's oldest son, Carlos Gracie. He was just 14 when he began. A dedicated student, he adapted Maeda's techniques to be more effective for fighting in the Brazilian streets. He not only entered competitions, but also advertised in newspapers to find opponents against whom he could test and improve his style. Later he taught his younger brothers, the youngest of whom, Helio, became especially adept. Helio had a slight build, and to accommodate his small size and lack of strength he further modified the jujitsu style, focusing more on technique and less on power or athleticism. Helio fought more than 600 matches with only two defeats, both occurring when he was past the age of 45. The techniques refined by Carlos and Helio were passed on to their children and through them to the next generation as well. Carlos had 21 children and Helio had 9, so a dynasty of Brazilian Jiu Jitsu fighters was born. Today there are more than 40 Gracie family members who either teach or compete, Renzo being one of their number.

The confidence of the Gracies was manifest in an open challenge: Anyone could walk into their academy and demand a fight on the spot, with no time limit and no rules. Any challenger's notion of engaging in a bare-knuckle slugfest would be quickly disarmed by the nature of Brazilian Jiu Jitsu. Little if any damage would occur before the Gracies would take the fight to the ground. Once there, it would end with submission from either a choke hold or a joint lock. The opponent, sensing he would shortly lose consciousness in the first instance or would have his arm broken in the second, would "tap out" on the floor to signal his desire to end the contest. It made the challenge workable both for the Gracies and for their foes. There was no harm done, but the superiority of the Gracie tech-

CONCLUSION

niques was made clear. Over time these sorts of one-on-ones expanded to take place in open, no-holds-barred competitions called *vale tudo*, Portuguese for “anything goes,” which became common throughout Brazil as a practical testing ground. These competitions came to the United States in 1994 when Rorion and Royce Gracie, two of Helio’s sons, promoted mixed martial arts tournaments in the *vale tudo* tradition to demonstrate to the world at large the domination that Brazilian Jiu Jitsu had enjoyed in Brazil for the better part of a century.

The very existence of *vale tudo* competition points to the key aspect of Brazilian Jiu Jitsu that allowed it to gain superiority over other martial arts. It was not just that it was actively improved through competition while other forms stayed rooted in tradition. It was that it had been developed in a way that allowed this competition and testing to occur in the first place. The techniques that formed the basis of Brazilian Jiu Jitsu could be practiced without causing harm. Many fighting techniques cannot be practiced live because they inevitably cause injury. The genius behind the development of Brazilian Jiu Jitsu was to select for inclusion only those techniques that are applied in a slow and measured way so that an opponent can stop before injury occurs.¹

What became apparent over time was that having a firm understanding of the actual application of a set of controlled techniques through live training and real fights was superior to having a quiver filled with techniques that were powerful in theory but could not be tried and refined until an actual fight occurred. You can see where this is leading when it comes to the markets. Does it make sense to do the same thing? Should we pull away from dangerous innovations, even if those innovations appear to be useful? And what constitutes those that are dangerous?

I believe the markets can better conquer their endogenous risks if we do not include every financial instrument that can be dreamed up, and take the time to gain experience with the standard instruments we already have. Just because you can turn some cash flow into a tradable asset doesn’t mean you should; just because you can create a swap or forward contract to trade on some state variable doesn’t mean it makes sense to do so. Well, in the efficient market paradigm it does, because there nirvana is attained when a position can be taken against every possible state of nature. But in the world of normal accidents and primal risk, limitless trading possibilities might cause more harm than good. Each innovation adds layers

CONCLUSION

of increasing complexity and tight coupling. And these cannot be easily disarmed through oversight or regulation. If anything, attempts at regulating a complex system just makes matters worse. Furthermore, if an innovation is predicated on behavior predicted by the efficient markets theory, then things may not operate as advertised: People just don't behave that way. The point is that these innovations have externalities for the entire financial system that are hard to measure but dominate their apparent value. Rather than adding complexity and then trying to manage its consequences with regulation, we should rein in the sources of complexity at the outset.

Linked to the need to reduce market complexity is the need to relax tight coupling. The easiest course for reducing tight coupling is to reduce the speed of market activity. This has been done in times of crisis through the imposition of bank holidays and so-called circuit breakers, but doing so on a day-to-day basis would turn back the clock on financial markets in an unacceptable way. A less disruptive course of action is to reduce the amount of leverage that comes as a result of the liquidity, since this is ultimately the culprit that high liquidity and speed of execution breeds. The externalities to high leverage are greater than they appear, because on most days everything runs smoothly. But as we have seen time and again, in the instances where it really matters the liquidity that is supposed to justify the leverage will disappear with a resulting spiral into crisis.

Simpler financial instruments and less leverage make up a painfully obvious prescription for fixing the design of our markets. These modifications will lead to a financial marketplace that will be apparently less finely tuned and less responsive to investor needs. But, like the coarse response mechanism of the cockroach, when faced with the inevitable march of events that we cannot even contemplate, simpler financial instruments and less leverage will create a market that is more robust and survivable.

NOTES

CHAPTER 1 Introduction

1. The problem of increasing market risk in the face of reduced economic risk, and its implications for greater endogenous risk in the financial markets, is presented in detail in Horace W. Brock, “The Transformation of Risk: Main Street versus Wall Street,” May 2002 Profile Report, Strategic Economic Decisions, Inc.

CHAPTER 2 The Demons of '87

1. Fischer Black and Myron Scholes, “The Pricing of Options and Corporate Liabilities,” *Journal of Political Economy* 81, no. 3 (May/June 1973). Robert C. Merton, “Theory of Rational Option Pricing,” *Bell Journal of Economics and Management Science* 4 (Spring 1973); reprinted as Chapter 8 in his book, *Continuous-Time Finance* (Malden, MA: Blackwell Publishing, 1992).

What is most remarkable about the Black-Scholes formula is that there is no need to know the expected course of the security price on which the option is based, the preferences of investors, or economic conditions, all things that are subjective and hard to predict. If you follow the dynamic hedging prescriptions dictated by the model, you generate a payoff that looks just like that of the call options you specified in setting the formula up.

For all the considerable math and computing power behind it, operationally option pricing theory and the portfolio insurance that was derived from it are really nothing more than a sophisticated version of a common stop-loss strategy, in which a stock is sold at a predetermined price to prevent further losses. The most common stop-loss strategy can leave something to be desired, because it makes no provision for short-term reversals. A slightly more sophisticated approach might entail dumping a stock that

drops below some exit or floor price and getting back in if the stock starts to move up in price again. The investor will be protected against large losses while still enjoying a share of any future gains, with the obvious drawback that the investor racks up a loss every time the stock bounces back and forth between his entry and exit points.

An even more sophisticated stop-loss strategy moves the investor out of the stock gradually as it drops toward the floor and then gradually increases the position again as it rises above the floor—essentially a dynamic stop-loss strategy. This will reduce the costs that are incurred from selling the entire position every time the portfolio hits the floor value and buying it all back after the stock recovers. The cost of this whipsawing cannot be fully eliminated, because the essence of the stop-loss structure, even in its dynamic form, is to sell on the way down and buy on the way up. As a result, the strategy always entails some slippage because it makes adjustments to the price change only after the fact. But the rate of selling as prices decline and of buying back into the position as prices move up can be structured to lead to a remarkable result: If the buying and selling of this dynamic stop-loss strategy are done in the mathematically correct way, the total cost of executing the strategy can be predicted with a surprising level of accuracy no matter what path prices take. Whether the stock drops like a rock, trends up, or bounces around the floor value, the total cost of this stop-loss strategy will be the same: the price of an option that has an exercise price equal to the floor value.

There are a few things that can throw a monkey wrench into the pricing but that are assumed to be sufficiently stable to not matter—although of course they then end up mattering at the worst possible times. The cost of the option will change if interest rates or the stock volatility changes. Interest rates matter because they affect the financing cost of the option hedge, and volatility matters because higher volatility will increase the frequency of the costly price adjustments. And the option model assumes stocks follow a smooth price path; sudden jumps in price end up being problematic unless they are explicitly built into the model.

2. Their approach is presented in M. Rubinstein and H. Leland, “Replicating Options with Positions in Stocks and Cash,” *Financial Analysts Journal* 37, no. 4 (July–August 1981): 63–72. The link between portfolio insurance and option theory should come as no surprise, because the payoff of this strategy—enjoying any upside of the market while giving a minimum floor value for a predetermined cost—is nothing more or less than a call option, which gives the holder the right, but not the obligation, to buy a stock or other asset at a predetermined price. Thus the dynamic stop-loss strategy at the core of portfolio insurance was designed to transform a portfolio into one giant call option; it replicated the payoff from a call option that had an exercise price equal to the floor.

NOTES

3. Another person who detected the emerging problem was Sandy Grossman, a brilliant researcher who was one of the early academics to move into the hedge fund world. In the months before the 1987 crash, Sandy noticed that option implied volatility was expanding relative to the volatility of the underlying stocks and that this differential was trying to tell the markets something: There was a lot of liquidity demand chasing these options, and the actual cost of hedging was probably higher than it appeared. The markets had something to say, but no one was listening. Instead, the higher option volatility meant more portfolio insurance moved toward dynamic futures hedges rather than doing the replication with the better-tracking, but apparently more expensive, exchange-traded options. His analysis was later published in Sanford Grossman, "An Analysis of the Implications for Futures and Stock Volatility of Program Trading and Dynamic Hedging Strategies," *Journal of Business* 61, no. 3 (1988): 275–298.
4. This blur of stocks was not unique to the day of the crash in October of 1987. Just over a decade earlier, a more sustained but far more severe market downturn left the stock market with a loss of more than 40 percent. It was the era of Watergate, the oil crisis, and Gerald Ford's "Whip Inflation Now" buttons. Here, too, the distinctions no longer mattered. The so-called one-decision stocks that had captivated the market during the go-go 1960s—Eastman Kodak, Polaroid, Avon, and Xerox—lost their sheen and deteriorated to look just like the others.
5. Richard Bookstaber and David P. Jacob, "The Composite Hedge: Controlling the Credit Risk of High-Yield Bonds," *Financial Analysts Journal* (March–April 1986): 25–35.

CHAPTER 4 How Salomon Rolled the Dice and Lost

1. During a nine-month period beginning that month, 1-year Treasury yields increased by 275 basis points, the 5-year yields by 160 basis points, the 10-year yields by 110 basis points, and 30-year yields by 80 basis points.
2. Mortgages have withdrawn from the risk limelight in recent years because rates have been low and stable. But the option-related risk remains, and in fact is far greater now than it was in the early 1990s. One reason is the proliferation of ever more creative mortgage structures, structures that appear to be reasonably designed just because the interest rates have cooperated. A second reason is that the long period with rates at a relatively constant level means there is a concentration of gamma exposure. The nonlinearities that will accompany a rise in rates will hit a substantial proportion of the existing mortgage supply.
3. LIBOR, the London Interbank Offered Rate, is the interest rate that the banks charge each other for short-term loans and serves as the benchmark for determining the rate of short-term loans to corporations.

CHAPTER 6 Long-Term Capital Management Rides the Leverage Cycle to Hell

1. One example of the unexpected problems is the losses by the III Offshore Advisors' High Risk Opportunities fund during the Russia crisis. The fund was in fact well hedged, but one side of the position required a mark to market and the other was based on forward contracts that were not marked to market—their value was formally determined only when the contracts matured. The banks on the mark-to-market side did not recognize the other side of the position as an offset for their collateral demands. So the fund was held to be in default and had to liquidate even though from an economic standpoint they had done everything right.
2. A repo, or repurchase agreement, is a method of obtaining a loan where securities are used as collateral. With a repo, you sell a bond you own and enter into an agreement with the purchaser that you will repurchase the bond at a specified future date for a specified price. You get cash for the sale of the bond that you can use for other investments. The interest on the loan is paid by setting the price for repurchase of the bond higher than the initial sale price. For example, an agreement might be to sell a bond for \$1 million and buy it back one week later for \$1,001,000, the additional \$1,000 representing the interest rate over the week on the million dollars.
3. The call on capital and the ability to wait out a liquidity crisis are one advantage a proprietary desk in a large investment bank has over a hedge fund, at least if the investment bank secures term lending agreements. With the liquidity problems that killed Drexel Burnham Lambert still freshly in mind, John MacFarlane, who ran the finance desk at the time, pushed more of the firm's financing out to longer terms, a farsighted move that was important in saving Salomon during its Treasury scandal. Recently, Citadel, a large Chicago-based hedge fund, began issuing bonds as a source of financing to reduce its dependence on bank loans.
4. Michael Lewis, "How the Eggheads Cracked," *New York Times*, January 24, 1999.
5. The specifics of many of the strategies that led to the downfall of UBS are described in some detail in Nicholas Dunbar, *Inventing Money* (New York: John Wiley & Sons, 2000). The political intrigues that allowed for the repeated lapses in risk control are described in Dirk Schutz, *The Fall of UBS* (New York: Pyramid, 2000).
6. One example of convergence trades, possibly the biggest convergence trades of all time, were those that permeated the fixed income world in anticipation of the emergence of the European Monetary Union (EMU). Taking as a starting point the assumption that the union would take place—an assumption that was not without risk and was carefully monitored—convergence trades on the government bonds of the union members were taken

NOTES

on the assumption that they would become identical at a known future time—once the union was consummated. They might differ slightly because, like the U.S. on-the-run and off-the-run bonds, they might have differing levels of liquidity or, for some of the members, higher or lower levels of default risk, but certainly the bonds of Germany and France would be identical. Yet they did not trade that way. They did not even trade that way due to any reasonable assumption that the EMU would not go through. The Salomon European fixed income arbitrage team placed convergence trades of this nature several years before the union, and made nearly half a billion dollars on them.

CHAPTER 7 Colossus

1. “Wall Street Settlement: Excerpt from Settlement with Citigroup on Salomon’s Recommendations,” *New York Times*, April 29, 2003.
2. “Oversight of Investment Banks’ Response to the Lessons of Enron,” Volume 1, Senate Committee on Government Affairs, December 11, 2002.
3. Statement of Senator Carl Levin, Chairman, Permanent Subcommittee on Investigations, Hearing on the Role of the Financial Institutions in Enron’s Collapse, July 30, 2002.
4. The Grubman-Weill-92nd Street Y triangle was a journalist’s dream. Three articles from the *New York Times* from three successive days covering the story are: “Wall Street and the Nursery School: A New York Story,” November 14, 2002; “More Details on Messages by Ex-Analyst for Citigroup,” November 15, 2002; and “Private Preschool Admissions: Grease and the City,” November 16, 2002. A detailed treatment of Jack Grubman’s world by an analyst who worked as his counterpart at Merrill Lynch is Dan Reingold, *Confessions of a Wall Street Analyst* (New York: Collins, 2006).
5. “It’s Cleanup Time at Citi,” *New York Times*, November 7, 2004.
6. “Citigroup Told to Fix Problems Before Any Mergers,” *New York Times*, March 18, 2005.
7. John C. Coffee Jr., “What Caused Enron? A Capsule Social and Economic History of the 1990’s,” Columbia Law School, Center for Law and Economic Studies, Working Paper No. 214. Also George Moriarty and Philip Livingston, “Quantitative Measures of the Quality of Financial Reporting,” *Financial Executive* (July/August 2001), and U.S. General Accounting Office, “Report to the Chairman, Committee on Banking, Housing, and Urban Affairs, U.S. Senate, Financial Statement Restatements: Trends, Market Impacts, Regulatory Responses and Remaining Challenges” (October 2002) (GAO-03-138).
8. William Weiss and David Tinius, “Luca Pacioli: Accounting’s Renaissance Man,” *Management Accounting* (July 1991): 54–56.

CHAPTER 8 Complexity, Tight Coupling, and Normal Accidents

1. The term *normal accident* was coined by Charles Perrow in his book, *Normal Accidents: Living with High-Risk Technologies* (New York: Basic Books, 1984; an updated edition is available from Princeton University Press, 1999). Another source, which focuses on the potential for normal accidents involving nuclear arms, is Scott Sagan, *The Limits of Safety* (Princeton University Press, 1995). These two books provide the details and examples that are the basis of the discussion in this chapter.
2. For a complete description of Three Mile Island, refer to the “Report of The President’s Commission on the Accident at Three Mile Island,” U.S. Government Printing Office (October 1979). The engineering and safety issues are presented in Daniel Ford, *Three Mile Island* (New York: Penguin Books, 1982). The most complete source for the ValuJet accident is National Transportation Safety Board Aircraft Accident Report 97/06 DCA96MA054. “The Lessons of ValuJet 592” by William Langewiesche in the March 1998 *Atlantic Monthly* relates this accident as an example of a system-related normal accident.
3. With Three Mile Island, for example, there was no consensus on whether the reaction of zirconium and water under conditions of extreme heat would lead to the creation of potentially explosive hydrogen within the containment facility.

CHAPTER 9 The Brave New World of Hedge Funds

1. Jack Treynor, “Bulls, Bears and Market Bubbles,” *Financial Analysts Journal* (March/April 1998): 69–74.
2. Distinctions must be made between individual investors and professional investors whose livelihood depends on their returns. The latter group, more often than not, has a tendency to reduce risk as winnings pile on so they can protect their annual payout. Indeed, it has been my observation that one of the greatest determinants of which portfolio managers will grow to manage multibillion-dollar portfolios and which equally talented ones will remain in the hinterlands is whether they have a positive or a negative wealth effect.
3. Ironically, the tulips that were the objects of such ardor attained their prized status because they were diseased. The flaming was the result of a virus, and the virus that gave them their coveted streaks of irregular colors also weakened them, making them less likely to reproduce and therefore increasing their rarity. The virus was also the reason a tulip with flames might have offspring that did not have them; that depended on whether or not the offspring became infected with the virus. For the history of tulip mania, see Mike Dash, *Tulipomania* (New York: Crown Publishers, 1999).

NOTES

4. Eli Ofek and Matthew Richardson, "DotCom Mania: A Survey of Market Efficiency in the Internet Sector" (working paper, New York University, May 2001). Also, Eli Ofek and Matthew Richardson, "The IPO Lock-up Period: Implications for Market Efficiency and Downward Sloping Demand Curves" (working paper, New York University, March 2000).
5. Of course, the potential buyers in the used car market are not going to be stupid forever. They will discover that they cannot determine whether any particular car is good or bad, putting them at a disadvantage relative to the seller. So fewer buyers will venture into the market, and those who do will pay less for a car because of the uncertainty about its quality. The lower price will further discourage those who have good used cars from going into the market. As a result, the market will be smaller than it might otherwise be, and the average quality of the cars in the used car market will be lower. In a variation of Gresham's law, where bad money drives out good, bad cars will drive out good. Things can get pretty perverse if adverse selection goes too far. Extending this good car/bad car case to where there is a continuum of quality of cars, we might find that not only do the bad cars drive out the good cars, but the really bad cars drive out the moderately bad ones, and so on, until no market exists at all. This basic argument, presented in George Akerlof's famous paper, "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism" (*Quarterly Journal of Economics*, August 1970, pages 488–500), was the basis for his award of the Nobel Prize in 2001.
6. Carol J. Loomis's *Fortune* article, "Warren Buffett's Wild Ride at Salomon" (October 27, 1997), provides an insider's view of Warren Buffett's reaction to this scandal. A detailed report of the events is the SEC report, *In the Matter of John H. Gutfreund, Thomas W. Strauss, and John W. Meriwether*, Administrative Proceedings File No. 3-7930, 51 SEC Dkt. 93 (December 3, 1992).

CHAPTER 10 Cockroaches and Hedge Funds

1. The ideal also assumes certain technical and often arbitrary mathematical conditions to assure equilibrium. It is only recently that people have begun to consider the implications of these conditions, and it turns out that if these are relaxed, the markets can become erratic and unstable.
2. William Manchester, *A World Lit Only by Fire: The Medieval Mind and the Renaissance—Portrait of an Age* (Boston: Little Brown & Company, 1993).
3. *The Duke of Norfolk's Case, or The Doctrine of Perpetuities Fully Set Forth and Explained* (London, 1688) and Amasa Walker, *The Science of Wealth: A Manual of Political Economy* (Boston: Little Brown & Company, 1866).
4. One of the more creative techniques used to circumvent primogeniture was a process called a "common recovery." In this process, the landowner would find a pair of conspirators. One would enter a claim that he was in fact the true owner of the land, and the other would act as a "vouchee" whom the

NOTES

landowner put forward to the court to state as a witness that the landowner had originally claimed the land from the vouchee's descendants. The parties would set up the scene for the court, then leave on some pretext. When the landowner and the claimant later returned to court, they would find that the vouchee had disappeared, "departed in contempt of the Court," as it would be formally recorded. With the vouchee gone, the judgment would go against the landowner, who then would have to go after the vouchee for any recompense.

The landowner's heirs no longer had recourse to the court; any claim to the land was replaced by a claim against the absent vouchee. If the objective of this charade was for the landowner to sell the land to the claimant, then the play was over; the claimant need only make the payment to the landowner based on their prearrangement. If the objective was for the landowner to free the land of the restrictions of the perpetuity, then the claimant would convey the land back to the landowner as a fee simple, a form in which the land was now salable. Of course, the claimant could simply disavow the previous arrangements and take the land free and clear. And the vouchee remained on the hook for a claim by the landowner and his heirs in perpetuity. So it was essential that the conspirators be of the utmost trustworthiness in the first instance and of limited means in the second.

5. In some cases this had the perverse effect of increasing the average size of estates. One reason was that the population in general was not growing, and few families lasted more than three or four generations before there was no male heir. Rather than leading to the dissolution of estates, however, this increased the consolidation, as the few surviving families acquired the land of the failing lines through either marital connections or purchase. The consolidation was the most extreme in Scotland, where a third, maybe even as much as half, of the land was encumbered by perpetuities.
6. Bertrand Russell, *The Autobiography of Bertrand Russell* (London: Routledge, 2000): 147–155.
7. Edward N. Lorenz, "Deterministic Nonperiodic Flow," *Journal of the Atmospheric Sciences* 20, no. 2 (March 1963): 130–141.
8. Richard Bookstaber and Joseph Langsam, "On the Optimality of Coarse Behavior Rules," *Journal of Theoretical Biology* 116 (1985): 161–193. The article provides more examples of behavior that appears to be suboptimal given the animal's view of the world but is superior to a fine-tuned rule when unforeseeable changes in the environment are considered. A brief description of the implications that this uncertainty has for humans, specifically in their market and organizational behavior, is presented in Richard Bookstaber, "Risk Management in Complex Organizations," *Financial Analysts Journal* 55, no. 2 (March/April 1999): 18–20.
9. Tijs Goldschmidt, *Darwin's Dreampond* (Cambridge, MA: MIT Press, 1998): 26–37.

NOTES

10. Ibid., 202.
11. The ecology of Lake Victoria, as well as the Nile perch in its newfound home, turned out to be more resilient and adaptable than that of the furu. Kingfishers that had fed on the furu changed their diet to eat dagaa, a type of sardine, instead. The dagaa proliferated because their main food source, plankton, was far more abundant now that the plankton-eating furu had disappeared. Little midge flies, free of the prime predator for their larvae, seemed to reproduce without bound. They started to dance along the shoreline of the lake in dark clouds that would momentarily disperse with every gust of wind. Prawns started to appear in abundance, even in the open waters of the lake. One species of furu had been a specialist at tracking down and eating prawns, while another had competed with the prawns for the organic waste. With these furu out of the picture, the prawns blanketed the lake bottom. The Nile perch, though devoid of their supply of furu, continued to thrive, replacing their furu-rich diet with prawns and dagaa.
12. There are other flavors of risk that, like primal risk, appear difficult to address. One is what is called tail risk, the uncertainty of rare events, events that might even be so rare as to never have been observed. The second is risk where the probabilities of events are difficult to quantify. But unlike this primal risk, both of these end up fitting into the existing structure of probability theory.

Once tail risks are recognized, they can be dealt with like any other event. They just have more zeros to the right of the decimal point. Indeed, there is an area of statistics called extreme value theory that deals with measuring the probability of tail events. Tail risk has been popularized as a topic by Nassim Taleb in *Fooled by Randomness* (New York: Texere, 2001) using John Stuart Mill's example of a black swan. Before they were discovered in Australia, black swans had never been observed and so their discovery was surprising—it was a tail event. Still, while it may be that few people thought much about whether black swans might exist, the possibility of their existence, though a tail event, was within the realm of the anticipatable. To see this, suppose in medieval times a guild of artisans created decorative mantelpieces by painting stuffed swans different colors. An artisan concerned about the viability of his craft could have harbored the risk that colored swans might be discovered, bred, or created through a magic potion. Based on his view of how often new species were discovered, his familiarity with the techniques of breeding, and his belief in magic, he might have found the risk of such an event worrisome on the one hand or dismissed it as far-fetched on the other.

The topic of the hard-to-quantify risks was first addressed by the economist Frank Knight in 1921 in *Risk, Uncertainty, and Profit* (Boston: Houghton Mifflin), and is called Knightian uncertainty. Knight distinguished the uncertainty of events that are random but where the probabilities are known, as would be the case for a game of roulette, from the uncertainty of economic or political events where the probabilities cannot

NOTES

be so mechanistically determined, such as the risk of a future war in Europe. It turns out that, as with tail risks, Knightian uncertainty can fit within the existing structure of probability. We may not be able to estimate the probability of some events with the same precision as we can for games of chance, but we can, whether through objective or subjective means, come up with a probability. Even in the case of no knowledge of the likelihood of the outcomes, analysis can still proceed with an assumption of a uniform distribution (i.e., with all the possible outcomes being assigned an equal probability of occurring). The probability might itself be subject to measurement error, but no matter, for that error in the probability estimate can itself be imbedded within the final probability distribution.

13. In his extensive works, Herbert A. Simon introduced the notion of “bounded rationality” to address our tendency to behave in a less than optimal manner. In his view, we “satisfice,” doing “good enough” rather than optimizing. His argument is based on our inability to fully process all the information that is already at our disposal. By contrast, the argument I am making for coarse behavior is that, as exemplified by the cockroach, it may be better to ignore some information even if it is possible to prepare a more complete and fine-tuned response, because the result may be more robust than a response that is optimized using the currently known states of nature.
14. M. Jensen and W. Meckling, “Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure,” *Journal of Financial Economics* 3 (1976): 305–360.

CHAPTER 11 Hedge Fund Existential

1. Hedge Fund Disclosure for Institutional Investors, July 27, 2001. This was used as a framework for the International Association of Financial Engineers document on hedge fund disclosure.

CONCLUSION Built to Crash?

1. This is a point made by John Danaher in the introduction to *Brazilian Jiu-Jitsu: Theory and Technique*, by Renzo Gracie and Royler Gracie with Kid Peligro and John Danaher (Montpelier, VT: Invisible Cities Press, 2001).

INDEX

- Accidents/organizations, 159–161
- Accountants, failure (reasons), 135
- Accounting conventions, problems, 138
- Accounting orientation, 137–138
- Adaptation, best measure, 232–233
- Adverse selection, 191–192
- American depository receipts (ADRs), 68
- America Online (AOL), 139
- Amex Major Market Index (XMI) futures, 12
- Analytically driven funds, 248
- Analytical Proprietary Trading (APT), 44–45
 - initiation, 189
 - remnant, form, 190
- A Programming Language (APL), 43–47
 - asset, problem, 45
- Armstrong, Michael, 130
- Arthur Andersen, failure, 135
- Artificial markets, 229
- Asia Crisis (1997), 3, 115
- Asian currency crisis, 114
- Asian economies, 118
- Asia-Pacific Economic Cooperation (APEC), 63
- Assets
 - class, hedge fund classification, 245
 - direction, hedge fund classification, 246
- Asynchronous pricing, 225
- AT&T Wireless Services IPO, SSB underwriting, 130

- Back-office functions, 39
- Bacon, Louis, 165
- Bamberger, Gerry, 185–187, 231
- Bankers Trust
 - lawsuit, 38
 - purchase announcement, 75
- Bank exposure, 146–147
- Bank failures, 146
- Bank of Japan, objectives/strategies, 166
- Baptist Foundation, restatements/liability, 135
- Barings (bank)
 - bankruptcy, 39
 - clerical trading error, 38–39
 - derivatives cross-trading, 143

- Beard, Anson, 13
- Beder, Tanya, 204
- Behavior, economic theory, 231
- Berens, Rod, 73
- Bernard, Lewis, 42, 52
- Biggs, Barton, 11
- Black, Fischer, 9
- Black Monday (1929), 17
- Black-Scholes formula, 9, 252
- Block desk, 184–185
 - trading positions, 186
- Bond positions, hedging, 30
- Booth, David, 29
- Breakdowns, explanation, 5–6
- Broker-dealer
 - block-trading desk, usage, 184
 - price setting role, 213–214
- Bucket shop era, 177
- Buffett, Warren, 62, 99, 181, 198
 - arb unit closure, 87–88
- Bushnell, Dave, 129–131
- Butterfly effect, essence, 227

- Capital cushions, 106
- Capitalism, 250
- Cash futures, 251
 - arbitrageurs, 19, 23
 - spread, 19
 - trade, 19
- Cerullo, Ed, 41
- Cheapest-to-deliver bond, 251
- Chicago Board Options Exchange (CBOE), 252
 - Black-Scholes formula, impact, 9–10
- Citigroup
 - Associates First Capital Corporation, 128
 - consolidation, impact, 132–134
 - Japanese private banking arm, 133
 - management change, Fed reaction, 133
 - organizational complexity/structural uncertainty, 126
- Citron, Robert, 38
- Coarse behavior
 - benefits, 232–233
 - consistency, 236–237

INDEX

- Coarse behavior (*Continued*)
 - decision rules, 233
 - in humans, 235–237
 - measurement of, 238–239
 - response based on, 236
 - rules, optimality, 238
- Cockroach example, 232–233, 235
- Collateral, usage, 218
- Collateralized mortgage obligations (CMOs), 71–75, 250
- Commercial Credit, Primerica purchase, 126
- Competitive prices, 36
- Complexity
 - by-product, 143
 - implications, 156
 - importance, 144–146
- Consumer lending violations, Federal Reserve fine, 132
- Control-oriented risk management, 200
- Convergence Capital, 80
- Convergence trades, 122
- Convertible bond (CB) strategy, 57–58
- Cooke, Bill, 185–187
- Corporate defaults, possibility, 29–30
- Corporate political risk, 140
- Corrigan, Gerald, 196–198
- Countervailing trades, 213
- Credit Suisse First Boston, 72–73
- Crises, causes, 240

- da Vinci, Leonardo, 136
- Denham, Bob, 62–63, 99, 195
- Derivatives customization, 143
 - trading strategy, 30
- Deterministic nonperiodic flow, 228
- Detroit Edison, Fermi-1 experimental breeder reactor, 161–164
- Deutsche Bank, investment banking (problems), 72–73
- Dimon, Jamie, 77–78, 91, 97–98, 126
- Distressed debt, event risk, 248–249
- Dow Jones Industrial Average (DJIA), 2, 12
- Dynamic hedge, 12, 161
- Dynamic system, 228–229

- Ebbers, Bernard, 70
- Economic catastrophe, 257
- Efficient markets hypothesis, 211
- Einstein, Albert, 224–226
- Emerging market bonds, 71
- Enron
 - restatements/liability, 135
 - U.S. Senate hearings, 129–130
- Epstein, Sheldon, 46–47, 49
 - index-amortizing swap, 116
- Equity trading
 - profitability, 71–75
 - proprietary reliance, 73–74
- European Monetary System currency crisis (1992), 3
- Event risk, 248–249

- Factor exposures, 202
- Fair value basis, 29

- Federal Deposit Insurance Corporation (FDIC), 113
- Federal Reserve
 - policy shifts, 85
 - rate hike, impact, 53
- Feduniak, Bob, 42, 52
- Feuerstein, Donald, 196
- Financial instability, aspects, 3–4
- Financial markets, 224–225
- Financial risk, 256–257
- Fisher, Andy, 59, 80
- Fixed income focus, 251–252
- Fixed income research (FIR), 8–9, 43–44
- Flood, Gene, 190
- Franklin, Mark, 97
- Free-floating anxiety, 235
- Frictionless markets, 209
- FrontPoint Partners, 204, 205
- FTSE Index, 117
- Fundamental data, 166
- Furu, example, 233–235
- Futures market, 17–19
- Futures shock (1635), 175–177

- Galbraith, John Kenneth, 16
- Gamma, problems, 24–25
- General Electric (GE), 41–42
- Generally accepted accounting principles (GAAP), 135
- Geographic regions, classification, 246
- Global Crossing, restatements/liability, 135
- Godel, Kurt, 222–224, 227–228
- Gold, Jeremy, 8–9
- Goldman Sachs
 - acquisition, 75
 - public offering, delay, 109
- Goldstein, Ramy, 116–118
- Gracie family, 258–259
- Gracie, Gastao, 258
- Greenhill, Bob, 73
- Greenmail, taxation, 13–14
- Gross Domestic Product (GDP), 3–4
- Growth bias, 202
- Grubman, Jack, 128–130, 134
 - MCI/BT involvement, 69–71
 - nursery school admissions, 131–132
- Gutfreund, John, 62–63, 105, 195–197
 - resignation, 199

- Haghani, Victor, 102–104, 110, 112
- Hall, Andy, 63–67
- Hawkins, Greg, 51
- Hedgefundedness, 243–244
- Hedge funds, 165, 207, 214–215, 243
 - classification, problem, 245
 - classification, 245–246
 - control, 252–253
 - defining, 245
 - economic service, 219
 - existence, question, 244
 - regulation, 247–250
- Heisenberg, Werner, 223–228
- Hilibrand, Larry, 79, 110, 113
- Human error, 149

INDEX

- Illiquidity, cost, 217–218
- Index-amortizing swap, 46–48
- Information
 - flow, process, 210
 - implications, derivation, 170–171
 - overload, 220–230
- Information-based trading, 166
- Information Technology (IT), support function, 185–186
- Initial public offerings (IPOs), 72
 - creation, 173–174
 - issuance, amount, 178–179
- Innovation, positive effects, 255–256
- INSEAD, 66
- Intangibles, 137–138
- Interactive complexity, 154–157
- Interest only (IO), 55
- Interest rate, 84–85, 87
- International Monetary Fund (IMF) package, 103
- Internet
 - bubble, 179–181
 - businesses, virtual nature, 172
 - stocks, run-up (1998), 178
- Interrelated markets, complexity (by-product), 143
- Intraday price movement, 183
- Inventory service, 71
- Investment
 - buyers, scare, 22
 - coverage, 249–250
 - investor behavior, 203–204
 - strategy, 247
 - type, classification, 246
- Investors, irrational behavior, 203–204
- Irrational markets, impact, 180–181
- Iverson, Keith, 48
- Iverson, Ken

- Japan, liquidity, 39
- Japanese swap spread strategy/profit, 100
- Jenkinson, Robert Banks, 89–90
- Jett, Joe, 39–41
- Jiu Jitsu Academy, 258
- Jones, Paul Tudor, 165
- Junk bonds, 71

- Kaplan, Joel, 44–45
- Kaplanis, Costas, 63, 79
- Kidder, Peabody, 39–42
- Knowledge, limits, 221–230
- Krasker, Bill, 86

- Langsam, Joe, 232, 236–237
- Laplace, Pierre-Simon, 223, 225
- Lead-lag strategy, 193–194
- Leeson, Nick (impact), 38–39
- Leibowitz, Marty, 8, 51, 53
- Leland, Hayne, 10
- Leverage, 244
 - amount, reduction, 260
 - crisis, occurrence, 111–113
 - regulations, imposition, 248
- Levin, Carl, 130
- Lewis, Michael, 52
- Liquidation ability, 93

- Liquidity
 - basics, 213–220
 - complexity, relationship, 145
 - demand, 26, 191
 - hedge fund classification, 246
 - history, 217–218
 - impact, 212–213
 - needs, 183
 - providers, 213–215
 - role, 215–220
 - squeeze, prospects, 105
 - suppliers, 22, 192–193
 - supply, price elasticity, 94–95
 - transparency, 226
- Liquidity crisis
 - cycle, 93–94
 - prevention, 94–95
 - providing, hedge funds (impact), 214–215
- Long-range forecasting, 228
- London Exchange, Rothschild visit, 90
- London Interbank Offered Rate (LIBOR)
 - government rates, parity, 57
 - higher-yielding LIBOR bond, 57
 - LIBOR-denominated debt, 56
- Long-dated call options, 57
- Long-Distance Discount Service (LDDS),
 - acquisitions, 70
- Long/short equity hedge funds, 200–205
- Long-Term Capital Management (LTCM)
 - capital reserves, assumption, 106–108
 - collapse, 93
 - decision point, 110
 - disaster, 57, 60, 92–93, 100, 145
 - hedge fund debacle/crisis, 1–3
 - leverage cycle, 97
 - liquidity risk, 107–108
 - losses, 108–111
 - management, initiation, 195–200
 - market price positions, feedback, 112
 - market risks, modeling/monitoring, 111–112
 - problems, public knowledge, 104–105
 - repurchase agreement, problem, 104
 - risk arbitrage position, 107
 - risk burden, 108
- Long-term rates, short-term rates (interaction), 47
- Loops, usage/impact, 45
- Loosely coupled system, 157
- Lorenz, Edward, 227–229
 - deterministic systems, 229–230

- Mack, John, 28, 29, 35, 37
 - trader emulation, 35
- Macro data, usage, 166–167
- Macro strategies, 202
- Maeda, Mitsuyo, 258
- Margin-induced sale, 94
- Market
 - aberrations, opportunities, 122
 - breakdown, reaction, 146
 - crises, worsening (aspects), 3–4
 - cycle, basis, 169
 - decline, respite, 23–24
 - exponential growth, 17

INDEX

- Market (*Continued*)
 failures, safeguards, 239–240
 illiquidity, portfolio insurance by-product, 14
 innovation, 11–12
 makers, problem, 191–192
 regulation, 146–154
 risk, paradox, 1
 volatility, 5, 25
 vulnerability, 224–225
- Market bubbles, 168–174
- Market-to-book ratio, 138
- Marx, Karl, 250
- Marxist backward market, exploitation, 250
- Material adverse change clause, 65
- Maughan, Deryck, 59, 73–77
- MCI Communications
 British Telecom (BT), merger/trade, 63–64, 67, 128
 conclusion, 74
 EPS, decline, 70
 renegotiation, willingness, 67–68
 stock, decline, 64
- Mean-reversion analysis, 190
- Mechanical failure, 149
- Mercury Asset Management, 196
- Mergers and acquisitions (M&A)
 advice/underwriting, 33
- Meriwether, John, 52, 100, 197
 resignation, 199
- Merrill Lynch, 42
- Merton, Robert, 9, 207
- Metallgesellschaft Refining and Marketing (MGRM), oil price risk (offloading), 37–38
- Mexican Brady bond/Eurobond spread, 107
- Mexican peso crisis (1994), 3
- Miller, Heidi, 78–80, 140
- Modigliani, Franco, 208–209
- Money flows, 167
- Morgan Stanley
 APL, usage, 44–45
 Dean Witter, merger, 75
 IT department, 43
 portfolio insurance, 10–12
 risk arbitrage department, 15
 risk manager, 42
- Morgan Stanley Asset Management (MSAM), 11
- Morgan Stanley Investment Management (MSIM), 205
- Mortgage-backed securities (MBSs), 54–56, 213
- Mortgage market, 35, 54–55, 102
- Mortgages, opportunities, 35
- Mozer, Paul, 195–198
- Munger, Charlie, 62, 99, 101, 197–198
- Myojin, Sugar, 59, 63, 78–79
- Natural catastrophe, 257
- New York Stock Exchange (NYSE)
 specialists, impact, 20–21
 stock sale, 13
- Noncash exchanges, 40
- Norman Conquest, 215
- Norris, Floyd (editorial), 91–92
- O'Brien, John, 10
- One-off events, 249
- Opportunistic strategies
 birth/death cycle, 252
 history, 251
- Optimal behavior, mathematical framework, 237–240
- Options, stripping, 117
- Option theory, 24
- Orange County, bankruptcy, 38
- Organizational dysfunction, 134–136
- Pacioli, Luca, 136–137
- Pairwise stock trades, 187
- Palmedo, Peter, 17, 28–29
- Paloma Partners Management Company, 42
- Pandit, Vikram, 12
- Parets, Andy, 63–69
- Parkhurst, Charlie, 85
- Partnership model, 37
- Perfect market paradigm, 209–210
 imperfections, 210–212
 liquidity, degree, 212–213
- Phibro, Salomon acquisition, 66
- Physical processes, modeling, 229
- Platt, Bob, 7–8
- Portfolio insurance, 10–15
 market crash, 22
- Portfolio managers, loss (risk), 204–205
- Position
 disclosure, problems, 225
 transparency, increase (financial market regulator advocacy), 225
- Preference shares, illiquidity, 115
- Price convergence, 121–122
- Primal risk, 235–237
 knowledge, limits, 230–232
- Primogeniture, 215–220
 implications, 216–217
 objective, impact, 216
- Principal only (PO), 54–55
- Principia Mathematica* (Russell), 221–223
- Procter & Gamble, losses, 38
- Program trading, absence, 24–25
- Protest bids, 195–196
- Quants, 8–9, 82–84
- Quantum Fund, 180–181
- Quattrone, Frank, 72
- Rational man approach, 231
- Real assets, valuation, 137–138
- Real-world risk, 237–238
- Reed, John, 127
- Relative strength index (RSI), 190
- Relative value trades, 101–102
- Rhoades, Loeb, 125
- RISC workstations, 191
- Risk
 control, 220
 knowledge, absence, 231–232
 management, 36
 nature, variation, 249
 reduction, 185
 progress/refinement, impact, 4
 tactical usage, 200

INDEX

- Risk arbitrage, 15–16, 65, 71
- Risk Architecture, 126
- Risk-controlled relative value trading, 102–103
- Risk-management structure, 238
- Robertson, Julian, 165, 179–182
- Rosenbluth, Jeff, 59, 83
- Rosenfeld, Eric, 51, 79, 86
- Rothschild, Nathan, 88–89
 - trading strategy, 90–93
 - Waterloo, relationship, 89–90
- Rubinstein, Mark, 10
- Russell, Bertrand, 221–223
- Russia default, 103–104
- Russian short-term bonds, 103

- Salomon Brothers
 - arbitrage units, 73–74, 80–82
 - closure, 88–89
 - tracking error, problems, 86–89
 - competition, 60–61
 - fixed income trading floor, 82
 - Japanese unit, 56–62
 - July Fourth massacre, 86–89
 - mortgage position, loss, 55–56
 - organization, trader involvement, 73
 - risk arbitrage group, mortgage position, 80–81
 - Travelers purchase, 77
- Salomon North, 81, 100, 199
- Salomon Smith Barney
 - convergence trades, 120–124
 - proprietary trading, reduction, 92
 - risk management committee, 98–101
 - risk measuring/monitoring, 126
 - Travelers, interaction, 125
 - U.S. fixed income arbitrage group, 91–93
 - U.S. Treasury spread position, on-the-run/off-the-run trade, 120–124
- Samuelson, Paul, 208–209
- Santa Fe Institute, 229
- Scharf, Charlie, 123–124
- Scholes, Myron, 9, 119
- Scribe Reports, 201–202
- Securities
 - price behavior, 36
 - volatility calculation, 14
- Self-referentiality, 222
- Shaw, David E., 165, 189, 243
- Shiller, Robert, 179
- Shopkorn, Stan, 181
- Shuttle disasters, 159–161
- Simon, Jim, 165
- Singapore International Monetary Exchange (SIMEX), Osaka Exchange (matched positions), 39
- Smith Barney, 75
- Solow, Robert, 208–209
- Soros, George, 179–180
- Speculative traders, economic service, 219
- Spread trades, holding periods, 83
- Standard & Poor's (S&P) futures contracts, sale, 20
- Statistical arbitrage
 - concept, 194
 - emergence, 182–188
 - origination, 184
 - traders, 191–192
- Statistics, objective, 134–135
- Stavis, Rob, 53, 59, 70–80, 84
- Stock clearance/finance, 27–28
- StockGeneration, web site, 168
- Stock market crash (1929), 137
- Stock market crash (1987), 1–2
 - demons, 7
 - impact/meditation, 18–20
 - mechanism, understanding, 28
 - meltdown, physics, 25–30
 - postmortem, 13–14
 - warning signs, 16–18
- Strauss, Thomas, 196
 - resignation, 199
- STRIPS, 40–41
- Sunbeam, restatements/liability, 135
- SuperDOT, linkage, 189
- Swap maturity/rate level, interaction, 47
- Systems, interactive complexity, 155

- Takeovers, impact, 15–16
- Tartaglia, Nunzio, 187–189
- Thorp, Ed, 188
- Three Mile Island, 159
 - safety systems, dangers, 147–150
- Tight coupling, 143
 - accidents, 154–157
 - complexity, combination, 256
 - failure, cascade, 145–146
 - interactive complexity, relationship, 157–159
 - origination, 144–145
 - reduction, 260
- Tilley, Jim, 8–9, 46
- Time disintermediation, 21
- Toevs, Alden, 9
- Too big to fail doctrine, 113
- Top-down approach, 167
- Travelers
 - brokerage unit, 75
 - Citigroup merger, 125, 140
 - Shearson American Express, merger, 78
- Tribeca Investments, 204
- Tulip mania, 174–177
- Turnover, demand satisfaction, 173
- Two-plus model, 85–86

- U.K. tax code, change (1997), 117, 118
- Ukraine Radiological Institute, 163
- Unanticipated events, 239
- Uncertainty, implications, 257
- Uncertainty Principle, 223–227
 - behavioral analogue, 226
- Undecidability Theorem, 223–224
- Underwriting, profitability, 33–34
- Union Bank of Switzerland (UBS), 48
 - Epstein, impact, 48–49
 - Goldstein, impact, 116–118
 - LTCM impact, 118–120
 - LTCM losses, 113–116, 120
 - monetary loss, 49
 - Swiss Bank Corporation, merger, 49
- Unwind/unrealized losses, 68

INDEX

- U.S. arbitrage, losses, 100
- U.S. swap market, 92
- U.S. Treasuries
 - auction, false bid, 198–199
 - stripping/reconstitution, 40
- U.S. government issues, 123
- U.S. yield curve trades, tracking error, 84–86

- Valuation model, 60
- ValuJet, 151–154, 159
- Virtual stocks, investment, 168–169
- Vranos, Michael, 81

- Warrants, 57–59
- Waste Management, restatements/liability, 135

- Wealth
 - change, market appreciation (relationship), 178
 - dilution, prevention, 216
- Weill, Sandy, 75–80, 125–129
- Whitehead, Alfred North, 221–223
- Whitney, Arthur, 45
- Window dressing adjustments, 202–203
- Workforce, diversification, 4
- WorldCom, 70
 - SEC inquiry, 128
- World Trade Organization (WTO) accord (1999), 171

- ZBI, 203–204