"I shall never believe that God plays dice with the world," Einstein famously said. Whether or not he was right about quantum mechanics and the cosmos, his statement is certainly not true of the games people play in their daily lives. Life is not chess but backgammon, with a throw of the dice at every turn. As a result, it is hard to make predictions, especially about the future (as Yogi Berra allegedly said). But in a universe with any regularities at all, decisions informed by the past are better than decisions made at random. That has always been true, and we would expect organisms, especially informavores such as humans, to have evolved acute intuitions about probability. The founders of probability theory, like the founders of logic, assumed that they were just formalizing common sense.

But then why do people often seem to be "probability-blind," in the words of Massimo Piattelli-Palmarini? Many mathematicians and scientists have bemoaned the innumeracy of ordinary people when they reason about risk. The psychologists Amos Tversky and Daniel Kahneman have amassed ingenious demonstrations of how people's intuitive grasp of chance appears to flout the elementary canons of probability theory. Here are some famous examples.

• People gamble and buy state lottery tickets, sometimes called "the stupidity tax." But since the house must profit, the players, on average, must lose.

• People fear planes more than cars, especially after news of a gory plane crash, though plane travel is statistically far safer. They fear nuclear power, though more people are crippled and killed by coal. Every year a thousand Americans are accidentally electrocuted, but rock stars don't campaign to reduce the household voltage. People clamor for bans on pesticide residues and food additives, though they pose trivial risks of cancer compared to the thousands of natural carcinogens that plants have evolved to deter the bugs that eat them.

• People feel that if a roulette wheel has stopped at black six times in a row, it's due to stop at red, though of course the wheel has no memory and every spin is independent. A large industry of self-anointed seers hallucinate trends in the random walk of the stock market. Hoop fans believe that basketball players get a "hot hand," making baskets in clusters, though their strings of swishes and bricks are indistinguishable from coin flips.

• This problem was given to sixty students and staff members at Harvard Medical School: "If a test to detect a disease whose prevalence is 1/1000 has a false positive rate of 5%, what is the chance that a person found to have a positive result actually has the disease, assuming you know nothing about the person's symptoms or signs?" The most popular answer was .95. The average answer was .56. The correct answer is .02, and only eighteen percent of the experts guessed it. The answer, according to Bayes' theorem, may be calculated as the prevalence or base rate (1/1000) times the test's sensitivity or hit rate (proportion of sick people who test positive, presumably 1), divided by the overall incidence of positive test results (the percentage of the time the test comes out positive, collapsing over sick and healthy people—that is, the sum of the sick people who test positive, $1/1000 \times 1$, and the healthy people who test positive, $999/1000 \times .05$). One bugaboo in the problem is that many people misinterpret "false positive rate" as the proportion of positive results that come from healthy people, instead of interpreting it as the proportion of healthy people who test positive. But the biggest problem is that people ignore the base rate (1/1000), which ought to have reminded them that the disease is rare and hence improbable for a given patient even if the test comes out positive. (They apparently commit the fallacy that because zebras make hoofbeats, hoofbeats imply zebras.) Surveys have shown that many doctors needlessly terrify their patients who test positive for a rare disease.

• Try this: "Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in antinuclear demonstrations. What is the probability that Linda is a bankteller? What is the probability that Linda is a bankteller and is active in the feminist movement?" People sometimes give a higher estimate to the probability that she is a feminist bankteller than to the probability that she is a bankteller. But it's impossible for "A and B" to be more likely than "A" alone.

When I presented these findings in class, a student cried out, "I'm ashamed for my species!" Many others feel the disgrace, if not about themselves, then about the person in the street. Tversky, Kahneman, Gould, Piattelli-Palmarini, and many social psychologists have concluded that the mind is not designed to grasp the laws of probability, even though the laws rule the universe. The brain can process limited amounts of information, so instead of computing theorems it uses crude rules of thumb. One rule is: the more memorable an event, the more likely it is to happen. (I can remember a recent gory plane crash, therefore planes are unsafe.) Another is: the more an individual resembles a stereotype, the more likely he is to belong to that category. (Linda fits my image of a feminist bankteller better than she fits my image of a bankteller, so she's more likely to be a feminist bankteller.) Popular books with lurid titles have spread the bad news: *Irrationality: The Enemy Within; Inevitable Illusions: How Mistakes of Reason Rule Our Minds; How We Know What Isn't So: The Fallibility of Human Reason in Everyday Life*. The sad history of human folly and prejudice is explained by our ineptness as intuitive statisticians.

Tversky and Kahneman's demonstrations are among the most thought-provoking in psychology, and the research has drawn attention to the depressingly low intellectual quality of our public discourse about societal and personal risk. But in a probabilistic world, could the human mind really be oblivious to probability? The solutions to the problems that people flub can be computed with a few keystrokes on a cheap calculator. Many animals, even bees, compute accurate probabilities as they forage. Could those computations really exceed the information-processing capacity of the trillion-synapse human brain? It is hard to believe, and one does not have to believe it. People's reasoning is not as stupid as it might first appear.

To begin with, many risky choices are just that, choices, and cannot be gainsaid. Take the gamblers, plane phobics, and chemical avoiders. Are they really *irrational*? Some people take pleasure in awaiting the outcomes of events that could radically improve their lives. Some people dislike being strapped in a tube and flooded with reminders of a terrifying way to die. Some people dislike eating foods deliberately laced with poison (just as some people might choose not to eat a hamburger fortified with harmless worm meat). There is nothing irrational in any of these choices, any more than in preferring vanilla over chocolate ice cream.

The psychologist Gerd Gigerenzer, along with Cosmides and Tooby, have noted that even when people's judgments of probability depart from the truth, their reasoning may not be illogical. No mental faculty is omniscient. Color vision is fooled by sodium vapor streetlights, but that does not mean it is badly designed. It is demonstrably well designed, far better than any camera at registering constant colors with changing illumination (see Chapter 4). But it owes its success at this unsolvable problem to tacit assumptions about the world. When the assumptions are violated in an artificial world, color vision fails. The same may be true of our probability-estimators.

Take the notorious "gambler's fallacy": expecting that a run of heads increases the chance of a tail, as if the coin had a memory and a desire to be fair. I remember to my shame an incident during a family vacation when I was a teenager. My father mentioned that we had suffered through several days of rain and were due for good weather, and I corrected him, accusing him of the gambler's fallacy. But long-suffering Dad was right, and his know-it-all son was wrong. Cold fronts aren't raked off the earth at day's end and replaced with new ones the next morning. A cloud cover must have some average size, speed, and direction, and it would not surprise me (now) if a week of clouds really did predict that the trailing edge was near and the sun was about to be unmasked, just as the hundredth railroad car on a passing train portends the caboose with greater likelihood than the third car.

Many events work like that. They have a characteristic life history, a changing probability of occurring over time which statisticians call a hazard function. An astute observer *should* commit the gambler's fallacy and try to predict the next occurrence of an event from its history so far, a kind of statistics called time-series analysis. There is one exception: devices that are *designed* to deliver events independently of their history. What kind of device would do that? We call them gambling machines. Their reason for being is to foil an observer who likes to turn patterns into predictions. If our love of patterns were misbegotten because randomness is everywhere, gambling machines should be easy to build and gamblers easy to fool. In fact, roulette wheels, slot machines, even dice, cards, and coins are precision instruments; they are demanding to manufacture and easy to defeat. Card counters who "commit the gambler's fallacy" in blackjack by remembering the dealt cards and betting they won't turn up again soon are the pests of Las Vegas.

So in any world but a casino, the gambler's fallacy is rarely a fallacy. Indeed, calling our intuitive predictions fallacious because they fail on gambling devices is backwards. A gambling device is, by definition, a machine designed to defeat our intuitive predictions. It is like calling our hands badly designed because they make it hard to get out of handcuffs. The same is true of the hot-hand illusion and other fallacies among sports fans. If basketball shots were easily predictable, we would no longer call basketball a sport. An efficient stock market is another invention designed to defeat human pattern detection. It is set up to let traders quickly capitalize on, hence nullify, deviations from a random walk.

Other so-called fallacies may also be triggered by evolutionary novelties that trick our probability calculators, rather than arising from crippling design defects. "Probability" has many meanings. One is relative frequency in the long run. "The probability that the penny will land heads is .5" would mean that in a hundred coin flips, fifty will be heads. Another meaning is subjective confidence about the outcome of a single event. In this sense, "the probability that the penny will land heads is .5" would mean that on a scale of 0 to 1, your confidence that the next flip will be heads is halfway between certainty that it will happen and certainty that it won't.

Numbers referring to the probability of a single event, which only make sense as estimates of subjective confidence, are commonplace nowadays: there is a thirty percent chance of rain tomorrow; the Canadiens are favored to beat the Mighty Ducks tonight with odds of five to three. But the mind may have evolved to think of probabilities as relative frequencies in the long run, *not* as numbers expressing confidence in a single event. The mathematics of probability was invented only in the seventeenth century, and the use of proportions or percentages to express them arose even later. (Percentages came in after the French Revolution with the rest of the metric system and were initially used for interest and tax rates.) Still more modern is the input to the formulas for probability: data gathered by teams, recorded in writing, checked for errors, accumulated in archives, and tallied and scaled to yield numbers. The closest equivalent for our ancestors would have been hearsay of unknown validity, transmitted with coarse labels like *probably*. Our ancestors' usable probabilities must have come from their own experience, and that means they were frequencies: over the years, five out of the eight people who came down with a purple rash died the following day.

Gigerenzer, Cosmides, Tooby, and the psychologist Klaus Fiedler noticed that the medical decision problem and the Linda problem ask for single-event probabilities: how likely is that *this patient* is sick, how likely is it that *Linda* is a bankteller. A probability instinct that worked in relative frequencies might find the questions beyond its ken. There's only one Linda, and either she is a bankteller or she isn't. "The probability that she is a bankteller" is uncomputable. So they gave people the vexing problems but stated them in terms of frequencies, not single-event probabilities. One out of a thousand Americans has the disease; fifty out of a thousand healthy people test positive; we assembled a thousand Americans; how many who test positive have the disease? A hundred people fit Linda's description; how many are banktellers; how many are feminist banktellers? Now a majority of people—up to ninety-two percent—behave like good statisticians.

This cognitive therapy has enormous implications. Many men who test positive for HIV (the AIDS virus) assume they are doomed. Some have taken extreme measures, including suicide, despite their surely knowing that most men don't have AIDS (especially men who do not fall into a known risk group) and that no test is perfect. But it is hard for doctors and patients to use that knowledge to calibrate the chance of being infected, even when the probabilities are known. For example, in recent years the prevalence of HIV in German men who do not belong to a risk group is 0.01%, the sensitivity (hit rate) of a typical HIV test is 99.99%, and the false positive rate is perhaps 0.01%. The prospects of a patient who has tested positive do not sound very good. But now imagine that a doctor counseled a patient as follows: "Think of 10,000 heterosexual men like you. We expect one to be infected with the virus, and he will almost certainly test positive. Of the 9,999 men who are not infected, one additional man will test positive. Thus we get two who test positive, but only one of them actually has the virus. All we know at this point is that you have tested positive. So the chance that you actually have the virus is about 50–50." Gigerenzer has found that when probabilities are presented in this way (as frequencies), people, including specialists, are vastly more accurate at estimating the probability of a disease following a medical test. The same is true for other judgments under uncertainty, such as guilt in a criminal trial.

Gigerenzer argues that people's intuitive equation of probability with frequency not only makes them calculate like statisticians, it makes them think like statisticians about the concept of probability itself—a surprisingly slippery and paradoxical notion. What does the probability of a single event even *mean*? Bookmakers are willing to make up inscrutable numbers such as that the odds that Michael Jackson and LaToya Jackson are the same person are 500 to 1, or that the odds that circles in cornfields emanate from Phobos (one of the moons of Mars) are 1,000 to 1. I once saw a tabloid headline announcing that the chances that Mikhail Gorbachev is the Antichrist are one in eight trillion. Are these statements true? False? Approximately true? How could we tell? A colleague tells me that there is a ninety-five percent chance he will show up at my talk. He doesn't come. Was he lying?

You may be thinking: granted, a single-event probability is just subjective confidence, but isn't it rational to calibrate confidence by relative frequency? If everyday people don't do it that way, wouldn't they be irrational? Ah, but the relative frequency of what? To count frequencies you have to decide on a class of events to count up, and a single event belongs to an infinite number of classes. Richard von Mises, a pioneer of probability theory, gives an example.

In a sample of American women between the ages of 35 and 50, 4 out of 100 develop breast cancer within a year. Does Mrs. Smith, a 49-year-old American woman, therefore have a 4% chance of getting breast cancer in the next year? There is no answer. Suppose that in a sample of women between the ages of 45 and 90—a class to which Mrs. Smith also belongs—11 out of 100 develop breast cancer in a year. Are Mrs. Smith's chances 4%, or are they 11%? Suppose that her mother had breast cancer, and 22 out of 100 women between 45 and 90 whose mothers had the disease will develop it. Are her chances 4%, 11%, or 22%? She also smokes, lives in California, had two children before the age of 25 and one after 40, is of Greek descent . . . What group should we compare her with to figure out the "true" odds? You might think, the more specific the class, the better—but the more specific the class, the smaller its size and the less reliable the frequency. If there were only two people in the world very much like Mrs. Smith, and one developed breast cancer, would anyone say that Mrs. Smith's chances are 50%? In the limit, the only class that is truly comparable with Mrs. Smith in all her details is the class containing Mrs. Smith herself. But in a class of one, "relative frequency" makes no sense.

These philosophical questions about the meaning of probability are not academic; they affect every decision we make. When a smoker rationalizes that his ninety-year-old parents have been puffing a pack a day for decades, so the nationwide odds don't apply to him, he might very well be right. In the 1996 presidential election, the advanced age of the Republican candidate became an issue. *The New Republic* published the following letter:

To the Editors:

In your editorial "Is Dole Too Old?" (April 1) your actuarial information was misleading. The average 72-year-old white man may suffer a 27 percent risk of dying within five years, but more than health and gender must be considered. Those still in the work force, as is Senator Bob Dole, have a much greater longevity. In addition, statistics show that greater wealth correlates to a longer life. Taking these characteristics into consideration, the average 73-year-old (the age that Dole would be if he takes office as president) has a 12.7 percent chance of dying within four years.

Yes, and what about the average seventy-three-year-old wealthy working white male who hails from Kansas, doesn't smoke, and was strong enough to survive an artillery shell? An even more dramatic difference surfaced during the murder trial of O.J. Simpson in 1995. The lawyer Alan Dershowitz, who was consulting for the defense, said on television that among men who batter their wives, only one-tenth of one percent go on to murder them. In a letter to *Nature*, a statistician then pointed out that among men who batter their wives *and whose wives are then murdered by someone*, more than *half* are the murderers.

Many probability theorists conclude that the probability of a single event cannot be computed; the whole business is meaningless. Single-event probabilities are "utter nonsense," said one mathematician. They should be handled "by psychoanalysis, not probability theory," sniffed another. It's not that people can believe anything they want about a single event. The statements that I am more likely to lose a fight against Mike Tyson than to win one, or that I am not likely to be abducted by aliens tonight, are not meaningless. But they are not *mathematical* statements that are precisely true or false, and people who question them have not committed an elementary fallacy. Statements about single events can't be decided by a calculator; they have to be hashed out by weighing the evidence, evaluating the persuasiveness of arguments, recasting the statements to make them easier to evaluate, and all the other fallible processes by which mortal beings make inductive guesses about an unknowable future.

So even the ditziest performance in the *Homo sapiens* hall of shame—saying that Linda is more likely to be a feminist bankteller than a bankteller—is not a fallacy, according to many mathematicians. Since a single-event probability is mathematically meaningless, people are forced to make sense of the question as best they can. Gigerenzer suggests that since frequencies are moot and people don't intuitively give numbers to single events, they may switch to a third, nonmathematical definition of probability, "degree of belief warranted by the information just presented." That definition is found in many dictionaries and is used in courts of law, where it corresponds to concepts such as probable cause, weight of evidence, and reasonable doubt. If questions about single-event probabilities nudge people into that definition—a natural interpretation for subjects to have made if they assumed, quite reasonably, that the experimenter had included the sketch of Linda for some purpose—they would have interpreted the question as, To what extent does the information given about Linda warrant the conclusion that she is a bankteller? And a reasonable answer is, not very much.

A final mind-bending ingredient of the concept of probability is the belief in a stable world. A probabilistic inference is a prediction today based on frequencies gathered yesterday. But that was then; this is now. How do you know that the world hasn't changed in the interim? Philosophers of probability debate whether *any* beliefs in probabilities are truly rational in a changing world. Actuaries and insurance companies worry even more—insurance companies go bankrupt when a current event or a change in lifestyles makes their tables obsolete. Social psychologists point to the schlemiel who avoids buying a car with excellent repair statistics after hearing that a neighbor's model broke down yesterday. Gigerenzer offers the comparison of a person who avoids letting his child play in a river with no previous fatalities after hearing that a neighbor's child was attacked there by a crocodile that morning. The difference between the scenarios (aside from the drastic consequences) is that we judge that the car world is stable, so the old statistics apply, but the river world has changed, so the old statistics are moot. The person in the street who gives a recent anecdote greater weight than a ream of statistics is not necessarily being irrational.

Of course, people sometimes reason fallaciously, especially in today's data deluge. And, of course, everyone should learn probability and statistics. But a species that had no instinct for probability could not learn the subject, let alone invent it. And when people are given information in a format that meshes with the way they naturally think about probability, they can be remarkably accurate. The claim that our species is blind to chance is, as they say, unlikely to be true.