ber, and speed of the molecules and the observable frequency and magnitude of the jiggling. Einstein had, for the first time, connected new and measurable consequences to statistical physics. That might sound like a largely technical achievement, but on the contrary, it represented the triumph of a great principle: that much of the order we perceive in nature belies an invisible underlying disorder and hence can be understood only through the rules of randomness. As Einstein wrote, "It is a magnificent feeling to recognize the unity of a complex of phenomena which appear to be things quite apart from the direct visible truth."[35]

In Einstein's mathematical analysis the normal distribution again played a central role, reaching a new place of glory in the history of science. The drunkard's walk, too, became established as one of the most fundamental—and soon one of the most studied—processes in nature. For as scientists in all fields began to accept the statistical approach as legitimate, they recognized the thumbprints of the drunkard's walk in virtually all areas of study—in the foraging of mosquitoes through cleared African jungle, in the chemistry of nylon, in the formation of plastics, in the motion of free quantum particles, in the movement of stock prices, even in the evolution of intelligence over eons of time. We'll examine the effects of randomness on our own paths through life in chapter 10. But as we're about to see, though in random variation there are orderly patterns, patterns are not always meaningful. And as important as it is to recognize the meaning when it is there, it is equally important not to extract meaning when it is not there. Avoiding the illusion of meaning in random patterns is a difficult task. It is the subject of the following chapter.

CHAPTER 9

Illusions of Patterns and Patterns of Illusion

IN 1848 TWO TEENAGE GIRLS, Margaret and Kate Fox, heard unexplained noises, like knocking or the moving of furniture. Their house, it happened, had a reputation for being haunted. As the story goes,[1] Kate challenged the source of the noises to repeat the snap of her fingers and to rap out her age. It rose to both challenges. Over the next few days, with their mother's and some neighbors' assistance, the sisters worked out a code with which they could communicate with the rapper (no pun intended). They concluded that the rapping originated with the spirit of a peddler who had been murdered years earlier in the home they now occupied. With that, modern spiritualism—the belief that the dead can communicate with the living—was born. By the early 1850s a particular type of spiritual contact, called table rapping, and its cousins, table moving and table turning, had become the rage in the United States and Europe. The enterprise consisted of a group of individuals arranging themselves around a table, resting their hands upon it, and waiting. In table rapping, after some time passed, a rap would be heard. In table moving and table turning, after time passed, the table would begin to tilt or move about, sometimes dragging the sitters along with it. One pictures serious bearded men with jackets reaching their midthigh and ardent women in hoop skirts, eyes wide in wonder as their hands followed the table this way or that.

Table moving became so popular that in the summer of 1853 scientists began to look into it. One group of physicians noted that during the silent sitting period a kind of unconscious consensus seemed to form about the direction in which the table would move.[2] They found that when they diverted the sitters' attention so that a common expectation could not form, the table did not move. In another trial they managed to create a condition in which half the sitters expected the table to move to the left and half expected it to move to the right, and again it did not move. They concluded that "the motion was due to muscular action, mostly exercised unconsciously." But the definitive investigation was performed by the physicist Michael Faraday, one of the founders of electromagnetic theory, inventor of the electric motor, and one of the greatest experimental scientists in history.[3] Faraday first discovered that the phenomenon would occur even with just one subject sitting at the table. Then, enrolling subjects who were both "very honorable" and accomplished table movers, he conducted a series of ingenious and intricate experiments proving that the movement of the sitters' hands preceded that of the table. Further, he designed an indicator that alerted the subjects in real time whenever that was occurring. He found that "as soon as the . . . [indicator] is placed before the most earnest [subject] . . . the power [of the illusion] is gone; and this only because the parties are made conscious of what they are really doing."[4]

Faraday concluded, as the doctors had, that the sitters were unconsciously pulling and pushing the table. The movement probably began as random fidgeting. Then at some point the sitters perceived in the randomness a pattern. That pattern precipitated a self-fulfilling expectation as the subjects' hands followed the imagined leadership of the table. The value of his indicator, Faraday wrote, was thus "the corrective power it possesses over the mind of the table-turner."[5] Human perception, Faraday recognized, is not a direct consequence of reality but rather an act of imagination.[6]

Perception requires imagination because the data people en-

counter in their lives are never complete and always equivocal. For example, most people consider that the greatest evidence of an event one can obtain is to see it with their own eyes, and in a court of law little is held in more esteem than eyewitness testimony. Yet if you asked to display for a court a video of the same quality as the unprocessed data captured on the retina of a human eye, the judge might wonder what you were trying to put over. For one thing, the view will have a blind spot where the optic nerve attaches to the retina. Moreover, the only part of our field of vision with good resolution is a narrow area of about 1 degree of visual angle around the retina's center, an area the width of our thumb as it looks when held at arm's length. Outside that region, resolution drops off sharply. To compensate, we constantly move our eyes to bring the sharper region to bear on different portions of the scene we wish to observe. And so the pattern of raw data sent to the brain is a shaky, badly pixilated picture with a hole in it. Fortunately the brain processes the data, combining the input from both eyes, filling in gaps on the assumption that the visual properties of neighboring locations are similar and interpolating.[7] The result—at least until age, injury, disease, or an excess of mai tais takes its toll—is a happy human being suffering from the compelling illusion that his or her vision is sharp and clear.

We also use our imagination and take shortcuts to fill gaps in patterns of nonvisual data. As with visual input, we draw conclusions and make judgments based on uncertain and incomplete information, and we conclude, when we are done analyzing the patterns, that our "picture" is clear and accurate. But is it?

Scientists have moved to protect themselves from identifying false patterns by developing methods of statistical analysis to decide whether a set of observations provides good support for a hypothesis or whether, on the contrary, the apparent support is probably due to chance. For example, when physicists seek to determine whether the data from a supercollider is significant, they don't eyeball their graphs, looking for bumps that rise above the noise; they apply mathematical techniques. One such technique, significance testing, was developed in the 1920s by R. A. Fisher, one of the greatest statisticians of the twentieth century (a man also known for his uncontrollable temper and for a feud with his fellow statistics pioneer Karl Pearson that was so bitter he continued to attack his nemesis long after Pearson's death, in 1936).

To illustrate Fisher's ideas, suppose that a student in a research study on extrasensory perception predicts the result of some coin tosses. If in our observations we find that she is almost always right, we might hypothesize that she is somehow skilled at it, for instance, through psychic powers. On the other hand, if she is right about half the time, the data support the hypothesis that she was just guessing. But what if the data fall somewhere in between or if there isn't much data? Where do we draw the line between accepting and rejecting the competing hypotheses? This is what significance testing does: it is a formal procedure for calculating the probability of our having observed what we observed *if* the hypothesis we are testing is true. If the probability is low, we reject the hypothesis. If it is high, we accept it.

For example, suppose we are skeptics and hypothesize that the student cannot accurately predict the results of coin tosses. And suppose that in an experimental trial she predicts the coin tosses correctly in a certain number of cases. Then the methods we analyzed in chapter 4 allow us to calculate the probability that she could have accomplished the predictions by chance alone. If she had guessed the coin-toss results correctly so often that, say, the probability of her being that successful by chance alone is only 3 percent, then we would reject the hypothesis that she was guessing. In the jargon of significance testing, we would say the significance level of our rejection is 3 percent, meaning that the chances are at most 3 percent that by chance the data has led us astray. A 3 percent level of significance is fairly impressive, and so the media might report the feat as new evidence of the existence of psychic powers. Still, those of us who don't believe in psychic powers might remain skeptical.

This example illustrates an important point: even with data significant at, say, the 3 percent level, if you test 100 nonpsychic people for psychic abilities—or 100 ineffective drugs for their effectiveness—

you ought to expect a few people to show up as psychic or a few ineffective drugs to show up as effective. That's one reason political polls or medical studies, especially small ones, sometimes contradict earlier polls or studies. Still, significance testing and other statistical methods serve scientists well, especially when they can conduct large-scale controlled studies. But in everyday life we don't conduct such studies, nor do we intuitively apply statistical analysis. Instead, we rely on gut instinct. When my Viking stove turned out to be a lemon and ~~by chance~~ an acquaintance told me she'd had the same experience, I started telling my friends to avoid the brand. When the flight attendants on several United Airlines flights seemed grumpier than those on other airlines I'd recently flown with, I started avoiding United's flights. Not a lot of data there, but my gut instinct identified patterns.

Sometimes those patterns are meaningful. Sometimes they are not. In either case, the fact that our perception of the patterns of life is both highly convincing and highly subjective has profound implications. It implies a kind of relativity, a situation in which, as Faraday found, reality is in the eye of the beholder. For example, in 2006 *The New England Journal of Medicine* published a $12.5 million study of patients with documented osteoarthritis of the knee. The study showed that a combination of the nutritional supplements glucosamine and chondroitin is no more effective in relieving arthritis pain than a placebo. Still, one eminent doctor had a hard time letting go of his feeling that the supplements were effective and ended his analysis of the study on a national radio program by reaffirming the possible benefit of the treatment, remarking that, "One of my wife's doctors has a cat and she says that this cat cannot get up in the morning without a little dose of glucosamine and chondroitin sulfate."[8]

When we look closely, we find that many of the assumptions of modern society are based, as table moving is, on shared illusions. Whereas chapter 8 is concerned with the surprising regularities exhibited by random events, in what follows, I shall approach the issue from the opposite direction and examine how events whose patterns appear to have a definite cause may actually be the product of chance.

. . .

It is human nature to look for patterns and to assign them meaning when we find them. Kahneman and Tversky analyzed many of the shortcuts we employ in assessing patterns in data and in making judgments in the face of uncertainty. They dubbed those shortcuts heuristics. In general, heuristics are useful, but just as our manner of processing optical information sometimes leads to optical illusions, so heuristics sometimes lead to systematic error. Kahneman and Tversky called such errors biases. We all use heuristics, and we all suffer from biases. But although optical illusions seldom have much relevance in our everyday world, cognitive biases play an important role in human decision making. And so in the late twentieth century a movement sprang up to study how randomness is perceived by the human mind. Researchers concluded that "people have a very poor conception of randomness; they do not recognize it when they see it and they cannot produce it when they try,"[9] and what's worse, we routinely misjudge the role of chance in our lives and make decisions that are demonstrably misaligned with our own best interests.[10]

Imagine a sequence of events. The events might be quarterly earnings or a string of good or bad dates set up through an Internet dating service. In each case the longer the sequence, or the more sequences you look at, the greater the probability that you'll find every pattern imaginable—purely by chance. As a result, a string of good or bad quarters, or dates, need not have any "cause" at all. The point was rather starkly illustrated by the mathematician George Spencer-Brown, who wrote that in a random series of $10^{1,000,007}$ zeroes and ones, you should expect at least 10 nonoverlapping subsequences of 1 million consecutive zeros.[11] Imagine the poor fellow who bumps into one of those strings when attempting to use the random numbers for some scientific purpose. His software generates 5 zeros in a row, then 10, then 20, 1,000, 10,000, 100,000, 500,000. Would he be wrong to send back the program and ask for a refund? And how would a scientist react upon flipping open a newly pur-

chased book of random digits only to find that all the digits are zeros? Spencer-Brown's point was that there is a difference between a process being random and the product of that process appearing to be random. Apple ran into that issue with the random shuffling method it initially employed in its iPod music players: true randomness sometimes produces repetition, but when users heard the same song or songs by the same artist played back-to-back, they believed the shuffling wasn't random. And so the company made the feature "less random to make it feel more random," said Apple founder Steve Jobs.[12]

One of the earliest speculations about the perception of random patterns came from the philosopher Hans Reichenbach, who remarked in 1934 that people untrained in probability would have difficulty recognizing a random series of events.[13] Consider the following printout, representing the results of a sequence of 200 tosses of a coin, with X representing tails and O representing heads: oooo xxxxoooxxxooooxxooxoooxxxooxxoooxxxxoooxooxoxoooooxooxoooo oxxooxxxoxxoxoxxxxoooxxooxxoxooxxxooxooxoxoxxoxoooxoxooooxx xxoooxxooxoxxoooxoooxxoxooxxooooxooxxxxooooxxoooxoooxxxxxx ooxxxooxooxoooooxxxx. It is easy to find patterns in the data—for instance, the four Os followed by four Xs at the beginning and the run of six Xs toward the end. According to the mathematics of randomness, such runs are to be expected in 200 random tosses. Yet they surprise most people. As a result, when instead of representing coin tosses, strings of Xs and Os represent events that affect our lives, people seek meaningful explanations for the pattern. When a string of Xs represents down days on the stock market, people believe the experts who explain that the market is jittery. When a string of Os represents a run of accomplishments by your favorite sports star, announcers sound convincing when they drone on about the player's "streakiness." And when, as we saw earlier, the Xs or Os stood for strings of failed films made by Paramount and Columbia Pictures, everyone nodded as the industry rags proclaimed just who did and who did not have a finger on the pulse of the worldwide movie audience.

Academics and writers have devoted much effort to studying patterns of random success in the financial markets. There is much evidence, for instance, that the performance of stocks is random—or so close to being random that in the absence of insider information and in the presence of a cost to make trades or manage your portfolio, you can't profit from any deviations from randomness.[14] Nevertheless, Wall Street has a long tradition of guru analysts, and the average analyst's salary, at the end of the 1990s, was about $3 million.[15] How do those analysts do? According to a 1995 study, the eight to twelve most highly paid "Wall Street superstars" invited by *Barron's* to make market recommendations at its annual roundtable merely matched the average market return.[16] Studies in 1987 and 1997 found that stocks recommended by the prognosticators on the television show *Wall $treet Week* did much worse, lagging far behind the market.[17] And in a study of 153 newsletters, a researcher at the Harvard Institute of Economic Research found "no significant evidence of stock-picking ability."[18]

By chance alone, some analysts and mutual funds will always exhibit impressive patterns of success. And though many studies show that these past market successes are not good indicators of future success—that is, that the successes were largely just luck—most people feel that the recommendations of their stockbrokers or the expertise of those running mutual funds are worth paying for. Many people, even intelligent investors, therefore buy funds that charge exorbitant management fees. In fact, when a group of savvy students from the Wharton business school were given a hypothetical $10,000 and prospectuses describing four index funds, each composed in order to mirror the S&P 500, the students overwhelmingly failed to choose the funds with the lowest fees.[19] Since paying even an extra 1 percent per year in fees could, over the years, diminish your retirement fund by as much as one-third or even one-half, the savvy students didn't exhibit very savvy behavior.

Of course, as Spencer-Brown's example illustrates, if you look long enough, you're bound to find someone who, through sheer luck, really has made startlingly successful predictions. For those who prefer real-world examples to mathematical scenarios involving

$10^{1,000,007}$ random digits, consider the case of the columnist Leonard Koppett.[20] In 1978, Koppett revealed a system that he claimed could determine, by the end of January every year, whether the stock market would go up or down in that calendar year. His system had correctly predicted the market, he said, for the past eleven years.[21] Of course, stock-picking systems are easy to identify in hindsight; the true test is whether they will work in the future. Koppett's system passed that test too: judging the market by the Dow Jones Industrial Average, it worked for eleven straight years, from 1979 through 1989, got it wrong in 1990, and was correct again every year until 1998. But although Koppett's predictions were correct for a streak of eighteen out of nineteen years, I feel confident in asserting that his streak involved no skill whatsoever. Why? Because Leonard Koppett was a columnist for *Sporting News*, and his system was based on the results of the Super Bowl, the championship game of professional football. Whenever the team from the (original) National Football League won, the stock market, he predicted, would rise. Whenever the team from the (original) American Football League won, he predicted, the market would go down. Given that information, few people would argue that Koppett was anything but lucky. Yet had he had different credentials—and not revealed his method—he could have been hailed as the most clever analyst since Charles H. Dow.

As a counterpoint to Koppett's story, consider now the story of a fellow who does have credentials, a fellow named Bill Miller. For years, Miller maintained a winning streak that, unlike Koppett's, was compared to Joe DiMaggio's fifty-six-game hitting streak and the seventy-four consecutive victories by the *Jeopardy!* quiz-show champ Ken Jennings. But in at least one respect these comparisons were not very apt: Miller's streak earned him each year more than those other gentlemen's streaks had earned them in their lifetimes. For Bill Miller was the sole portfolio manager of Legg Mason Value Trust Fund, and in each year of his fifteen-year streak his fund beat the portfolio of equity securities that constitute the Standard & Poor's 500.

For his accomplishments, Miller was heralded "the Greatest Money Manager of the 1990s" by *Money* magazine, "Fund Manager of the Decade" by Morningstar, and one of the top thirty most influential people in investing in 2001, 2003, 2004, 2005, and 2006 by *SmartMoney*.[22] In the fourteenth year of Miller's streak, one analyst was quoted on the CNNMoney Web site as putting the odds of a fourteen-year streak by chance alone at 372,529 to 1 (more on that later).[23]

Academics call the mistaken impression that a random streak is due to extraordinary performance the hot-hand fallacy. Much of the work on the hot-hand fallacy has been done in the context of sports because in sports, performance is easy to define and measure. Also, the rules of the game are clear and definite, data are plentiful and public, and situations of interest are replicated repeatedly. Not to mention that the subject gives academics a way to attend games and pretend they are working.

Interest in the hot-hand fallacy began around 1985, in particular with a paper by Tversky and his co-workers that was published in the journal *Cognitive Psychology*.[24] In that paper, "The Hot Hand in Basketball: On the Misperception of Random Sequences," Tversky and his colleagues investigated reams of basketball statistics. The players' talent varied, of course. Some made half their shots, some more, some less. Each player also had occasional hot and cold streaks. The paper's authors asked the question, how do the number and length of the streaks compare with what you would observe if the result of each shot were determined by a random process? That is, how would things have turned out if rather than shooting baskets, the players had tossed coins weighted to reflect their observed shooting percentages? The researchers found that despite the streaks, the floor shots of the Philadelphia 76ers, the free throws of the Boston Celtics, and the experimentally controlled floor shots of the Cornell University men's and women's varsity basketball teams exhibited no evidence of nonrandom behavior.

In particular, one direct indicator of "streakiness" is the conditional probability of success (that is, making a basket) *if* on the prior attempt the player had achieved success. For a streaky player, the chance of a success on the heels of a prior success should be higher

than his or her overall chance of success. But the authors found that for each player a success following a success was just as likely as a success following a failure (that is, a missed basket).

A few years after Tversky's paper appeared, the Nobel Prize–winning physicist E. M. Purcell decided to investigate the nature of streaks in the sport of baseball.[25] As I mentioned in chapter 1, he found, in the words of his Harvard colleague Stephen Jay Gould, that except for Joe DiMaggio's fifty-six-game hitting streak, "nothing ever happened in baseball above and beyond the frequency predicted by coin-tossing models." Not even the twenty-one-game losing streak experienced at the start of the 1988 season by Major League Baseball's Baltimore Orioles. Bad players and teams have longer and more frequent streaks of failure than great players and great teams, and great players and great teams have longer and more frequent streaks of success than lesser players and lesser teams. But that is because their average failure or success rate is higher, and the higher the average rate, the longer and more frequent are the streaks that randomness will produce. To understand these events, you need only to understand the tossing of coins.

What about Bill Miller's streak? That a streak like Miller's could result from a random process may seem less shocking in light of a few other statistics. For instance, in 2004 Miller's fund gained just under 12 percent while the average stock in the S&P gained more than 15 percent.[26] It might sound like the S&P trounced Miller that year, but actually he counted 2004 in his "win" column. That is because the S&P 500 is not the simple average of the prices of the stocks it comprises; it is a weighted average in which stocks exert influence proportional to each company's capitalization. Miller's fund did worse than the simple average of S&P stocks but better than that weighted average. Actually, there were more than thirty twelve-month periods during his streak in which he lost to the weighted average, but they weren't calendar years, and the streak was based on the intervals from January 1 to December 31.[27] So the streak in a sense was an artificial one to start with, one that by chance was defined in a manner that worked for Miller.

But how can we reconcile these revelations with those 372,529-to-1 odds against him? In discussing Miller's streak in 2003, writers for *The Consilient Observer* newsletter (published by Credit Suisse–First Boston) said that "no other fund has ever outperformed the market for a dozen consecutive years in the last 40 years." They raised the question of the probability of a fund's accomplishing that by chance and went on to give three estimates of that probability (the year being 2003, they referred to the chances of a fund's beating the market for only twelve consecutive years): 1 in 4,096, 1 in 477,000, and 1 in 2.2 billion.[28] To paraphrase Einstein, if their estimates were correct, they would have needed only one. What were the actual chances? Roughly 3 out of 4, or 75 percent. That's quite a discrepancy, so I'd better explain.

Those who quoted the low odds were right in one sense: if you had singled out Bill Miller *in particular* at the start of 1991 *in particular* and calculated the odds that by pure chance *the specific person* you selected would beat the market *for precisely the next fifteen years,* then those odds would indeed have been astronomically low. You would have had the same odds against you if you had flipped a coin once a year for fifteen years with the goal of having it land heads up each time. But as in the Roger Maris home run analysis, those are not the relevant odds because there are thousands of mutual fund managers (over 6,000 currently), and there were many fifteen-year periods in which the feat could have been accomplished. So the relevant question is, if thousands of people are tossing coins once a year and have been doing so for decades, what are the chances that one of them, for some period of fifteen years or longer, will toss all heads? That probability is far, far higher than the odds of simply tossing fifteen heads in a row.

To make this explanation concrete, suppose 1,000 fund managers—certainly an underestimate—had each tossed a coin once a year starting in 1991 (the year Miller began his streak). After the first year about half of them would have tossed heads; after two years about one-quarter of them would have tossed two heads; after the

third year one-eighth of them would have tossed three heads; and so on. By then some who had tossed tails would have started to drop out of the game, but that wouldn't affect the analysis because they had already failed. The chances that, after fifteen years, a *particular coin tosser* would have tossed all heads are then 1 in 32,768. But the chances that *someone among the 1,000* who had started tossing coins in 1991 would have tossed all heads are much higher, about 3 percent. Finally, there is no reason to consider only those who started tossing coins in 1991—the fund managers could have started in 1990 or 1970 or any other year in the era of modern mutual funds. Since the writers for *The Consilient Observer* used forty years in their discussion, I calculated the odds that by chance *some manager* in the last four decades would beat the market each year for *some period of fifteen years or longer.* That latitude increased the odds again, to the probability I quoted earlier, almost 3 out of 4. So rather than being surprised by Miller's streak, I would say that if no one had achieved a streak like Miller's, you could have legitimately complained that all those highly paid managers were performing worse than they would have by blind chance!

I've cited some examples of the hot-hand fallacy in the context of sports and the financial world. But in all aspects of our lives we encounter streaks and other peculiar patterns of success and failure. Sometimes success predominates, sometimes failure. Either way it is important in our own lives to take the long view and understand that streaks and other patterns that don't appear random can indeed happen by pure chance. It is also important, when assessing others, to recognize that among a large group of people it would be very odd if one of them *didn't* experience a long streak of successes or failures.

No one credited Leonard Koppett for his lopsided successes, and no one would credit a coin tosser. Many people did credit Bill Miller. In his case, though the type of analysis I performed seems to have escaped many of the observers quoted in the media, it is no news to those who study Wall Street from the academic perspective. For example, the Nobel Prize–winning economist Merton Miller (no relation to Bill) wrote, "If there are 10,000 people looking at the stocks and trying to pick winners, one in 10,000 is going to score, by chance alone, and that's all that's going on. It's a game, it's a chance operation, and people think they are doing something purposeful but they're really not."[29] We must all draw our own conclusions depending on the circumstances, but with an understanding of how randomness operates, at least our conclusions need not be naive.

IN THE PRECEDING I've discussed how we can be fooled by the patterns in random sequences that develop over time. But random patterns in space can be just as misleading. Scientists know that one of the clearest ways to reveal the meaning of data is to display them in some sort of picture or graph. When we see data exhibited in this manner, meaningful relationships that we would likely have missed are often made obvious. The cost is that we also sometimes perceive patterns that in reality have no meaning. Our minds are made that way—to assimilate data, fill in gaps, and look for patterns. For example, look at the following arrangement of grayish squares in the figure below.

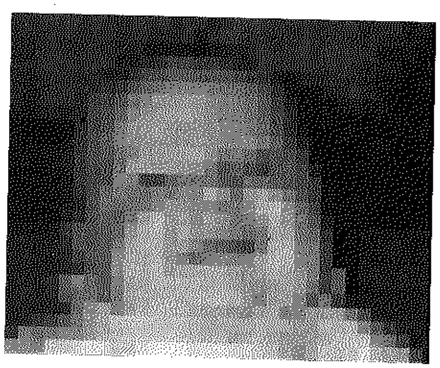

Photo from Frank H. Durgin, "The Tinkerbell Effect," *Journal of Consciousness Studies* 9, nos. 5–6 (May to June 2002).

The image does not literally look like a human being. Yet you can make enough sense of the pattern that if you saw in person the baby pictured, you would probably recognize it. And if you hold this book at arm's length and squint, you might not even perceive the imperfections in the image. Now look at this pattern of Xs and Os:

ooooxxxxoooxxxooooxxooxoooxxxooxxooo**xxxx**
oooxooxoxoooooxooxoooooxxooxxxoxxoxo**xxxx**
oooxxooxxoxooxxxooxooxoxoxxox**ooo**xoxoooox
xxxoooxxooxoxxoooxoooxxoxooxx**oooo**xoo**xxxx**
ooooxxxoooxoooxxxxxxooxxxooxooxooooo**xxxx**

Here we see rectangular clusters, especially in the corners. I have put them in boldface. If the Xs and Os represented events of interest, we might be tempted to wonder if those clusters signified something. But any meaning we assigned them would be misconceived because these data are identical to the earlier set of 200 random Xs and Os, except for the geometric 5-by-40 arrangement and the choice of which letters to put in boldface.

This very issue drew much attention toward the end of World War II, when V2 rockets started raining down on London. The rockets were terrifying, traveling at over five times the speed of sound, so that one heard them approach only after they had hit. Newspapers soon published maps of the impact sites, which seemed to reveal not random patterns but purposeful clusters. To some observers the clusters indicated a precision in the control of the rockets' flight path that, given the distance the rockets had to travel, suggested that German technology was much more advanced than anyone had dreamed possible. Civilians speculated that the areas spared were home to German spies. Military leaders worried that the Germans could target crucial military sites, with devastating consequences.

In 1946 a mathematical analysis of the bombing data was published in the *Journal of the Institute of Actuaries*. Its author, R. D. Clarke, divided the area of interest into 576 parcels half a kilometer on each side. Of these, 229 parcels sustained no hits while, despite their minuscule size, 8 parcels had four or five hits. Still, Clarke's analysis showed that, like the coin-toss data above, the overall pattern was consistent with a random distribution.[30]

Similar issues arise frequently in reports of cancer clusters. If you divide any city or county into parcels and randomly distribute incidents of cancer, some parcels will receive less than average and some more. In fact, according to Raymond Richard Neutra, chief of the Division of Environmental and Occupational Disease Control in California's Department of Health, given a typical cancer registry—a database on local rates for dozens of different cancers—for California's 5,000 census tracts, you could expect to find 2,750 with statistically significant but random elevations of some form of cancer.[31] And if you look at a large enough number of such parcels, you'll find some regions in which cancer occurred at many times the normal rate.

The picture looks even worse if you draw the parcel boundaries *after* the cancers are distributed. What you get then is called the sharpshooter effect, after the apocryphal fellow who excels in his aim because he shoots at blank paper and draws the target afterward. Unfortunately that is how it usually happens in practice: first some citizens notice neighbors with cancer; then they define the boundaries of the area at issue. Thanks to the availability of data on the Internet, America these days is being scoured for such clusters. Not surprisingly, they are being found. Yet the development of cancer requires successive mutations. That means very long exposure and/or highly concentrated carcinogens. For such clusters of cancer to develop from environmental causes and show themselves in concert and before the victims have moved away from the affected area is quite a long shot. According to Neutra, to produce the kind of cancer clusters epidemiologists are typically called on to investigate, a population would have to be exposed to concentrations of carcinogens that are usually credible only in patients undergoing chemotherapy or in some work settings—far greater concentrations than people receive in contaminated neighborhoods and schools. Nevertheless, people resist accepting the explanation that the clusters are random

fluctuations, and so each year state departments of health receive thousands of residential cancer-cluster reports, which result in the publication of hundreds of exhaustive analyses, none of which has convincingly identified an underlying environmental cause. Says Alan Bender, an epidemiologist with the Minnesota Department of Health, those studies "are an absolute, total, and complete waste of taxpayer dollars."[32]

So far in this chapter we have considered some of the ways in which random patterns can fool us. But psychologists have not contented themselves to merely study and categorize such misperceptions. They have also studied the reasons we fall prey to them. Let's now turn our attention to some of those factors.

PEOPLE LIKE TO EXERCISE CONTROL over their environment, which is why many of the same people who drive a car after consuming half a bottle of scotch will freak out if the airplane they are on experiences minor turbulence. Our desire to control events is not without purpose, for a sense of personal control is integral to our self-concept and sense of self-esteem. In fact, one of the most beneficial things we can do for ourselves is to look for ways to exercise control over our lives—or at least to look for ways that help us feel that we do. The psychologist Bruno Bettelheim observed, for instance, that survival in Nazi concentration camps "depended on one's ability to arrange to preserve some areas of independent action, to keep control of some important aspects of one's life despite an environment that seemed overwhelming."[33] Later studies showed that a prior sense of helplessness and lack of control is linked to both stress and the onset of disease. In one study wild rats were suddenly deprived of all control over their environment. They soon stopped struggling to survive and died.[34] In another study, in a group of subjects who were told they were going to take a battery of important tests, even the pointless power to control the order of those tests was found to reduce anxiety levels.[35]

One of the pioneers in the psychology of control is the psycholo-

But is this sample big enough?

gist and amateur painter Ellen Langer, now a professor at Harvard. Years ago, when she was at Yale, Langer and a collaborator studied the effect of the feeling of control on elderly nursing home patients.[36] One group was told they could decide how their rooms would be arranged and were allowed to choose a plant to care for. Another group had their rooms set up for them and a plant chosen and tended to for them. Within weeks the group that exercised control over their environment achieved higher scores on a predesigned measure of well-being. Disturbingly, eighteen months later a follow-up study shocked researchers: the group that was not given control experienced a death rate of 30 percent, whereas the group that was given control experienced a death rate of only 15 percent.[37]

Why is the human need to be in control relevant to a discussion of random patterns? Because if events are random, we are *not* in control, and if we are in control of events, they are *not* random. There is therefore a fundamental clash between our need to feel we are in control and our ability to recognize randomness. That clash is one of the principal reasons we misinterpret random events. In fact, inducing people to mistake luck for skill, or pointless actions for control, is one of the easiest enterprises a research psychologist can engage in. Ask people to control flashing lights by pressing a dummy button, and they will believe they are succeeding even though the lights are flashing at random.[38] Show people a circle of lights that flash at random and tell them that by concentrating they can cause the flashing to move in a clockwise direction, and they will astonish themselves with their ability to make it happen. Or have two groups simultaneously compete in a similar enterprise—one strives for clockwise motion along the circle, and the other attempts to make the lights travel counterclockwise—and the two groups will simultaneously perceive the lights traveling around the circle in the direction of their intention.[39]

Langer showed again and again how the need to feel in control interferes with the accurate perception of random events. In one of her studies, participants were found to be more confident of success when competing against a nervous, awkward rival than when com-

peting against a confident one even though the card game in which they competed, and hence their probability of succeeding, was determined purely by chance.[40] In another study she asked a group of bright and well-educated Yale undergraduates to predict the results of thirty random coin tosses.[41] The experimenters secretly manipulated the outcomes so that each student was correct exactly half the time. They also arranged for some of the students to have early streaks of success. After the coin tosses the researchers quizzed the students in order to learn how they assessed their guessing ability. Many answered as if guessing a coin toss were a skill they could cultivate. One quarter reported that their performance would be hampered by distraction. Forty percent felt that their performance would improve with practice. And when asked directly to rate their ability at predicting the tosses, the students who achieved the early streaks of success rated themselves better at the task than did the others even though the number of successes was the same for all the subjects.

In another clever experiment, Langer set up a lottery in which each volunteer received a sports trading card with a player's picture on it.[42] A card identical to one of the distributed cards was placed in a bag with the understanding that the participant whose card it matched would be declared the winner. The players were divided into two groups. Those in one group had been allowed to choose their card; those in the other had been handed a card at random. Before the drawing each participant was given the opportunity to sell his or her card. Obviously, whether participants chose their cards or were handed them had no effect on their chances of winning. Yet those who had chosen their own cards demanded more than four times as much money for them as those selling the randomly assigned cards.

The subjects in Langer's experiments "knew," at least intellectually, that the enterprises in which they were engaging were random. When questioned, for example, none of the participants in the trading-card lottery said they believed that being allowed to choose their card had influenced their probability of winning. Yet they had *behaved* as if it had. Or as Langer wrote, "While people may pay lip service to the concept of chance, they behave as though chance events are subject to control."[43]

In real life the role of randomness is far less obvious than it was in Langer's experiments, and we are much more invested in the outcomes and our ability to influence them. And so in real life it is even more difficult to resist the illusion of control.

One manifestation of that illusion occurs when an organization experiences a period of improvement or failure and then readily attributes it not to the myriad of circumstances constituting the state of the organization as a whole and to luck but to the person at the top. That's especially obvious in sports, where, as I mentioned in the Prologue, if the players have a bad year or two, it is the coach who gets fired. In major corporations, in which operations are large and complex and to a great extent affected by unpredictable market forces, the causal connection between brilliance at the top and company performance is even less direct and the efficacy of reactionary firings is no greater than it is in sports. Researchers at Columbia University and Harvard, for example, recently studied a large number of corporations whose bylaws made them vulnerable to shareholders' demands that they respond to rough periods by changing management.[44] They found that in the three years after the firing there was no improvement, on average, in operating performance (a measure of earnings). No matter what the differences in ability among the CEOs, they were swamped by the effect of the uncontrollable elements of the system, just as the differences among musicians might become unapparent in a radio broadcast with sufficient noise and static. Yet in determining compensation, corporate boards of directors often behave as if the CEO is the *only* one who matters.

Research has shown that the illusion of control over chance events is enhanced in financial, sports, and especially, business situations when the outcome of a chance task is preceded by a period of strategizing (those endless meetings), when performance of the task requires active involvement (those long hours at the office), or when competition is present (this never happens, right?). The first step in battling the illusion of control is to be aware of it. But even then it is

difficult, for, as we shall see in the following pages, once we think we see a pattern, we do not easily let go of our perception.

Suppose I tell you that I have made up a rule for the construction of a sequence of three numbers and that the sequence 2, 4, 6 satisfies my rule. Can you guess the rule? A single set of three numbers is not a lot to go on, so let's pretend that if you present me with other sequences of three numbers, I will tell you whether or not they satisfy my rule. Please take a moment to think up some three-number sequences to test—the advantage of reading a book over interacting in person is that in the book the author can display infinite patience.

Now that you have pondered your strategy, I can say that if you are like most people, the sequences you present will look something like 4, 6, 8 or 8, 10, 12 or 20, 24, 30. Yes, those sequences obey my rule. So what's the rule? Most people, after presenting a handful of such test cases, will grow confident and conclude that the rule is that the sequence must consist of increasing even numbers. But actually my rule was simply that the series must consist of increasing numbers. The sequence 1, 2, 3, for example, would have fit; there was no need for the numbers to be even. Would the sequences you thought of have revealed this?

When we are in the grasp of an illusion—or, for that matter, whenever we have a new idea—instead of searching for ways to prove our ideas wrong, we usually attempt to prove them correct. Psychologists call this the confirmation bias, and it presents a major impediment to our ability to break free from the misinterpretation of randomness. In the example above, most people immediately recognize that the sequence consists of increasing even numbers. Then, seeking to confirm their guess, they try out many more sequences of that type. But very few find the answer the fast way—through the attempt to falsify their idea by testing a sequence that includes an odd number.[45] As philosopher Francis Bacon put it in 1620, "the human understanding, once it has adopted an opinion, collects any instances that confirm it, and though the contrary instances may be more numerous and more weighty, it either does not notice them or else rejects them, in order that this opinion will remain unshaken."[46]

To make matters worse, not only do we preferentially seek evidence to confirm our preconceived notions, but we also interpret ambiguous evidence in favor of our ideas. This can be a big problem because data are often ambiguous, so by ignoring some patterns and emphasizing others, our clever brains can reinforce their beliefs even in the absence of convincing data. For instance, if we conclude, based on flimsy evidence, that a new neighbor is unfriendly, then any future actions that might be interpreted in that light stand out in our minds, and those that don't are easily forgotten. Or if we believe in a politician, then when she achieves good results, we credit her, and when she fails, we blame circumstances or the other party, either way reinforcing our initial ideas.

In one study that illustrated the effect rather vividly, researchers gathered a group of undergraduates, some of whom supported the death penalty and some of whom were against it.[47] The researchers then provided all the students with the same set of academic studies on the efficacy of capital punishment. Half the studies supported the idea that the death penalty has a deterrent effect; the other half contradicted that idea. The researchers also gave the subjects clues hinting at the weak points in each of the studies. Afterward the undergraduates were asked to rate the quality of the studies individually and whether and how strongly their attitudes about the death penalty were affected by their reading. The participants gave higher ratings to the studies that confirmed their initial point of view even when the studies on both sides had supposedly been carried out by the same method. And in the end, though everyone had read all the same studies, both those who initially supported the death penalty and those who initially opposed it reported that reading the studies had strengthened their beliefs. Rather than convincing anyone, the data polarized the group. Thus even random patterns can be interpreted as compelling evidence if they relate to our preconceived notions.

The confirmation bias has many unfortunate consequences in the real world. When a teacher initially believes that one student is smarter than another, he selectively focuses on evidence that tends to

confirm the hypothesis.[48] When an employer interviews a prospective candidate, the employer typically forms a quick first impression and spends the rest of the interview seeking information that supports it.[49] When counselors in clinical settings are advised ahead of time that an interviewee is combative, they tend to conclude that he is even if the interviewee is no more combative than the average person.[50] And when people interpret the behavior of someone who is a member of a minority, they interpret it in the context of preconceived stereotypes.[51]

The human brain has evolved to be very efficient at pattern recognition, but as the confirmation bias shows, we are focused on finding and confirming patterns rather than minimizing our false conclusions. Yet we needn't be pessimists, for it is possible to overcome our prejudices. It is a start simply to realize that chance events, too, produce patterns. It is another great step if we learn to question our perceptions and our theories. Finally, we should learn to spend as much time looking for evidence that we are wrong as we spend searching for reasons we are correct.

Our journey through randomness is now almost at its end. We began with simple rules and went on to learn how they reflect themselves in complex systems. How great is the role of chance in that most important complex system of all—our personal destiny? That's a difficult question, one that has infused much of what we have considered thus far. And though I can't hope to answer it fully, I do hope to shed light on it. My conclusion is evident from the following chapter's title, which is the same as that of this book: "The Drunkard's Walk."

FIRST VINTAGE BOOKS EDITION, MAY 2009

The Library of Congress has cataloged the Pantheon edition as follows:
Mlodinow, Leonard.
The drunkard's walk : how randomness rules our lives / Leonard Mlodinow.
p. cm.
Includes Index.
1. Random variables. 2. Probabilities. 3. Chance. I. Title.
QA273.M63 2008
519.2—dc22 2007042507

Vintage ISBN: 978-0-307-27517-2

Author photograph © Marcio Fernandes
Book design by Virginia Tan

www.vintagebooks.com

Printed in the United States of America
20 19 18 17 16 15 14 13 12 11

To my three miracles of randomness:
Olivia, Nicolai, and Alexei . . .
and for Sabina Jakubowicz