Mandelbroit and Hudson
The (Mis) Behavior of Markets

Chorale: The computer "bug" as artist, opus 2. (Overleaf) Computer-generated art from Mandelbrot 1982. This design was created by a "bug" in a software program while I was investigating various fractal forms—and it nicely demonstrates the creative power of chance, in art, finance and life.

CHAPTER I

Risk, Ruin, and Reward

IN THE SUMMER OF 1998, the improbable happened.

On Wall Street, the historic bull market of the New Gay '90s was looking tired. There was no single, overwhelming problem—just a series of worries: recession in Japan, possible devaluation in China, and in Washington a president battling impeachment. Then came news that Russia, just two years earlier the world's hottest emerging market, was hitting a cash crunch. Western banks and debt-traders would suffer; a few, it later emerged, were already near ruin. So on August 4, the Dow Jones Industrial Average fell 3.5 percent. Three weeks later, as news from Moscow worsened, stocks fell again, by 4.4 percent. And then again, on August 31, by 6.8 percent. Other markets reeled: Bank bonds plummeted a third from their usual value against government bonds. The hammer blows were shocking—and for many investors, inexplicable. It was a panic, irrational and unpredictable; "the culmination of a meltdown," one analyst told the *Wall Street Journal*. It might, said another, "take a lifetime for investors to ever recoup some of those losses."

So much for conventional market wisdom. As we know now, the International Monetary Fund patched Russia, the Federal Reserve stabilized Wall Street, and the bull market ran another few years. In fact, by the conventional wisdom, August 1998 simply should never have happened; it was, according to the standard models of the financial industry, so improbable a sequence of events as to have been impossible. The standard theories, as taught in business schools around the world, would estimate the odds of that final, August 31, collapse at one in 20 million—an event that, if you traded daily for nearly 100,000 years, you would not expect to see even once. The odds of getting three such declines in the same month were even more minute: about one in 500 billion. Surely, August had been supremely bad luck, a freak accident, an "act of God" no one could have predicted. In the language of statistics, it was an "outlier" far, far, far from the normal expectation of stock trading.

Or was it? The seemingly improbable happens all the time in financial markets. A year earlier, the Dow had fallen 7.7 percent in one day. (Probability: one in 50 billion.) In July 2002, the index recorded three steep falls within seven trading days. (Probability: one in four trillion.) And on October 19, 1987, the worst day of trading in at least a century, the index fell 29.2 percent. The probability of that happening, based on the standard reckoning of financial theorists, was less than one in 10^{50}—odds so small they have no meaning. It is a number outside the scale of nature. You could span the powers of ten from the smallest subatomic particle to the breadth of the measurable universe—and still never meet such a number.

So what's new? Everyone knows: Financial markets are risky. But in the careful study of that concept, risk, lies knowledge of our world and hope of a quantitative control over it.

For more than a century, financiers and economists have been striving to analyze risk in capital markets, to explain it, to quantify it, and, ultimately, to profit from it. I believe that most of the theorists have been going down the wrong track. The odds of financial ruin in a free, global-market economy have been grossly underestimated. In

this sense, the common man is wise in his prejudice that—especially after the collapse of the Internet bubble—markets are risky. But financial theorists are not so wise. Over the past century, they devised an intricate mathematical apparatus for appraising risk. It was adopted wholesale by Wall Street in the 1970s. The likes of Merrill Lynch, Goldman Sachs, and Morgan Stanley made it a part of intricate trading strategies. They tried tuning investment portfolios to different frequencies of risk and reward, as one might tune a radio. But the financial bumps and lurches of the 1980s and 1990s have forced a rethink, among financiers as well as among economists. Black Monday of 1987, the Asian economic crisis of 1997, the Russian summer of 1998, and the bear market of 2001 to 2003—surely, many now realize, something is not right. If reward and risk make a ratio, the standard arithmetic must be wrong. The denominator, risk, is bigger than generally acknowledged; and so the outcome is bound to disappoint. Better assessment of that risk, and better understanding of how risk drives markets, is a goal of much of my work.

My life has been a study of risk. I learned about it firsthand in the brutal school of World War II, as a Polish refugee hiding in the French countryside with a borrowed identity and touched-up ration coupons, masquerading (badly) as a simple country boy in an occupied land. I faced it in my career, rejecting the safety of French academia for the intellectual wanderings of an industrial scientist in a more free-wheeling America. As a scientist, all of my research has, in one way or another, veered between the two poles of human experience: deterministic systems of order and planning, and stochastic, or random, systems of irregularity and unpredictability. My key contribution was to found a new branch of mathematics that perceives the hidden order in the seemingly disordered, the plan in the unplanned, the regular pattern in the irregularity and roughness of nature. This mathematics, called fractal geometry, has much to say in the natural sciences. It has helped model the weather, study river flows, analyze brainwaves and seismic tremors, and understand the distribution of galaxies. It was immediately embraced as an essential mathematical tool in the 1980s by "chaos" theory, the study of order in the seeming-chaos of a whirlpool or a hurricane. It is routinely used today in the realm of man-made structures, to measure Internet traffic, compress computer files, and make movies. It was the mathematical engine behind the computer animation in the movie, *Star Trek II: the Wrath of Khan*.

I believe it has much to contribute to finance, too. For forty years in fits and starts, as allowed by my personal interests, by unfolding events, and by the availability of colleagues to talk to, the development of fractal geometry has continually interacted with my studies of financial markets and economic systems. I have investigated them not as an economist or financier, but as a mathematical and experimental scientist. To me, all the power and wealth of the New York Stock Exchange or a London currency-dealing room are abstract; they are analogous to physical systems of turbulence in a sunspot or eddies in a river. They can be analyzed with the tools science already has, and new tools I keep adding to the old ones as need and ability allow. With these tools, I have analyzed how income gets distributed in a society, how stock-market bubbles form and pop, how company size and industrial concentration vary, and how financial prices move—cotton prices, wheat prices, railroad and Blue Chip stocks, dollar-yen exchange rates. I see a pattern in these price movements—not a pattern, to be sure, that will make anybody rich; I agree with the orthodox economists that stock prices are probably not predictable in any useful sense of the term. But the risk certainly does follow patterns that can be expressed mathematically and can be modeled on a computer. Thus, my research could help people avoid losing as much money as they do, through foolhardy underestimation of the risk of ruin. Thinking about markets as a scientific system, we may eventually craft a stronger financial industry and a better system of regulation.

A warning to readers here and now: Some of what I say has been embraced as economic orthodoxy in the past decade—but some of it remains contested, ridiculed, even vilified. When I publish in academic journals, as a scientist must, I often stir intense controversy. Each time, I have listened to the critics, rephrased my claims, gone back to my study to think and to my computers to analyze, and devised better, more-accurate models. Result: progress. Unavoidable side-effect: an element of complication. Indeed, I did not conceive of just one model of price variation, but several. Starting in 1963 and 1965 I devised two separate but incompatible models of behavior, succeeding at last in reconciling them in 1972. After a long detour through other fields of science, I resumed my financial research in 1997. This book guides the reader along the same winding journey of scientific discovery as I took. The goal: a better understanding of financial markets.

My oldest, best-corroborated insights now influence some of the mathematical models by which traders price options and banks evaluate risk. My scientific approach to markets has been emulated by a new generation of those who call themselves "econophysicists." And my latest models have been studied by a small but growing band of mathematicians, economists, and financiers in Zurich, Paris, London, Boston, and New York. I have no financial interest in their success or failure; I am a scientist, not a money man. But I wish them good fortune.

And I hope readers of this book, whether they agree or disagree with everything I say, will forsake, at least for a moment, the practical details of *why*. Instead, I hope they emerge from the book's pages with a greater fundamental understanding of *how* financial markets work, and of the great risk we run when we abandon our money to the winds of fortune.

The Study of Risk

There are many ways of handling risk. In the financial markets, the oldest is the simplest: "fundamental" analysis. If a stock is rising, seek the cause in a study of the company behind it, or of the industry and economy around it. Study harder, and predict the stock's next move. "Because" is the key word here: The price of a stock, bond, derivative, or currency moves "because" of some event or fact that more often than not comes from outside the market. World wheat prices rise because a heat wave desiccates Kansas or Ukraine. The dollar sinks because talk of war raises oil prices. This is all common sense. Financial newspapers thrive on it; they sell news and rank the importance of all the "becauses." Financial firms make an industry of it; they employ thousands of fundamental analysts, classified by genus into macroeconomic and sectoral, "top-down" and "bottom-up." Regulators codify and enforce it; they dictate what a company must tell its investors. The implicit assumption in all this: If one knows the cause, one can forecast the event and manage the risk.

Would it were so simple. In the real world, causes are usually obscure. Critical information is often unknown or unknowable, as when the Russian economy trembled in August 1998. It can be concealed or misrepresented, as during the Internet bubble or the Enron and Parmalat corporate scandals. And it can be misunderstood: The precise market mechanism that links news to price, cause to effect, is mysterious and seems inconsistent. Threat of war: Dollar falls. Threat of war: Dollar rises. Which of the two will actually happen? After the fact, it seems obvious; in hindsight, fundamental analysis can be reconstituted and is always brilliant. But before the fact, both outcomes may seem equally likely. So how can one base an investment strategy and a risk profile entirely on this one dubious principle: I can know more than anybody else?

In response, the financial industry has developed other tools. The second-oldest form of analysis, after fundamental, is "technical." This is a craft of recognizing patterns, real or spurious—of studying reams of price, volume, and indicator charts in search of clues to buy or sell. The language of the "chartists" is rich: head and shoulders, flags and pennants, triangles (symmetrical, ascending, or descending). The discipline, in disfavor during the 1980s, expanded in the 1990s as thousands of neophytes took to the Internet to trade stocks

and insights. It truly thrives, however, in currency markets. There, all major "forex" houses employ technical analysts to find "support points," "trading ranges," and other patterns in the tick-by-tick data of the world's biggest and fastest market. And in the fun-house mirror logic of markets, the chartists can at times be correct. Sterling/dollar quotes really can approach a level advertised by the technical analysts, and then pull back as if hitting a solid wall—or accelerate as if bursting through a barrier. But this is a confidence trick: Everybody knows that everybody else knows about the support points, so they place their bets accordingly. It beggars belief that vast sums can change hands on the basis of such financial astrology. It may work at times, but it is not a foundation on which to build a global risk-management system.

And so was born what business schools now call "modern" finance. It emerged from the mathematics of chance and statistics. The fundamental concept: Prices are not predictable, but their fluctuations can be described by the mathematical laws of chance. Therefore, their risk is measurable, and manageable. This is now orthodoxy to which I subscribe—up to a point.

Work in this field began in 1900, when a youngish French mathematician, Louis Bachelier, had the temerity to study financial markets at a time "real" mathematicians did not touch money. In the very different world of the seventeenth century, Pascal and Fermat (he of the famous "last theorem" that took 350 years to be proved) invented probability theory to assist some gambling aristocrats. In 1900, Bachelier passed over fundamental analysis and charting. Instead, he set in motion the next big wave in the field of probability theory, by expanding it to cover French government bonds. His key model, often called the "random walk," sticks very closely indeed to Pascal and Fermat. It postulates prices will go up or down with equal probability, as a fair coin will turn heads or tails. If the coin tosses follow each other very quickly, all the hue and cry on a stock or commodity exchange is literally static—white noise of the sort you hear on a radio when tuned between stations. And how much the prices vary is measurable. Most changes, 68 percent, are small moves up or down, within one "standard deviation"—a simple mathematical yardstick for measuring the scatter of data—of the mean; 95 percent should be within two standard deviations; 98 percent should be within three. Finally—this will shortly prove to be very important—extremely few of the changes are very large. If you line all these price movements up on graph paper, the histograms form a bell shape: The numerous small changes cluster in the center of the bell, the rare big changes at the edges.

The bell shape is, for mathematicians, *terra cognita*, so much so that it came to be called "normal"—implying that other shapes are "anomalous." It is the well-trodden field of probability distributions that came to be named after the great German mathematician Carl Friedrich Gauss. An analogy: The average height of the U.S. adult male population is about 70 inches, with a standard deviation around two inches. That means 68 percent of all American men are between 68 and 72 inches tall; 95 percent between 66 and 74 inches; 98 percent between 64 and 76 inches. The mathematics of the bell curve do not entirely exclude the possibility of a 12-foot giant or even someone of negative height, if you can imagine such monsters. But the probability of either is so minute that you would never expect to see one in real life. The bell curve is the pattern ascribed to such seemingly disparate variables as the height of Army cadets, IQ test scores, or—to return to Bachelier's simplest model—the returns from betting on a series of coin tosses. To be sure, at any particular time or place extraordinary patterns can result: One can have long streaks of tossing only "heads," or meet a squad of exceptionally tall or dim soldiers. But averaging over the long run, one expects to find the mean: average height, moderate intelligence, neither profit nor loss. This is not to say fundamentals are unimportant; bad nutrition can skew Army cadets towards shortness, and inflation can push bond prices down. But as we cannot predict such external influences very well, the only reliable crystal ball is a probabilistic one.

Genius, in any time or clime, is often unrecognized. Bachelier's

doctoral dissertation was largely ignored by his contemporaries. But his work was translated into English and republished in 1964, and thence was developed into a great edifice of modern economics and finance (and five Nobel Memorial Medals in economic science). A broader variant of Bachelier's thinking often goes by the title one of my doctoral students, Eugene F. Fama of the University of Chicago, gave it: the Efficient Market Hypothesis. The hypothesis holds that in an ideal market, all relevant information is already priced into a security today. One illustrative possibility is that yesterday's change does not influence today's, nor today's, tomorrow's; each price change is "independent" from the last.

With such theories, economists developed a very elaborate toolkit for analyzing markets, measuring the "variance" and "betas" of different securities and classifying investment portfolios by their probability of risk. According to the theory, a fund manager can build an "efficient" portfolio to target a specific return, with a desired level of risk. It is the financial equivalent of alchemy. Want to earn more without risking too much more? Use the modern finance toolkit to alter the mix of volatile and stable stocks, or to change the ratio of stocks, bonds, and cash. Want to reward employees more without paying more? Use the toolkit to devise an employee stock-option program, with a tunable probability that the option grants will be "in the money." Indeed, the Internet bubble, fueled in part by lavish executive stock options, may not have happened without Bachelier and his heirs.

Alas, the theory is elegant but flawed, as anyone who lived through the booms and busts of the 1990s can now see. The old financial orthodoxy was founded on two critical assumptions in Bachelier's key model: Price changes are statistically independent, and they are normally distributed. The facts, as I vehemently argued in the 1960s and many economists now acknowledge, show otherwise.

First, price changes are not independent of each other. Research over the past few decades, by me and then by others, shows that many financial price series have a "memory," of sorts. Today does, in fact, influence tomorrow. If prices take a big leap up or down now, there is a measurably greater likelihood that they will move just as violently the next day. It is not a well-behaved, predictable pattern of the kind economists prefer—not, say, the periodic up-and-down procession from boom to bust with which textbooks trace the standard business cycle. Examples of such simple patterns, periodic correlations between prices past and present, have long been observed in markets—in, say, the seasonal fluctuations of wheat futures prices as the harvest matures, or the daily and weekly trends of foreign exchange volume as the trading day moves across the globe.

My heresy is a different, fractal kind of statistical relationship, a "long memory." This is a delicate point to which a full chapter will be devoted later. For the moment, think about it by observing that different kinds of price series exhibit different degrees of memory. Some exhibit strong memory. Others have weak memory. Why this should be is not certain; but one can speculate. What a company does today—a merger, a spin-off, a critical product launch—shapes what the company will look like a decade hence; in the same way, its stock-price movements today will influence movements tomorrow. Others suggest that the market may take a long time to absorb and fully price information. When confronted by bad news, some quick-triggered investors react immediately while others, with different financial goals and longer time-horizons, may not react for another month or year. Whatever the explanation, we can confirm the phenomenon exists—and it contradicts the random-walk model.

Second, contrary to orthodoxy, price changes are very far from following the bell curve. If they did, you should be able to run any market's price records through a computer, analyze the changes, and watch them fall into the approximate "normality" assumed by Bachelier's random walk. They should cluster about the mean, or

average, of no change. In fact, the bell curve fits reality very poorly. From 1916 to 2003, the daily index movements of the Dow Jones Industrial Average do not spread out on graph paper like a simple bell curve. The far edges flare too high: too many big changes. Theory suggests that over that time, there should be fifty-eight days when the Dow moved more than 3.4 percent; in fact, there were 1,001. Theory predicts six days of index swings beyond 4.5 percent; in fact, there were 366. And index swings of more than 7 percent should come once every 300,000 years; in fact, the twentieth century saw forty-eight such days. Truly, a calamitous era that insists on flaunting all predictions. Or, perhaps, our assumptions are wrong.

The Power of Power Laws

Examine price records more closely, and you typically find a different kind of distribution than the bell curve: The tails do not become imperceptible but follow a "power law." These are common in nature. The area of a square plot of land grows by the power of two with its side. If the side doubles, the area quadruples; if the side triples, the area rises nine-fold. Another example: Gravity weakens by the inverse power of two with distance. If a spaceship doubles its distance from Earth, the gravitational pull on it falls to a fourth its original value. In economics, one classic power law was discovered by Italian economist Vilfredo Pareto a century ago. It describes the distribution of income in the upper reaches of society. That power law concentrates much more of a society's wealth among the very few; a bell curve would be more equitable, scattering incomes more evenly around an average. Now we reach one of my main findings. A power law also applies to positive or negative price movements of many financial instruments. It leaves room for many more big price swings than would the bell curve. And it fits the data for many price series. I provided the first evidence in a 1962 research report, summarized by a brief published paper. The report showed that in the distribution of cotton price movements over the past century, the tails followed a power law; there were far too many big price swings to fit a bell curve. The same report continued with wheat prices, many interest rates, and railroad stocks—in other words, all the data I could locate in dusty library corners. Since then, a similar pattern has been found in many other financial instruments.

Economics is faddish. As in many scientific fields, so in the dismal science a consensus emerges about what is right and what is wrong, what research is worthy a doctoral thesis and what is not. I have run counter-trend most of my professional career. In the 1960s, most theoretical economists were lionizing Bachelier and his heirs. The next decade, Wall Street embraced their theories. They were the intellectual foundation for stock-index funds, options exchanges, executive stock options, corporate capital-budgeting, bank risk-analysis, and much of the world financial industry as we know it today. Throughout this time, I was being heard, but as a near-lone voice denouncing the flaws in the logic. By the late 1980s and 1990s, however, I was no longer alone in seeing those flaws. The financial dislocations convinced many professional financiers that something was wrong. Warren E. Buffett, the famously successful investor and industrialist, jested that he would like to fund university chairs in the Efficient Market Hypothesis, so that the professors would train even more misguided financiers whose money he could win. He called the orthodox theory "foolish" and plain wrong. Yet none of its proponents "has ever said he was wrong, no matter how many thousands of students he sent forth misinstructed. Apparently, a reluctance to recant, and thereby to demystify the priesthood, is not limited to theologians."

However dogmatic the professors, the practical men of Wall Street did eventually open to new ideas. My principal objections—that prices do not follow the bell curve and are not independent—were heeded, and hundreds of economists and market analysts have

by now documented their validity. But despite recognition of the problem, the old methods have surprising staying-power. The "classical" formulae of Bachelier and his heirs—how to build an investment portfolio, to evaluate the financial value of a new factory, to judge the riskiness of a stock—remain on the curriculum at hundreds of business schools around the world and are a standard part of the Chartered Financial Analyst exams administered to thousands of young brokers and bankers. They remain part of the orthodoxy of Wall Street professionals, too. For instance, the "Black-Scholes" formula for valuing a Merrill or GM executive's stock options was long the gold standard; only in 2004 did U.S. regulators officially countenance other formulae. Why such reluctance to change? The old methods are easy and convenient. They work fine, it is argued, for most market conditions. It is only in the infrequent moments of high turbulence that the theory founders—and at such moments, who can guard against a hostile takeover, a bankruptcy or other financial act of God? Such reasoning, of course, is little comfort to those wiped out on one of those "improbable," violent trading days.

But the financial industry is supremely pragmatic. While it may genuflect to the old icons, it invests its research dollars in the search for newer, better gods. "Exotic" options, "guaranteed-return" products, "value-at-risk" analysis, and other Wall Street creations have all benefited from this search. Central bankers, too, are pragmatic. After years of accepting the old ways, they have been pushing since 1998 for new, more realistic mathematical models by which a bank should evaluate its risk. These so-called Basle II rules will force many banks to change the way they calculate how much capital they set aside as a cushion against financial catastrophe. In response, economists have been rushing to oblige with new ideas and new models. Many, with such unattractive names as GARCH and FIGARCH, just patch the old models. Others start from scratch, rejecting all the old assumptions. Behavioral economists study markets as B. F. Skinner studied humans: as organisms that input information and output behavior according to rules to be deduced. In this spirit, some researchers have wired professional traders to measure skin resistance, EEG patterns, and pulse rates, in search of the biological imperative behind a "buy" order. And there is computer-intensive finance. Wall Street has long been the computer industry's biggest customer, unleashing "genetic algorithms," "neural networks," and other computational techniques on the market in hopes that silicon intelligence can find profitable patterns where carbon-based life forms cannot.

This "post-modern" finance has yet to yield real success. Nobody has hit the jackpot.

A Game of Chance

So, as Lenin's revolutionary manifesto put it: What is to be done?

As preparation, play a game.

On the facing page you see four price charts of the kind you would find in a brokerage-house report, but with the identifying dates and values removed. Two of the charts are real chronicles of the price of a real financial instrument—name also removed. Two are forgeries, entirely fictitious series of numbers, generated using different theoretical models of how markets work. Ignore whether they trend up or down. Focus on how they vary from one moment to the next. Which are real? Which fake? What rules were used to draw the fake?

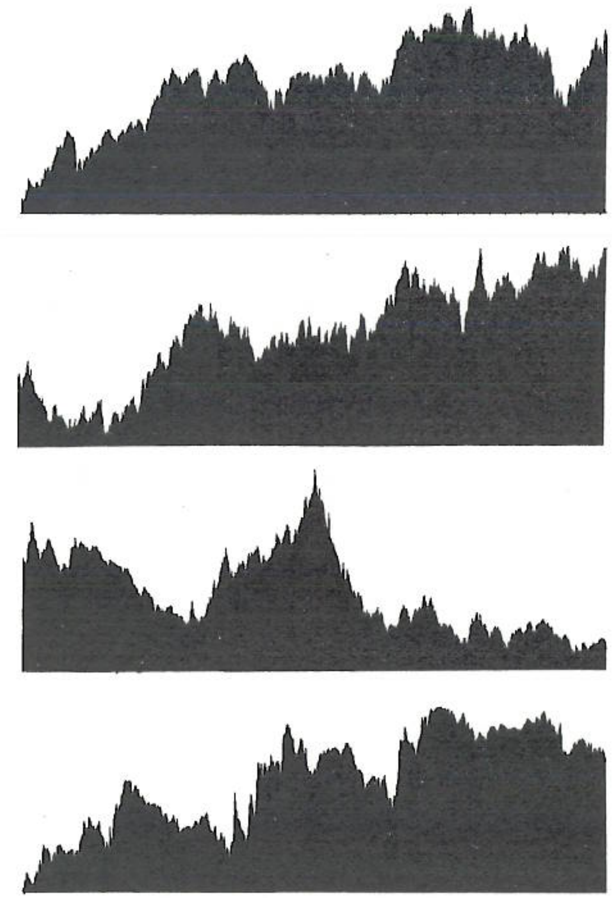

Four charts: Which are real, which are fake?

All fairly similar, many readers will say. Indeed, stripped of legends, axis labels, and other clues to context, most price "fever charts," as they are called in the financial press, look much the same. But pictures can deceive better than words.

For the truth, look at the next set of charts. These show, rather than the prices themselves, the changes in price from moment to moment. Now, a pattern emerges, and the eye is smarter than we normally give it credit for—especially at perceiving how things change.

The worst fake stands out from the rest, like a criminal in a police line-up. It is the second chart, which shows prices varying more or less uniformly over time. It was generated by the orthodox random-walk model. The size of most price changes varies within a narrow range, corresponding to the central portion of the bell curve mentioned earlier. True, the chart also shows bigger fluctuations, or outliers—but they barely stand up from the bulk of changes, as taller strands of grass rise above the average height of an unmown lawn.

Compare this fake chart with the two real ones, numbers 1 and 3. The top-most charts the relative price changes of IBM stock from 1959 to 1996; the third one charts the relative changes in the dollar/Deutschemark exchange rate. In these and all other real charts, price swings are highly erratic. The large ones are numerous and cluster together. Here, the appropriate analogy is no longer to grass, but to a forest of trees of all sizes—some gigantic. Another analogy is to the distribution of stars. They are not uniformly distributed throughout the universe. Instead they cluster into galaxies, then into galaxy clusters, in a hierarchy both random and ordered. Mathematically speaking, much the same thing is going on in these stock-price charts.

That leaves Chart No. 4—the ringer in this game. It is a fictitious series of price changes generated using my latest model of how financial markets work. It faithfully simulates the "volatile volatility" of the real charts—and, whether in financial modeling or weather forecasting, the proof of any model lies in its results. In times past, the predictions of models were expressed in a few numbers or diagrams. I pioneered the use of the computer to express the predictions of my models in this unique graphical form, a kind of forgery of reality. Here, the underlying model is called fractional Brownian motion in multifractal time. Though the name is forbid-

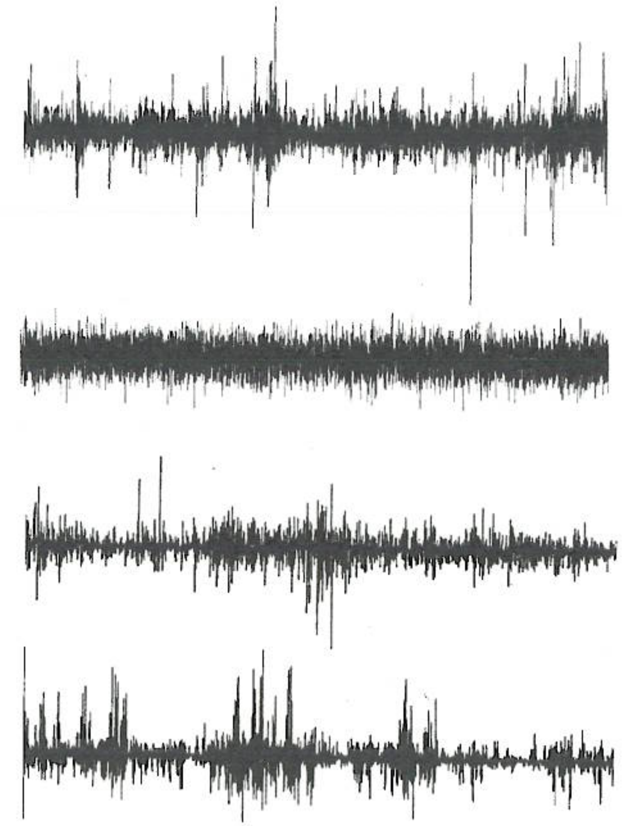

The "daily changes" in the four charts. Again, which are fake?

ding, later chapters will elaborate and show the model to be extremely parsimonious.

How does it work? It is based on my fractal mathematics, which subsequent chapters will elucidate. It is a model still in development. What I know cannot yet be used to pick stocks, trade derivatives, or value options; time, and further research by others, will determine whether it ever can. But to be able to imitate reality is a form of understanding, and as such, the multifractal model already offers some immediate insights into how markets work. Like the popular-finance press, I can boil some of them down to five "rules" of market behavior—concepts that, if grasped and acted upon, can help lessen our financial vulnerability.

Rule I. Markets are risky.

Extreme price swings are the norm in financial markets—not aberrations that can be ignored. Price movements do not follow the well-mannered bell curve assumed by modern finance; they follow a more violent curve that makes an investor's ride much bumpier. A sound trading strategy or portfolio metric would build this cold, hard fact into its foundations. Exactly how depends on the resources, talents, and stomach for risk of the individual; as ever, differing opinions make a market. But already, the mere knowledge that markets vary wildly is useful. It can be—and increasingly is—used in computer simulations to "stress-test" a portfolio, to play a wider and darker range of "what-if?" games on paper, before committing hard cash to a trading strategy. Thus, a cautious investor can build a portfolio with greater security than the standard models suggest. An aggressive trader can be better prepared to pounce on moments of high volatility. And a prudent market regulator can be more alert to urgent problems—thereby averting financial catastrophe and macroeconomic harm. Some commentators have called for a "Richter scale" of market turbulence; like that famous measure of earthquake intensity, its financial analog would rank market tremors and provide a scale for regulators to judge the severity of impending problems. Forewarned is forearmed.

Rule II. Trouble runs in streaks.

Market turbulence tends to cluster. This is no surprise to an experienced trader. In financial dealing-rooms across the world, the first fifteen minutes of trading each morning are critically important; it is when experienced traders, staring at their screens, take the tem-

perature of the market. They know that when a market opens choppily, it may well continue that way. They know that a wild Tuesday may well be followed by a wilder Wednesday. And they also know that it is in those wildest moments—the rare but recurring crises of the financial world—that the biggest fortunes of Wall Street are made and lost. They need no economists to tell them all this. But their intuition, not included in the standard model of efficient markets, is entirely validated by the multifractal model.

Rule III. Markets have a personality.

Prices are not driven solely by real-world events, news, and people. When investors, speculators, industrialists, and bankers come together in a real marketplace, a special, new kind of dynamic emerges—greater than, and different from, the sum of the parts. To use the economists' terms: In substantial part, prices are determined by *endogenous* effects peculiar to the inner workings of the markets themselves, rather than solely by the *exogenous* action of outside events. Moreover, this internal market mechanism is remarkably durable. Wars start, peace returns, economies expand, firms fail—all these come and go, affecting prices. But the fundamental process by which prices react to news does not change. A mathematician would say market processes are "stationary." This contradicts some would-be reformers of the random-walk model who explain the way volatility clusters by asserting that the market is in some way changing, that volatility varies because the pricing mechanism varies. Wrong. A striking example: My analysis of cotton prices over the past century shows the same broad pattern of price variability at the turn of the last century when prices were unregulated, as there was in the 1930s when prices were regulated as part of the New Deal.

Rule IV. Markets mislead.

Patterns are the fool's gold of financial markets. The power of chance suffices to create spurious patterns and pseudo-cycles that, for all the world, appear predictable and bankable. But a financial market is especially prone to such statistical mirages. My mathematical models can generate charts that—purely by the operation of random processes—appear to trend and cycle. They would fool any professional "chartist." Likewise, bubbles and crashes are inherent to markets. They are the inevitable consequence of the human need to find patterns in the patternless.

Rule V. Market time is relative.

There is what one may call a relativity of time in financial markets. Early on, but mostly when developing the multifractal model, I came to think of markets as operating on their own "trading time"—quite distinct from the linear "clock time" in which we normally think. This trading time speeds up the clock in periods of high volatility, and slows it down in periods of stability. Mathematically, I can write an equation showing how one time frame relates to the other and use it to generate the same kind of jagged price series that we observe in real life. This is how the successful forgery shown among the previous charts was made. It is almost as if dealing rooms need, besides the standard row of wall-clocks showing the time in Tokyo, London, and New York, a fourth clock showing "Greenwich Market Time."

This last point highlights an important subtext of this book: Market professionals know far more than they even realize. Professional traders often speak of a "fast" market or a "slow" one, depending on how they judge the volatility at that moment. They would quickly recognize, and affirm, the concept of trading time. Likewise, a bit of market folk-wisdom holds that all charts look alike: Without the identifying legends, one cannot tell if a price chart covers eighteen minutes, eighteen months, or eighteen years. This will be expressed by saying that markets scale. Even the financial press scales: There are annual reviews, quarterly bulletins, monthly newsletters, weekly magazines, daily newspapers, and tick-by-tick electronic newswires and Internet services. Market folklore and anecdote, of course, cannot confirm the multifractal model; only rigorous statistical analysis can do that. But the folklore does signal that the model is on the right track.

The multifractal model also has many implications for practical finance. As indicated, portfolio theory needs rethinking; options need revaluing; trading strategies need review. A small example: "stop-loss" orders are imperfect, to put it mildly. Many investors or traders leave instructions to close a position when a price hits a particular target. But as many have learned to their grief, when prices are really flying, they typically whiz past the target so fast that even the most attentive broker cannot execute the "sell" orders fast enough. Result: Greater losses, or smaller profit, than the investor intended. Another example: the mathematics of this model offers some potentially new yardsticks to measure volatility and risk. Instead of the standard deviations and "betas" of conventional finance, one can imagine new scales based on two new variables to be described later in this book: the H exponent of price dependence, and the α parameter characterizing volatility. A few fund managers have experimented with these concepts. They often call it chaos theory—though strictly speaking, that is marketing language riding on the coattails of a popular scientific trend. In reality, the mathematics is still young, the research barely begun, and reliable applications still distant.

So *caveat emptor*: This book will not make you rich. Bookseller: Do not put it on the same shelf with the "How to Make a Million in the Market" volumes. If it fits any genre, it is that of popular science. It explains a new, and important, way of looking at the world—in this case, the financial world. It attempts to do so using common English, with as few formulae and as little mathematical jargon as possible—or at least, with no jargon unexplained. That is because I aim to stimulate broader debate about financial-market modeling. It is a debate that has, hitherto, been confined to the rarefied circles of economics-minded mathematicians, or of mathematically inclined economists. The underlying mathematics is, frankly, forbidding—the primary reason why, when I first began publishing in the 1960s and 1970s, few mainstream economists were inclined to listen. But the extraordinary tumult and noise of this fin de siècle market turmoil are opening the ears of many who previously affected deafness.

Research in this field has far to go. It took more than sixty years after Bachelier's thesis for economists to formulate properly the Efficient Market Hypothesis, and another decade beyond that for their work to find valuable applications in the real world of zero-coupons and call options. With fractals, we are only a few short decades from the origin. But they already illumine some profound truths of finance and economics. Chief among these is the paramount importance of risk.

We have been mis-measuring risk.

Greater knowledge of a danger permits greater safety. For centuries, shipbuilders have put care into the design of their hulls and sails. They know that, in most cases, the sea is moderate. But they also know that typhoons arise and hurricanes happen. They design not just for the 95 percent of sailing days when the weather is clement, but also for the other 5 percent, when storms blow and their skill is tested. The financiers and investors of the world are, at the moment, like mariners who heed no weather warnings. This book is such a warning.

CHAPTER II

By the Toss of a Coin or the Flight of an Arrow?

FOR MOST PEOPLE, chance is a familiar but unexamined idea, a word with many separate meanings. They speak of the chance of winning the lottery, or the chance of being in a plane crash; they mean a simple number, the odds of something happening. Or they speak of a chance encounter, by which they mean unplanned, unanticipated. When they are investing, they have yet another meaning. They speak of the chance of losing money; here, chance is a menace, a risk. It is the thing that upsets their investment plans, makes them poor where they hoped to be rich. They try to weigh risks, comparing stocks with bonds, real estate with Treasuries. Most people have no idea how to do that systematically and numerically, but they accept that chance is, somehow, involved in their personal investments. Considering the alternative—that they have only themselves to blame for a lousy investment—bad luck makes a handy scapegoat.

But can chance describe not just their personal misfortunes, but the operations of the market overall? Bunk, say some. We live in the a horse moving headlong, with a wild motion that the rider could neither predict nor control. Another example: In Basque, "chance" is translated as *zoria*, a derivative of *zhar*, or bird. The flight of a bird, like the whims of a horse, cannot be predicted or controlled.

We can think of financial prices in much the same way: not predictable, not controllable. Under such circumstances, the best we can do is evaluate the odds for or against some outcome: a stock rising a certain amount this year, an option coming into the money, or an exchange rate holding steady through the next corporate budget cycle. To use the tools of probability is not to say chance governs global commerce and finance. Sure, after the fact, with enough time and effort, we can piece together a tolerable cause-and-effect story of why a price moved the way it did. But who cares? It is too late by then. Fortunes have been gained and lost. Before the fact, in the real world of fast markets, veiled motives, and uncertain outcomes, probability is the only tool at our disposal.

Chance, Simple or Complex

But how, you may ask, can the tools of probability describe the amazing richness of a stock chart?

First and foremost, random need not mean simple. There is more to probability than coins and dice. In the hands of a mathematician, even the most trivial random process—for example, a coin game—can generate surprising complexity, baroque detail, and highly structured behavior. One of the founders of modern probability theory, the late Russian mathematician Andrei Nikolaievitch Kolmogorov, wrote, "the epistemological value of probability theory is based on the fact that chance phenomena, considered collectively and on a grand scale, create a non-random regularity." Sometimes this regularity can be direct and awesome, at other times strange and wild.

For example, consider the old game of tossing a coin. It has been popular among theoreticians since the days of the Bernoulli brothers, a prolific family of eighteenth-century mathematicians from Basel whose studies helped found the field of probability. Imagine that Harry wins a Swiss franc on heads, and his brother Tom wins one on tails. (Past mathematicians called them Peter and Paul. But I could never remember which was which.) Each toss is pure luck. But after these three centuries of playing the game, millions and millions of times, each brother has every reason to expect to have won half of the time. Such is the dictate of the law of large numbers, a common-sense notion also approved by mathematicians: If you repeat a random experiment often enough, the average of the outcomes will converge towards an expected value. With a coin, heads and tails have equal odds. With a die, the side with one spot will come up about a sixth of the time. This is what Kolmogorov meant.

But other aspects of the game get more complicated. At any particular moment, one brother may have accumulated far more winnings than the other. Look at the full record of a coin-tossing experiment on the following page—10,000 simulated tosses. It is due to an eminent mathematician I knew well, Willy Feller, who in 1950 wrote a probability textbook widely used at one time. After each toss, he charted Harry's cumulative winnings or debts. An erratic, but pronounced, pattern appears: A few long, up-and-down cycles stand out, while many shorter cycles ride on top of them. The "zero-crossings"—the moments when the imaginary purses of Harry or Tom go back to the empty state at which they started—are not uniformly spread but cluster together. It is structure of an irregular kind.

All those years ago, when this diagram was first published, few readers heeded it. But I spent hours examining it, dreaming on it, trying to discern the chance patterns and processes behind it. At first glance, how much like a stock chart is this? "Chartists" spend their days studying financial graphs, spotting head-and-shoulder patterns, identifying compression periods or support levels, and then confidently advising their clients to buy or sell. Would they spot the difference if I slipped one of these coin-tossing charts into their folders? Should I expect a call from one, advising me to buy?

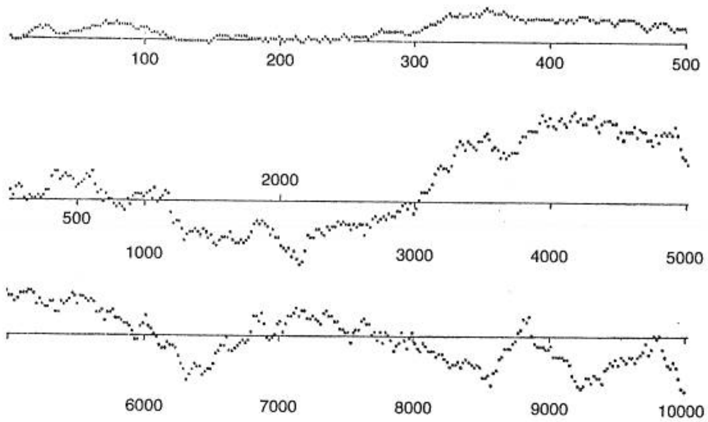

The record of 10,000 coin tosses. These charts, adapted from Feller 1950, show how far a coin-tosser's winnings can rise or fall from the expected average of zero (the horizontal lines). The top diagram shows the first 500 throws in detail. The lower two, placed end to end, cover 10,000 throws. The main point: A complex pattern can appear to emerge from even the simplest random process.

A key point in my work: Randomness has more than one "state," or form, and each, if allowed to play out on a financial market, would have a radically different effect on the way prices behave. One is the most familiar and manageable form of chance, which I call "mild." It is the randomness of a coin toss, the static of a badly tuned radio. Its classic mathematical expression is the bell curve, or "normal" probability distribution—so-called because it was long viewed as the norm in nature. Temperature, pressure, or other features of nature under study are assumed to vary only so much, and not an iota more, from the average value. At the opposite extreme is what I call "wild" randomness. This is far more irregular, more unpredictable. It is the variation of the Cornish coastline—savage promontories, craggy rocks, and unexpectedly calm bays. The fluctuation from one value to the next is limitless and frightening. In between the two extremes is a third state, which I call "slow" randomness.

Think about the three—mild, slow, and wild—as if the realm of chance were a world in its own right, with its own peculiar laws of physics. Mild randomness, then, is like the solid phase of matter: low energies, stable structures, well-defined volume. It stays where you put it. Wild randomness is like the gaseous phase of matter: high energies, no structure, no volume. No telling what it can do, where it will go. Slow randomness is intermediate between the others, the liquid state. I first proposed some of my views of chance in 1964 in Jerusalem, at an International Congress of Logic and Philosophy of Science. Since then, I have much expanded the theory and shown it to be critical to understanding financial markets in their proper light. As will be seen, the standard theories of finance assume the easier, mild form of randomness. Overwhelming evidence shows markets are far wilder, and scarier, than that.

The "Mild" Form of Chance

The most familiar type of randomness, expressed by the bell curve, first came into focus two centuries ago. From the start, its theory was both influential and controversial. Indeed, its discovery stirred a dispute over authorship—oft-told but worth repeating here—between an especially eminent mathematician, Adrien-Marie Legendre, and one of the greatest of all times, Carl Friedrich Gauss.

As the nineteenth century began, the calculation of celestial orbits was at the cutting edge of mathematical research. Improved telescopes were yielding new data on the heavens; and Newton's law of gravity provided the lens to interpret that data. But, as had been known as far back as Tycho Brahe in the late sixteenth century, telescope observations were prone to grievous error. There was the systematic error that arose from flaws in the instruments: an imperfectly ground lens, an uneven mounting. This kind of error could be explained, measured, and compensated for. Then there was the occasional error that could not be controlled: varying atmospheric conditions, tremors in the earth, or inebriated observatory assistants. This uncontrollable kind of error greatly complicated the task of calculating an orbit of a newly sighted comet or planet.

Like most great mathematicians until comparatively recent times, Legendre and Gauss had broad professional interests. Legendre in Paris rewrote Euclid's famous principles of geometry into what became a standard text in the field, wrote the first full-length treatise on number theory, and in the Napoleonic age helped precisely draw the map of France. Gauss in the north German Kingdom of Hanover (whose monarch had risen to the far richer throne in London) had been a child prodigy, a laborer's son who could count before he could speak and who developed his first famous mathematical proof, in geometry, when he was eighteen. Nearly every field he touched was the better for it: prime numbers, algebraic functions, infinite series, probability, topology. With a colleague, he designed the first electric telegraph. Like Legendre, he was a busy map surveyor. He calculated from meager data the orbits of several newly discovered planetoids. Indeed, his computational speed was legendary: The ten hours in which he determined and checked the orbit of one planetoid, Vesta, would have been, for a lesser man, several days of laborious calculation, tabular reference, and proofreading.

It was in astronomy that the two men clashed. In 1806, Legendre published a treatise on the calculation of orbits that included a supplement entitled, "On the method of least squares." It dealt with a common problem: how to find the "true" value of an orbit, or any other natural phenomenon, from a scattering of error-prone observations. The method was simple: Take a guess at the true value, and calculate how far away from it each observation is—the error. Then square each error and add them all together. Then take another guess at the true value, and see if the new squared errors are any smaller. Then do it again, and again. The "least-squares" estimate

yields errors with the smallest sum of squares; it is the value that fits closest to all the observations. It was an effective method, immediately recognized as handy and, even today, used regularly in every form of physical research from astronomy to biology. But three years after Legendre, Gauss wrote about a similar method without acknowledging the Frenchman's work. Legendre protested. Gauss was always loath to waste time quarreling with other mathematicians. He did not respond directly but assured colleagues that he had thought of the method himself when he was eighteen years old and had used it repeatedly in his astronomical calculations. Laplace tried to mediate, to no avail. In the end, both men were credited with the discovery. Proof of Gauss's priority, found later in his voluminous notebooks, is somewhat controversial, but he clearly saw much deeper meaning in the method than did Legendre.

Let us return to the coin-toss game. Say Harry or Tom keeps a record—such as Feller's diagram reproduced earlier—of the deviations from the expected average of zero. Like in tennis, divide the game into "sets," each made of one million tosses, and record how

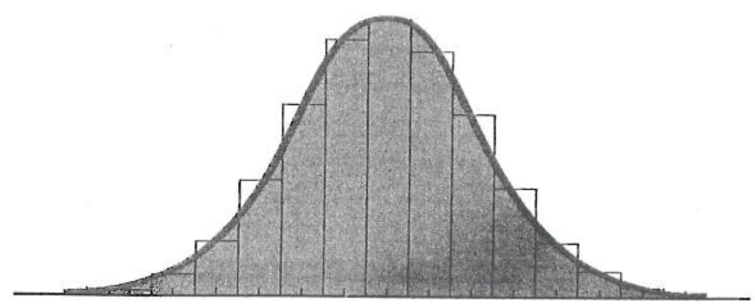

The bell curve. Harry is betting on heads coming up, and at each "set" of one million tosses keeps a record of his cumulative gain or loss. The height of the curve represents how often each type of outcome occurs. Most of the time, his winnings per set are small and get plotted into the fat center of this curve. Only occasionally, they are enormous—and appear on the skinny positive and negative "tails" of the curve. This is the distribution of a random process often called "normal."

much Harry won during the first set, the second, and so on. The size of the per-set purse will vary greatly, of course. It will often be about zero. But often, theory suggests, it will range in the favor of one brother or another—"typically," by 1,000 tosses. And on rare occasion, the "error," or deviation from the average they expect the coin to produce, will be far, far greater. If the brothers then graph the results in a "histogram" with a different-height bar for the number of times each score occurred, then the bars will start to form a familiar pattern. The numerous small winnings group around the expected average, zero—the tall center of the chart. The rare, fat purses go to the two extreme edges. Trace across the tops of all the bars, and you see the profile of the bell curve emerging.

If you study that bell curve, as did Gauss, some surprising facts arise. First, assume several games are going at once. While Harry and Tom play with the coin, their cousins are throwing dice and their friends are dealing cards. The players in each game expect a different average outcome; but for each, the graph of how their winnings per set differ from that average has the same general bell shape. Some bells may be squatter, and some narrower. But each has the same mathematical formula to describe it, and requires just two numbers to differentiate it from any other: the mean, or average, error, and the variance or standard deviation, an arbitrary yardstick that expresses how widely the bell spreads.

Now, this is all very convenient, in fact, simpler than most situations that occur in physics. One formula that includes two numbers as parameters can describe a vast range of human experience. Indeed, the common IQ test is deliberately designed to produce a bell curve of scores. The average IQ is, by definition, 100 points, the center of the bell. Then, 68 percent of the population has an IQ within one ten-point standard deviation of the mean, or between 90 and 110 points. About 95 percent are within two standard deviations of the mean, between 80 and 120 points. And 98 percent are within three standard deviations, also called sigma for the Greek letter, σ, used to write it. As sigma grows, the odds of being inside

the bell rapidly approach 100 percent, while the odds of being outside—an "outlier"—approach zero; an equation can estimate those odds. But that is not all. If you charted the IQ of every other person in a country rather than the entire population, you would still get a bell curve. If the verbal and mathematical test scores are independent and each follows a bell curve, so does the sum of the scores. Of course, the combined average score and its spread would have changed, but the basic properties of the curve would be the same.

In short, the normal curve is indestructible. It is mathematical alchemy. It is what you inevitably get if you combine lots of little variations, each one independent from the last, and each one negligible when compared to the total. No one individual matters much to the total IQ curve; no one coin toss matters much to Harry and Tom's game. But cumulatively, over time or across a population, the way the results vary forms a regular and predictable pattern. The data points are grains of sand on a shoreline, blades of grass in a lawn, electrons moving along a copper wire.

The Blindfolded Archer's Score

Now, this is a convenient way to look at the world, but is it the only way? Not at all. Late in his long life, the nineteenth-century French mathematician Augustin-Louis Cauchy thought of an especially tricky one. It was, when I was younger, viewed as interesting—but unrealistic and contrived. My work made it very real.

I think the theory best imagined in terms of an archer standing before a target painted on an infinitely long wall. He is blindfolded and consequently shoots at random, in any direction. Most of the time, of course, he misses. In fact, half of the time he shoots away from the wall, but let us not even record those cases. Now, had his recorded misses followed the mild pattern of a bell curve, most would be fairly close to the mark, and very few would be very wide of it. Suppose he shot arrows long enough, in successive "sets." For each set, he could calculate an average error and standard deviation—even give himself a score for blindfolded archery. But our archer is not in the land of the bell curve; his misses are not mild. All too often, his aim is so bad that the arrow flies almost parallel to the wall and strikes hundreds of yards from the target, or even a mile, if his arm is strong enough. Now, after each shot, let him try to work out his average target score. In the Gaussian environment, even the wildest shots have a negligible contribution to the average. After a certain number of strikes, the running average score will have settled down to one stable value, and there is practically no chance the next shot will change that average perceptibly. But the Cauchy case is completely different. The largest shot will be nearly as large as the sum of all the others. One miss by a mile completely swamps 100 shots within a few yards of the target. His scores for blindfolded archery never settle down to a nice, predictable average and a consistent variation around that average. In the language of probability, his errors do not converge to a mean. They have infinite expectation, hence also infinite variance.

Cauchy's is a totally different way of thinking of the world than Gauss's. The errors are not distributed as near-uniform grains of sand; they are a composite of grains, pebbles, boulders, and mountains. The practical importance of the distinction first became recognized through my work, but its existence was noted long ago. In 1853, the weekly printed proceedings of the French Academy of Sciences record a debate on the subject between Cauchy and another mathematician, Irénée Bienaymé. In effect, Cauchy observed that our archer's score challenged what was already by his time a casual, unreflective use of Gauss's formulae for nearly every measurement problem in science. Bienaymé retorted that the method was not merely convenient, but also reflected a fundamental truth about probability. Cauchy's fanciful error formula, he argued, was an unnatural oddity; if it ever occurred, a scientist would spot it immediately:

> The observations themselves would warn the least attentive observer. As the large-value errors would have to have a noticeably large probability, from the start they would present themselves, if not as often as the others, then at least in as large a proportion. Thus, you would have frighteningly discordant observations. And there is no doubt that they would be rejected and that the instruments, or the observation process, would be submitted to profound correction. . . . An instrument governed by such a [Cauchy] law of probability would never be put on sale by an ordinary craftsman. One could not even imagine a firm that would manufacture one.
>
> *Comptes Rendus de l'Académie des Sciences, Aug. 29, 1853*

Such has been the argument of most mathematicians and scientists ever since: Gaussian math is easy and fits most forms of reality, or so it seems. But with the sharp hindsight provided by fractal geometry, the Gaussian case begins to look not so "normal," after all. It was so-called only because science tackled it first; as ever in science, there is a healthy opportunism to begin with the problems easiest to handle. But the difference between the extremes of Gauss and of Cauchy could not be greater. They amount to two different ways of seeing the world: one in which big changes are the result of many small ones, or another in which major events loom disproportionately large. "Mild" and "wild" chance, described earlier, are my generalizations from Gauss and Cauchy.

You can see analogs of this dichotomy all around. In history, modernists argue that the course of human events is shaped by many trends, economic and social, enacted in the lives of millions of forgotten individuals; the historian's task is to trace these trends. By contrast, traditionalists, now coming back into fashion, contend that history was shaped and dominated by a few great men, Caesar or Napoleon, Newton or Einstein, for example. In the first, mild view, the birth or death of no single individual is crucial to the story of mankind; in the second, wild view, it most certainly is. Another

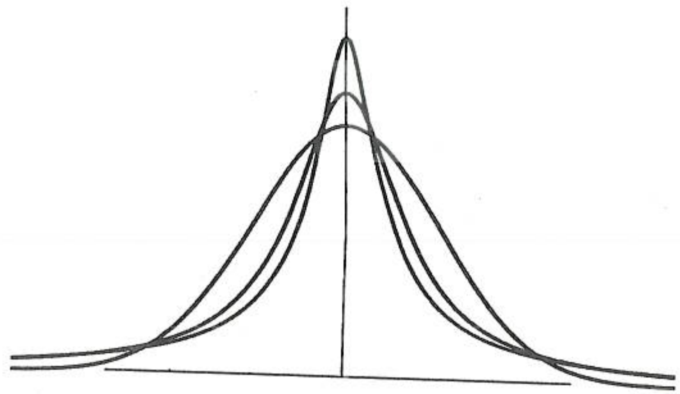

Three diagrams. In this figure, the famous bell curve is combined with two others of clearly different properties. The flattest curve is the bell, and its tails are so short that the horizontal axis had to be cut to avoid hiding them. The most peaked curve has extremely long tails. This is the Cauchy curve, that gives the distribution of the scores of our blindfolded archer. The intermediate curve will serve later in this book to represent the distribution of price increments of cotton.

example: Under a microscope, the edge of a sharp razor blade looks a bit ragged. It has random pits and bumps, but they appear to be minor imperfections on an approximately straight edge. You can easily spot the dominant trend. This is mild variation. By contrast, consider the rugged coastline of Brittany: Does it really have an "average" outline, like that of the razor blade? Only from the very great height of a satellite, where the familiar map shape can be imagined; but from closer up, in an airplane or from a tower, the tortuous, random details of promontories and bays, crags and hollows obscure the image. This coastline is wild. Yet a third example, this time in electronics. If you run a steady electrical current through a copper wire, you can "hear" it on a loudspeaker as a steady, white noise—the static of mild variation, due to the thermal excitation of the electrons. But if you try to run computer data down a very long wire, you will pick up irregular, intermittent "pops" and crackles on the line. Engineers call this 1/*f* noise, and it is the bane of

computer communications, causing transmission errors. It cannot be predicted or prevented; it can only be accommodated, with error-correcting software. That is wild variation.

Wild randomness is uncomfortable. Its mathematics is unfamiliar and in many cases remains to be developed. It looks difficult, often requiring elaborate computer simulations rather than a quick punch on a calculator. Unfortunately, the world has not been designed for the convenience of mathematicians. There is much in economics that is best described by this wilder, unpleasant form of randomness—perhaps because economics is about not just the physics of wheat, weather, and crop yields, but also the mercurial moods and unmeasurable anticipations of wheat farmers, traders, bakers, and consumers.

This makes for strange conundra. Suppose you are asked to calculate the average size of companies in the software industry. So you go down a list, counting the firms, adding up their reported revenues, and dividing one number into the other to get a simple average. But how long should the list be? Just the top fifty publicly traded firms? Every company in an industry directory? Every firm that files a tax return and says it is in software? Impossible to say: Each time you lengthen the list and add more, smaller firms, your calculated average drops. And what about Microsoft? It is the colossus of the industry, dwarfing every other firm. Try to survey the industry: If you include Microsoft in the sample, it grotesquely inflates what the survey suggests the "typical" company value is. But if you exclude it, you ignore the most important company in the industry. In short, the distribution of company size is wild—Wild West, in the view of Microsoft's critics.

Back to Finance

Now, having wandered rather far, we come to financial markets. Suppose you can simulate on your computer an artificial stock mar-

mark of a great pedagogue. On one occasion, he admonished a class that it must know the Greek alphabet perfectly—and then proceeded to forget it himself. A 1921 letter by an official in the education ministry called him "a malcontent." It said he got one of his first teaching jobs at the intercession of the ministry, in recognition of his service as a lieutenant in World War I, but over the objections of other mathematicians. He clearly lacked the political finesse required to get ahead. "He is not an eagle," the functionary wryly observed.

The shortage of university chairs in mathematics at the time meant that, on the rare occasions a vacancy occurred, a scramble ensued. In 1926 such an opening arose in Dijon, where Bachelier had worked before. His rival for the chair, Georges Cerf, was a brilliant young mathematician with all the right connections in Paris and an ally in Dijon, Maurice Gevrey, a sitting math professor. Gevrey appears to have taken a passionate dislike to Bachelier. Scouring the latter's work, Gevrey soon spotted a glaring mathematical error. When the academic committee met to decide the professorship, Gevrey brandished a letter from the eminent French probabilist, Paul Lévy in Paris, confirming the fault. Result: "Bachelier was blackballed," as Lévy ruefully recalled years later, in correspondence with me. By then, Lévy regretted the incident. He had read only the passage highlighted by Gevrey rather than the entire treatise; and in the full context of Bachelier's work the error appears benign. Lévy later apologized to Bachelier that "an impression, produced by a single initial error, should have kept me from going on with my reading of a work in which there were so many interesting ideas."

Apologies were too late, however. Bachelier's response had been immediate and intemperate. He circulated a letter (of which several copies survive) exposing the "deplorable and iniquitous" manner in which his career had been sabotaged at Dijon. His rival, he whined, was eighteen years his junior "and did not serve in the war, when at 48 I was on the front as an officer." Gevrey's partisan conduct, he

wrote, "would not astonish anyone who knew the weakness of his character." Another of the committee members, he sneered, "is well known for his ingenuity: He was able to make a vacuum of his physics course" by boring the students away. And he let fly a fusillade of more than four hundred vitriolic words at Lévy. He called Lévy's critique "violent and unjustified" and based on total ignorance of his work. The Parisian, who had just finished a book on probability, had not even bothered "opening my book" on the subject before writing his own, Bachelier complained. He concluded with an insinuation typical of the time: "Without doubt, it is inconceivable that M. Paul Lévy had wanted, by a sort of last-minute trick, to favor *un coreligionnaire*." Lévy was a Jew.

Given Bachelier's temper, it is remarkable that he ever won the security of a professorial chair—which he ultimately did, at Besançon. But that was twenty-seven years after his doctoral thesis, the work for which he is so well remembered today.

The Coin-Tossing View of Finance

The Bourse, the bustling Paris exchange, was at that time a world capital of bond trading. After the French Revolution, the government made restitution to some of the returning nobility by issuing a billion francs in perpetual bonds—a certificate that pays fixed interest indefinitely but never repays the principal. These *rentes*, as they were called, became a financial hit, widely held and actively traded: By 1900, 70 billion francs of domestic and international bonds were outstanding, compared to a government budget of four billion francs. As with U.S. Treasury bonds and U.K. Gilts today, such was the depth of the French bond market that parallel trading developed in related futures, options, and other derivatives with such exotic jargon as "call o' more's," "put o' more's," "spreads," and "contangoes." Bachelier was intimately familiar with the arcana of these markets, and he devoted part of his sixty-eight-page thesis to a detailed description of them. Indeed, some historians suggest he may have worked at the exchange for a while. His goal was to develop formulae to price these complicated derivatives. For that, he needed first to work out how the underlying bond prices themselves moved. Such "a formula which expresses the likelihood of a market fluctuation does not appear to have been published to date," he wrote. And for good reason:

> The factors that determine activity on the Exchange are innumerable, with events, current or expected, often bearing no apparent relation to price variation. Beside the somewhat natural causes for variation come artificial causes: The Exchange reacts to itself, and the current trading is a function, not only of prior trading, but also of its relationship to the rest of the market. The determination of this activity depends on an infinite number of factors: It is thus impossible to hope for mathematical forecasting. Contradictory opinions about these variations are so evenly divided that at the same instant buyers expect a rise and sellers a fall.
>
> The calculus of probability can doubtless never be applied to market activity, and the dynamics of the Exchange will never be an exact science. But it is possible to study mathematically the state of the market at a given instant—that is to say, to establish the laws of probability for price variation that the market at that instant dictates. If the market, in effect, does not predict its fluctuations, it does assess them as being more or less likely, and this likelihood can be evaluated mathematically
>
> *Opening lines of "Théorie de la Spéculation"*

Now, there had been a few earlier mathematical sorties at the market. A French stockbroker, Jules Regnault, had observed in 1863 that the longer you hold a security, the more you can win or lose on its price variations. He even worked out a formula for it. But most market analysis looked at stock and bond prices in the conventional way: Something happens and prices react—a story of cause and effect, easy to work out afterwards, difficult to forecast beforehand. But this approach was futile; one can never know everything. Instead, Bachelier tried to estimate the odds that prices will move—then, a novel approach. And he did so brilliantly by observing "a strange and unexpected" analogy between the "diffusion" of heat through a substance and how a bond price wanders up and down. Both, he saw, are processes that you cannot precisely forecast. At the level of particles in matter or of individuals in markets, the details are just too complicated; you can never discriminate and describe every relevant factor or analyze exactly how they all interrelate to spread energy or energize spreads. But in both fields, you can back away from the messy details of how or who and see the broad pattern of probability that describes the whole system. So, in the most specific of his models, Bachelier adapted the equations of one field to the problems of another.

In this model, he started by looking at the bond market as what he called "a fair game." Recall the old pastime of tossing a coin, discussed earlier. If the coin is fair, or unweighted, it is as likely to come up heads as tails on each toss. If you win a dollar for every head and lose a dollar for every tail, at the end of a string of tosses the language of probability says you should "expect" a profit of zero. Moreover, each time you toss the coin the odds of heads or tails remain 50-50, regardless of what happened on the prior toss. Put another way, a key idea behind a fair coin is that it has no memory: While you can get long runs of heads or tails, at each toss the run is as likely to end as to continue.

Now, God does not play games of chance with French bond prices; but it can certainly look that way to anybody in the thick of the trading, unable to see exactly what is driving the market up or down. And it can be described that way mathematically, with the resultant formulae used to make probabilistic statements about what could happen next. That was another of Bachelier's key insights. He assumed the split-brain thinking so common among economists today, two different ways of looking at the same event. One was after the fact or *ex post facto*, and the other before the fact or *ex ante*. After a price move, you can examine it and deduce a cause-and-effect "story" of why it happened; for instance, bond prices fell because of a new, gloomy inflation report, or because of new rumors that a big bond dealer was insolvent. But before the price move, it would have been difficult to predict those news events and even harder to forecast how the market would react. So, in your ignorance, you would have simply looked at the then-current bond prices and assumed they were fair, that the market had already taken account of all relevant information, and that prices were in equilibrium with supply matched to demand, and seller paired with buyer. Unless some new information came along to change that fine balance, you would have no reason to expect any change in price. The next move would as likely be up as down, left as right, north as south.

In effect, prices follow a random walk, the metaphor adopted by Bachelier's successors. The term comes from a quaint puzzler in probability theory. Suppose you see a blind drunk staggering across an open field. If you pass by again later on, how far will he have gotten? Well, he could go two steps left, three right, four backwards, and so on in an aimless, jagged path. On average—just as in the coin-toss game—he gets nowhere. So if you consider only that average, his random walk across the field will be forever stuck at his starting point. And that would be the best possible forecast of his future position at any time, if you had to make such a guess. The same reasoning applies to a bond price: In the absence of new information that might change the balance of supply and demand, what is the best possible forecast of the price tomorrow? Again, the price can go up or down, by big increments or small. But, with no new information to push the price decisively in one direction or another, the price on average will fluctuate around its starting point. So again, the best forecast is the price today. Moreover, each variation in price is unrelated to the last, and is generated by the same

unchanging but mysterious process that drives the markets. The price-changes, in the language of statistics, form a sequence of independent and identically distributed random variables.

In fact, it is even simpler than that, Bachelier reasoned. If you plot all of a bond's price-changes over a month or a year onto a graph, they would spread out across the paper in the familiar bell-curve shape—the many small changes clustered in the center of the bell, the few big changes at the edges. This opened the whole kit of common mathematical tools for the normal, or Gaussian, distributions mentioned earlier. And thus, through the agency of Bachelier, Gauss's theoretical curve came to be applied to the analysis of financial markets.

But Bachelier also ventured into new mathematical territory. Nearly a century before, the great French mathematician Jean Baptiste Joseph Fourier had devised equations to describe the way heat spreads. Bachelier knew the formulae well from his physics lectures. He adapted them to calculate the probability of bond prices moving up or down, and called the technique "radiation of probability." Strangely, it worked. Also, as fate would have it, very different motivations had sent other scientists on this trail. Long before, the invention of the microscope led to observations of the erratic way that tiny pollen grains jiggled about in a sample of water. A Scottish botanist, Robert Brown, studied this motion, observed that it is not a manifestation of life but a physical phenomenon, and received (possibly inflated) credit for the discovery through the term "Brownian motion." In 1905, Albert Einstein developed for it equations very similar to Bachelier's own equations of bond-price probability—though Einstein never knew that. Regardless, one cannot help but marvel that the movement of security prices, the motion of molecules, and the diffusion of heat could all be of the same mathematical species. As will be seen, it is one of many such strange liaisons in nature.

Bachelier did not stop at theory: He also tested his equations against real prices for options and futures contracts. The theories worked. For instance, he calculated that the buyer of a forty-five-day option at half a franc has 40 percent odds of earning a profit. He was uncannily close: Looking back at real trading data, he found 39 percent of such options had in fact yielded a profit to their buyers. "The market, unwittingly, obeys a law which governs it, the law of probability," he concluded.

The Efficient Market

Alas, Bachelier's economic insights went largely unnoticed for many years. In those days, finance theory was an oxymoron; finance was a distasteful trade, not a subject fit for academic inquiry. That attitude did not start changing until the Crash of 1929. Then, more economists began trying to understand financial markets. Independently of Bachelier, some started to think about a random walk. Alfred Cowles III, a wealthy investor frustrated by the imprecision of what passed for financial advice, established a foundation to gather and analyze market data. In one 1933 paper, he found what Bachelier would have predicted: Among twenty-four stock-market forecasters whom Cowles systematically studied, he found "no evidence of skill." They might as well have been shooting craps. Twenty years later, a British statistician, Maurice G. Kendall, took a long look at London shares, New York cotton, and Chicago wheat—more than a century of data—in search of conventional patterns upon which an investor could turn an easy buck or quid. "On the whole," he laconically concluded after pages of fruitless regression analysis, "I regard this experiment as a failure. . . . There is no hope of being able to predict movements on the exchange."

But it was not until 1956 that Bachelier's name reappeared in economics, this time, as an acknowledged forerunner, in a thesis on options-pricing by a student of MIT economist Paul A. Samuelson. Bachelier's idea of a "fair game" caught on; and economists recognized the practical virtues of describing markets by the laws of chance and Brownian motion. They were, in the 1960s and 1970s, put into a broader theoretical framework by Eugene F. Fama. As a student at the University of Chicago, Fama contacted me at IBM and Harvard; I became his thesis adviser, by telephone, mail, and repeated visits. His dissertation was on my views of market dynamics (of which more, later). But we often discussed Bachelier's ideas beyond the model of independent increments, and in subsequent years Fama elaborated them into what is now called the Efficient Markets Hypothesis. It is the intellectual bedrock on which orthodox financial theory today sits.

At its heart: In an ideal market, security prices fully reflect all relevant information. A financial market is a fair game in which buyer balances seller. Given that, the price at any particular moment must be the "right" one. Buyer and seller may differ in opinion; one may be a bear, and another a bull. But they both agree on the price, or there would be no deal. Multiply this thinking by the millions of daily deals of a bustling market, and you conclude that the general market price must be "right," as well—that is, that the published price reflects the market's overall best guess, given the information to hand, of what a stock is likely to profit its owner. And if that is true—and here is the bitterest pill for an investor to swallow—then you cannot beat the market.

Consider three cases. First, suppose a clever chart-reader thinks he has spotted a pattern in the old price records—say, every January, stock prices tend to rise. Can he get rich on that information, by buying in December and selling in January? Answer: No. If the market is big and efficient then others will spot the trend, too, or at least spot his trading on it. Soon, as more traders anticipate the January rally, more people are buying in December—and then, to beat the trend for a December rally, in November. Eventually, the whole phenomenon is spread out over so many months that it ceases to be noticeable. The trend has vanished, killed by its very discovery. In fact, in 1976 some economists spotted just such a pattern of regular January rallies in the stocks of small companies. Many investors close their losing positions towards the end of the year so they can book the loss as a tax deduction—and the market rebounds when they reinvest early in the new tax year. The effect is most pronounced on small stocks, which are more sensitive to small money movements. Alas, before you rush out to trade on this trend, you should know that its discovery seems to have killed it. After all the academic hoopla over it, it no longer shows up as clearly in price charts.

Second case: Suppose a financial analyst, poring over France Telecom's annual reports and chatting with its bankers and competitors, concludes that the company's debt is getting too large. To keep paying it, it will have to cut its dividend, borrow more, or sell an important asset. Can the analyst get rich on that insight? Not if the market is efficient. Other analysts will swiftly spot the problem, too, and advise their clients to sell France Telecom short. Or the bankers, who reached the conclusion first and now fear a loan default, will start charging France Telecom extra for its routine credit lines. The market will notice that, and the stock will fall. Again, the information is quickly priced into the stock.

Third and final case: Suppose the France Telecom CEO starts cashing in his stock options, because he knows the debt is a time bomb. How long can he profit on his inside information? In an efficient market, not very long. Traders will notice the captain is abandoning ship, and figure something bad is about to happen. So they sell, too, and the stock falls.

That is the theory, anyway. In all three cases—reading price charts, analyzing public information, and acting on inside information—the market quickly discounts the new information that results. Prices rise or fall to reach a new equilibrium of buyer and seller; and the next price change is, once again, as likely to be up as down. That does not mean you cannot win, ever. In fact, by the simple odds of a fair game, you can expect to win half the time and lose half the time. And if you have special insights into a stock, you

could profit from being the first in the market to act on it. But you cannot be sure you are right or first; after all, the market is full of people at least as smart as you. So, in sum, it may not be worth your while to spend all that time and money getting the information in the first place. Cheaper and safer to ride with the market. Buy a stock index fund. Relax. Be passive. Or as Samuelson at MIT put it: "They also serve who only sit and hold." His advice, then:

> A respect for evidence compels me to incline toward the hypothesis that most portfolio decision makers should go out of business—take up plumbing, teach Greek, or help produce the annual GNP by serving as corporate executives. Even if this advice to drop dead is good advice, it obviously is not counsel that will be eagerly followed. Few people will commit suicide without a push.
>
> *From* The Journal of Portfolio Management, *1974.*

A dark, nihilistic message. But then, Wall Street is nothing if not flexible—and so what could have been its epitaph was recast as a rallying cry. Bachelier's thesis was elaborated into a mature theory of how prices vary and how markets work. It came in the 1970s and 1980s to be the guiding principle for many of the standard tools of modern finance, the orthodox line, taught in business schools and shrink-wrapped into financial software packages, for how to value securities, how to build portfolios, how to measure financial performance, and how to judge the merits of a financial project. As will be seen, it is a house built on sand.

motto of one of the discipline's first reference texts, the 1890 *Principles of Economics* by Alfred Marshall. If you assume continuity, you can open the well-stocked mathematical toolkit of continuous functions and differential equations, the saws and hammers of engineering and physics for the past two centuries (and the foreseeable future). You can also draw important, useful inferences. For instance, as discussed in the preceding chapter, Markowitz's central idea was to reduce all investment decisions to two simple numbers, the mean and variance of expected prices, mathematical proxies for return and risk. In 1970 MIT's Samuelson offered a proof for Markowitz's simplification predicated on the assumption that prices change continuously.

Reality:

Clearly, prices do jump, both trivially and significantly. The trivial: Brokers often quote prices in round numbers, skipping intermediate values. Thus in the currency market, professional traders observe, about 80 percent of quotes end in a 0 or a 5, skipping the intermediate digits. The usual odds would suggest those values, being just two of the ten possible final digits in a number, should occur only about 20 percent of the time. Then there is the significant: Almost every day on the New York Stock Exchange, "order imbalances" occur in one stock or another. On one typical day, January 8, 2004, Reuters News Service reported imbalances happening eight times. Here, major news—approval of a medicine by the Food and Drug Administration, an unexpected takeover offer, or a windfall legal victory—caused market indigestion; sell and buy orders did not match, and market-makers had to raise or lower their price quotes until they did. To cope, some exchanges license "specialist" broker-dealers to step into the breach and trade when others will not. These specialists, while risking much, also profit greatly. Discontinuity, far from being an anomaly best ignored, is an essential ingredient of markets that helps set finance apart from the natural sciences.

4) Assumption: Price changes follow a Brownian motion.

Theory:

Brownian motion, again, is a term borrowed from physics for the motion of a molecule in a uniformly warm medium. Bachelier had suggested that this process can also describe price variation. Several critical assumptions come together in this idea.

First, *independence*: Each change in price—whether a five-cent uptick or a $26 collapse—appears independently from the last, and price changes last week or last year do not influence those today. That means any information that could be used to predict tomorrow's price is contained in today's price, so there is no need to study the historical charts.

A second assumption: statistical *stationarity* of the price changes. That means the process generating price changes, whatever it may be, stays the same over time. If you assume coin tosses decide prices, the coin does not get switched or weighted in the middle of the game. All that changes are the number of heads or tails as the coin is tossed; not the coin itself.

And a third assumption: the *normal* distribution. Price changes follow the proportions of the bell curve—most changes are small, an extremely few are large, in predictable and rapidly declining frequency.

Reality:

Life is more complex. This third set of assumptions is the one most clearly contradicted by the facts. Because it underpins almost every tool of modern finance, it gets special attention in the following chapter-in-a-chapter.

Pictorial Essay: Images of the Abnormal

PICTURES ARE UNDERVALUED in science. They are not trusted. That is partly the 200-year-old legacy of the French mathematicians Lagrange and Laplace, who scrupulously labored to reduce all logical thought to precise formulae and carefully chosen words; sloppy diagrams were suspect. Their motivation was, I believe, partly technological: At that time drawings were imprecise and costly, a product of human hands. But in our lifetime the computer has changed all that. A modern diagram or chart can be as precise as desired, and is no more costly than the computer that draws it. The picture can now aid, not mislead (or replace!) the scientist. It permits instant comparison, instant comprehension. Thus we begin this assault on the normal with pictures, not numbers.

Start by looking more closely at real price charts. They are so common in newspapers and on television that, by their familiarity, their intricacy can be easily missed. First, consider the most frequently published chart of all, the Dow Jones Industrial Average. It is a simple average of the stock prices of the thirty most-highly val-

ued companies in the United States. There are scores of other indices, with more or fewer stocks, varying criteria for inclusion, and different weighting systems. But the Dow, due to its age, simplicity, and wide following, is a good place to look first. It is the Mona Lisa of pictures in financial markets. So we examine it in this pictorial essay. In successive steps, we clean the years of accumulated grime from its surface to show the real information it conveys. And so we come to understand its enigmatic smile.

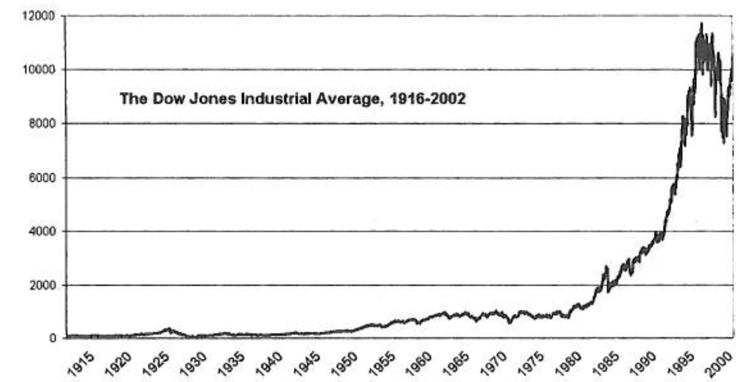

The Old Master. Here is the Dow in its most familiar form: The actual, daily index values, from 1916 through to its peak of 11,722 in January 2000 and the few years of bear market that followed.

Prominent features: Few, aside from the broad upward trend. The spike downwards, of October 19, 1987, is visible. But what stands out is the rocket rise of the 1990s. For the most part, this way of drawing the Dow makes it appear as if history did not begin until about the 1980s, when the index finally left the 1,000-mark behind.

Try another approach, to see more.

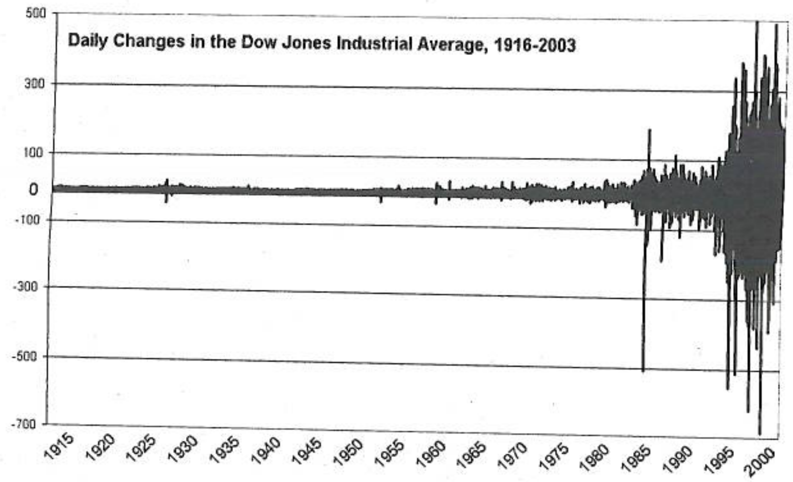

Looking closer. Here is the same data—but rather than showing the index values, this displays the index changes from one day to the next.

Prominent features: The magnitude of the index fluctuations increased towards the end of the twentieth century, as you would expect with a rapidly rising index value: Whereas in 1900 a 1 percent increase was a one-point rise in the Dow, by 2000 it was a 100-point rise. But you can see that, even in a market that rises overall, you can still get many vertiginous, one-day falls.

Now, draw the index a different way, to see more.

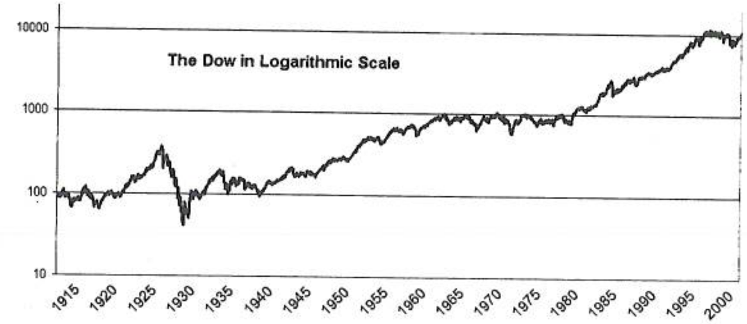

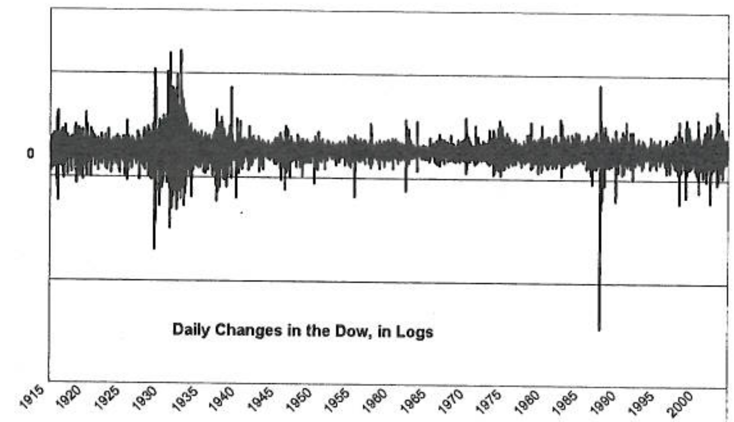

Looking under the varnish. Two charts here: the same daily index values (top chart) and changes (bottom chart) as before—but drawn to a more useful scale, the logarithmic. Logarithms rescale everything, so that a 1 percent change in 1900 will look about the same on our charts as a 1 percent change in 2000. That is just a different way of looking at the data. It makes the charts look the way the market actually felt to someone living through it.

Prominent features: The overall change in the magnitude of the index is no longer overwhelming. The Crash of 1929, the Great Depression, and World War II dominate the picture—just as they dominate our understanding of twentieth-century American economic history. Only the Crash of 1987 rivals those turbulent years. But most price changes merge into a broad strip, which varies in some sort of irregular pattern. The strip alternately narrows and widens, in some apparently haphazard cycle of thin and broad. Also, the spikes seem most likely to cluster together when the strip is wide.

Now we put the Dow to one side, and look at some new data.

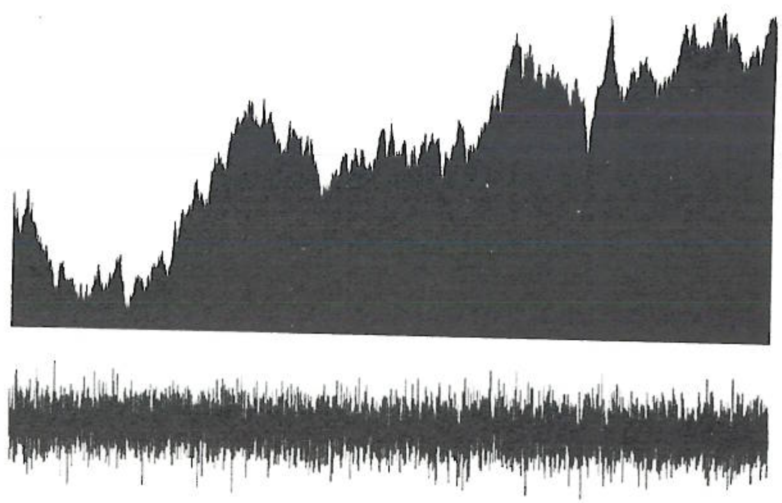

The reproduction. These two charts use the same drawing methods we applied to the Dow—but the picture is very different. These are price charts according to the Bachelier Brownian motion model. As discussed earlier, this is in the catechism of orthodox financial theory. It assumes each day's price change is independent of the last, and follows the mildly random pattern predicted by the bell curve. The top chart shows a computer-simulated Brownian price series—the silicon version of the previous, ordinary Dow chart. The bottom charts the changes from one Brownian moment to the next.

Prominent features: By "eyeball" comparison with the Dow, this is not merely different from it; it is an entirely distinct species. While the topmost chart could pass for reality, the bottommost chart is obviously aberrant. Compared to the Dow, this chart's spikes rise and fall within a small range like the blades of grass in a lawn. Its tallest spikes are interspersed across the entire chart, rather than concentrated into moments of high drama. Let us magnify the contrast, with a new scale.

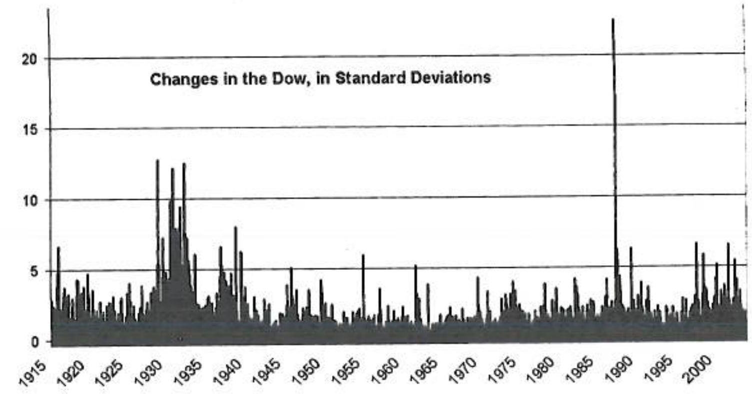

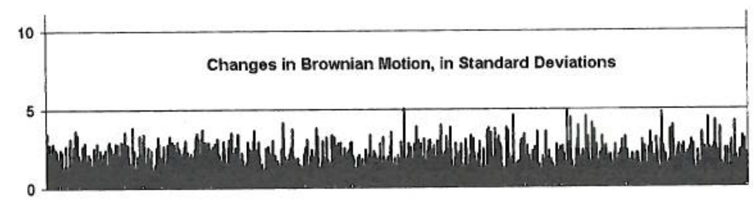

Original vs. reproduction—through the analyzer. Here you can see the differences between the Brownian (bottom) and Dow (top) charts more clearly. Instead of using a log scale as before, here we translate each index change into the number of standard deviations it is beyond the average change—in other words, how unusual it is. A very large, rare index movement will have a tall bar on this chart; the common, small changes have short bars.

Prominent features: In the Brownian chart, most changes—in fact, about 68 percent—are small. They are within one standard deviation of the average index change, zero. Mathematicians use the Greek letter sigma, σ, for standard deviation. About 95 percent of the changes are within 2σ, 98 percent within 3σ, and very, very few values are any larger. Next look at the Dow variations. The spikes are huge. Some are 10σ; one, in 1987, is 22σ. The odds of that are something less than one in 10^{50}—so minute that the standard Gaussian tables do not even contemplate it. In other words, virtually impossible. Yet there it is.