CHANCES ARE...

ADVENTURES IN PROBABILITY

Michael Kaplan and
Ellen Kaplan

1 | Thinking

The present is a fleeting moment, the past is no more; and our prospect of futurity is dark and doubtful. This day may possibly *be my last: but the laws of probability, so true in general, so fallacious in particular, still allow about fifteen years.*

—Gibbon, *Memoirs*

We search for certainty and call what we find destiny. Everything is possible, yet only one thing happens—we live and die between these two poles, under the rule of probability. We prefer, though, to call it Chance: an old familiar embodied in gods and demons, harnessed in charms and rituals. We remind one another of fortune's fickleness, each secretly believing himself exempt. *I* am master of my fate; *you* are dicing with danger; *he* is living in a fool's paradise.

Until the 1660s, when John Graunt, a bankrupt London draper, proposed gauging the vitality of his city by its Bills of Mortality, there were only two ways of understanding the world: inductively, by example; or deductively, by axiom. Truths were derived either from experience—and thus hostages to any counterexample lying in ambush—or were beautiful abstractions: pure, consistent, crystalline, but with no certain relevance to the world of mortals. These two modes of reasoning restricted not just the answers we had about life, but the questions we could ask. Beyond, all else was chance, fortune, fate—the riddle of individual existence.

Graunt was the first to unearth truths from heaps of data. His invention, eventually known as statistics, avoided alike the basic question of Being ("all things are possible") and the uniqueness of individual existence ("only one thing happens"). It got around the problem of uncertainty by asking: "Well, exactly how right do you need to be just now?"

In that same inventive decade, Blaise Pascal was working both on dice-throwing puzzles and on his own, far more compelling, problem: "What shall I do to be saved?" Again, neither induction nor deduction could provide the answers: God and the dice alike refuse to be bound by prior behavior. And yet, and yet . . . in the millennia since Creation, the world has *tended* to be a certain way, just as, over a thousand throws of a die, six *tends* to come up a certain proportion of times. Pascal was the first to see that there could be laws of probability, laws neither fit for Mosaic tablets nor necessarily true for any one time and place, but for life en masse; not for me today, but for mankind through all of time.

The combination of the tool of statistics and the theory of probability is the underpinning of almost all modern sciences, from meteorology to quantum mechanics. It provides justification for almost all purposeful group activity, from politics to economics to medicine to commerce to sports. Once we leave pure mathematics, philosophy, or theology behind, it is the unread footnote to every concrete statement.

And yet it goes against all human instinct. Our natural urge in seeking knowledge is either for deductive logical truth ("Happiness is the highest good") or for inductive truths based on experience ("Never play cards with a man called Doc"). We want questions to fall into one of these categories, which is one of many reasons most of us find probability alien and statisticians unappealing. They don't tell us what we want to know: the absolute truth. Their science is right everywhere but wrong in any one place: like journalism, it is true except in any case of which we have personal knowledge. And, while we may be willing to take a look at the numbers, we rebel at the idea of *being* one—a "mere statistic."

But there are people in the world who can dance with data, people for whom this mass of incomplete uncertainties falls beautifully into place, in patterns as delightful and revealing as a flock of migrating swans. Graunt, Pascal, the Reverend Thomas Bayes, Francis Galton, R. A. Fisher, John von Neumann—all are figures a little to the side of life, perhaps trying to puzzle their way toward a grasp of human affairs that the more sociable could acquire glibly through maxim and proverb. The people who use probability today—market-makers, cardplayers, magicians, artificial-intelligence experts, doctors, war-game designers—have an equally interesting and oblique view of the human affairs they analyze and, to an extent, control.

If you have ever taken a long car journey with an inquisitive child, you will know most of the difficulties with formal reasoning. Questions like "How do you know that?" and "What if it's not like that?" pose real problems—problems that have kept philosophy hard at work for over two thousand years. "How do you know?" is particularly tricky: you "know" that protons are composed of quarks, or that the President spoke in Duluth because he's courting the union vote—but is this "knowing" the same as knowing that the angles of a triangle add up to 180 degrees, or that you have a slight itch behind your left ear? Intuition says they are *not* the same—but how do you know?

This was the great question in Plato's time, particularly because the Sophists insisted that it was no question at all: their idea was that *persuasion* was the basis of knowledge, and that therefore rhetoric was the form of proof. The Sophist Gorgias promised to give his students "such absolute readiness for speaking, that they should be able to convince their audience *independently of any knowledge of the subject.*" Conviction was enough, since, he believed, nothing actually existed; or if it did, it could not be known; or if it could, it was inexpressible. This view offered the advantage that we could know everything the same way—protons to presidents—but had the disadvantage that we knew nothing very well.

Plato and his circle hated the Sophists for their tort-lawyer cockiness and their marketing of wisdom, but most of all for their relativism. Platonists never accept that things are so just because someone has had the last word; some things are so because they *have* to be. A well-constructed pleading does not make 3 equal 5. Plato's student Euclid arranged his books of geometry into *definitions* of objects; *axioms*, the basic statements of their relations; and *theorems*, statements that can be proved by showing how they are only logical extensions of axioms. A demonstration from Euclid has a powerful effect on any inquiring mind; it takes a statement, often difficult to believe, and in a few steps turns this into knowledge as certain as "I am."

So why can't all life be like Euclid? After all, if we could express every field of inquiry as a consistent group of axioms, theorems, and proofs, even a president's speech would have one, incontrovertible meaning. This was Aristotle's great plan. The axioms of existence, he said, were matter and form. All things were the expression of form on matter, so the relationship between any two things could be defined by the forms they shared. Mortality, dogness, being from Corinth, or being the prime mover of the universe—all were aspects of being that could be set in their proper, nested order by logical proof. Thus, in the famous first syllogism of every textbook:

All men are mortal.
Socrates is a man.
Therefore Socrates is mortal.

This *must* be so; the conclusion is built into the definitions. Aristotle's syllogisms defined the science of reasoning from his own time right up to the beginning of the seventeenth century. But there is an essential flaw in deductive reasoning: the difference between the *valid* and the *true*. The rules for constructing a syllogism tell you whether a statement is logically consistent with the premises, but they tell you nothing about the premises themselves. The Kamchatkans believe that volcanoes are actually underground feasting places where demons barbecue whales: if a mountain is smoking, the demons are having a party; there is nothing *logically* wrong with this argument. So deductive logic is confined to describing relations between labels, not necessarily truths about things. It cannot make something from nothing; like a glass of water for Japanese paper flowers, it simply allows existing relationships to unfold and blossom. Today, its most widespread application is in the logic chip of every computer, keeping the details of our lives from crashing into contradiction. But, as computer experts keep telling us, ensuring that the machines are not fed garbage is our responsibility, not theirs. The premises on which automated logic proceeds are themselves the result of human conclusions about the nature of the world—and those conclusions cannot be reached through deduction alone.

You'll remember that the other awkward question from the back seat was "What if it's not like that?" Instinctively, we reason from example to principle, from objects to qualities. We move from seeing experience as a mere bunch of random stuff to positing the subtly ordered web of cause and effect that keeps us fascinated through a lifetime. But are we justified in doing so? What makes our assumptions different from mere prejudice?

Sir Francis Bacon fretted over this question at the turn of the seventeenth century, projecting a new science, cut loose from Aristotle's apron strings and ready to see, hear, feel, and conclude for itself using a method he called "induction." Bacon was Lord Chancellor, the senior judge of the realm, and he proceeded in a lawyerly way, teasing out properties from experience, then listing each property's positive and negative instances, its types and degrees. By cutting through experience in different planes, he hoped to carve away all that was inessential. Science, in his scheme, was like playing "twenty questions" or ordering a meal in a foreign language: the unknown relation was defined by indirection, progressively increasing information by attempting to exclude error.

Induction actually has three faces, each turned a slightly different way. The homely village version is our most natural form of reasoning: the proverb. "Don't insult an alligator until you're over the creek"; "A friend in power is no longer your friend." Everything your daddy told you is a feat of induction, a crystal of permanent wisdom drawn out of the saturated solution of life.

Induction's second, more exalted face is mathematical: a method of amazing power that allows you to fold up the infinite and put it in your pocket. Let's say you want to prove that the total of the first n odd numbers, starting from 1, is n^2. Try it for the first three: $1 + 3 + 5 = 9 = 3^2$; so far, so good. But you don't want to keep checking individual examples; you want to know if this statement is true or false over *all* examples—the first billion odd numbers, the first googol odd numbers.

Why not start by proving the case for the first odd number, 1? Easy:

$1 = 1^2$. Now *assume* that the statement is true for an abstract number n; that is: $1 + 3 + 5 + \ldots$ up to n odd numbers will equal n^2. It would probably help if we defined what the nth odd number is: well, the nth even number would be $2n$ (since the evens are the same as the 2 times table), so the nth odd number is $2n - 1$ (since the first odd, 1, comes before the first even, 2). Now we need to show that *if* the statement is true for the nth odd number, it will also be true for the $n + 1$st; that is:

Assuming that

$$1 + 3 + 5 + \ldots + 2n - 1 = n^2$$

show that

$$(1 + 3 + 5 + \ldots + 2n - 1) + (2n + 1) = (n + 1)^2$$

Let's look more closely at that $(n + 1)^2$ on the right. If we do the multiplication, it comes out as $n^2 + 2n + 1$. But wait a minute: that's the same as n^2, the sum of the first n odd numbers, plus $2n + 1$, the next odd number. So if our statement is true for n odd numbers it *must* be true for $n + 1$.

But, you may be asking, aren't you just proving a relation between two imaginary things? How is this different from deduction? It's different because we already know the statement is true for the first odd number, 1. Set n equal to 1; now we know it's true for the second odd, 3; so we can set n equal to 2, proving the statement for the next odd, 5—and so on. We don't need to look at every example, because all the examples are equivalent; we have constructed a rule that governs them all under the abstract title n. Away they go, like a row of dominoes, rattling off to infinity.

The third, inscrutable, face of induction is scientific. Unfortunately, very little in the observable world is as easily defined as an odd number. Science would be so much simpler if we could consider protons, or prions, or pandas under an abstract p and show that anything true for one is bound to be true for $p + 1$. But of course we can't—and this is where probability becomes a necessity: the things we are talking about, the forms applied to matter, are, like Aristotle's axioms, defined not by themselves but by us. A number or a geometrical form is its own definition—a panda isn't.

Our approach to science follows Bacon's: look and see, question and test. But there are deep questions hiding below these simple instructions. What are you looking for? Where should you look? How will you know you've found it? How will you find it again? Every new observation brings with it a freight of information: some of it contains the vital fact we need for drawing conclusions, but some is plain error. How do we distinguish the two? By getting a sense of likely variation.

This makes scientific induction a journey rather than an arrival; while every example we turn up may confirm the assumption we have made about a cause, we will still never reach ultimate truth. Without repetition we could never isolate qualities from experience, but repetition on its own proves nothing. Simply saying "The sun is bright" requires, in all honesty, the New Englander's reply "Yep—so far."

All swans are white—until you reach Australia and discover the black swans paddling serenely. For science built on induction, the counterexample is always the ruffian waiting to mug innocent hypotheses as they pass by, which is why the scientific method now deliberately seeks him out, sending assumptions into the zone of maximum danger. The best experiments deduce an effect from the hypothesis and then isolate it in the very context where it may be disproved. This falsifiability is what makes a hypothesis different from a belief—and science distinct from the other towers of opinion.

For everyone, not just scientists, induction poses a further problem: we need to act on our conclusions. For those of us who must venture out into the world, wagering our goods on uncertain expectations, the counterexample could well be the storm that sinks our ship, the war that wrecks our country. In human affairs, we cannot hope either to predict with certainty or to test with precision, so we instead try to match the complexity of the moment with the complexity of memory, of imagination, and of character. In studying history we are doing induction of a kitchen rather than laboratory style. When Plutarch contrasted the characters of great Greeks and Romans, or Thomas à Kempis urged us to imitate Christ in all things, they were setting out a line of reasoning by which the complexities of life, seen through the equally complex filter of a virtuous character, could resolve into a simpler decision.

But now that our village walls encompass the whole world, we have exemplars ranging from Mahatma Gandhi to General Patton, which shows the weakness of a purely humanist form of induction. We need a method of reasoning that can offer both the accountability of science and the humanities' openness to untidy, fascinating life. If it is to be accountable, it needs a way to make clear, falsifiable statements, or we are back wrangling with the Sophists. If it is to reflect life, it needs to embrace uncertainty, since that, above all else, is our lot.

Woe's me! woe's me! In Folly's mailbox
Still laughs the postcard, Hope:
Your uncle in Australia
Has died and you are Pope,
For many a soul has entertained
A mailman unawares—
And as you cry, Impossible,
A step is on the stairs.

—Randall Jarrell

The science of uncertainty is probability; it deals with what is repeated but inconsistent. Its statements are not the definitive *all* or *no* of deductive logic but the nuanced: *most*, *hardly*, *sometimes*, and *perhaps*. It separates *normal* from *exceptional*, *predictable* from *random* and determines whether an action is "worth it." It is the science of risk, conjecture, and expectation—that is, of getting on with life.

Yes, but why does probability have to be numerical? Both laypeople and mathematicians groan at the mere mention of probability—the mathematicians because the messiness of the subject seems to sully the discipline itself, leaving it provisional and tentative, a matter of recipes rather than discoveries; and laypeople for the excellent reason that it's hard to see the value for real life in an expression like:

$$P(A|B) = P(A) \times \frac{P(B|A)}{\{P(B|A) \times P(A)\} + \{P(B|\bar{A}) \times P(\bar{A})\}}$$

And yet this is an important statement about the way we come to believe things.

Abstraction, modeling—putting interesting things in numerical terms—can seem like freeze-drying, leaving the shape of life without the flavor. But there is no avoiding number. It is needed to set real things in order, to compare across variety of experience, to handle extremes of scale, and to explore regions our intuition cannot easily enter. It is not intrinsically more *true* than other kinds of discourse—a mortality table is no closer to life than is *Death in Venice*. Nor is speaking numerically a cure for speaking nonsense, although it does offer a more convenient way to detect nonsense once it is said. Numbers make statements about likelihood falsifiable, extend our understanding of experience beyond our local habitation to the extremes of time and space, and give us an elastic frame of reference, equally suitable to this room and the universe, this instant and eternity.

Probability, meanwhile, gives us a method of defining a belief as it ought to exist in a reasonable mind: Truth within known limits—and here, too, number offers a transferable standard by which we can judge that truth.

Why do we need such an abstract standard? Because our senses can fail us and our intuition is often untrustworthy. Our perception of normal and abnormal depends crucially on our field of attention: In a recent experiment, subjects who had been asked to count the number of times basketball players on one team passed the ball failed to spot a man in a gorilla suit running around the court. Even when we are trying to concentrate on an important matter of likelihood—in a doctor's office, in a court of law—our instincts can lead us astray, but probability can get us back on track. The economists Tversky and Kahneman devised a scenario closer to real life: a taxi sideswiping a car on a winter night. There are two taxi companies in town: Blue and Green. The latter owns 85 percent of the cabs on the road. A witness says she saw a blue taxi. Independent tests suggest she makes a correct identification 80 percent of the time. So, what color was the taxi? Almost everyone says that it was blue, because people concentrate on the reliability of the witness. But the real issue is how her reliability affects the base fact that a random taxi has an 85 percent chance of being green. When those two probabilities are combined, the chance that the taxi in question was *green* is actually 59 percent—more likely than not. It's a conclusion we could never reach through intuition—it requires calculation.

If we want a numerical model of uncertainty, we need a way of counting the things that *can* happen and comparing that total with what actually *does* happen. "How do I love thee? Let me count the ways." I can love you or not—that's two possibilities—but Elizabeth Barrett could love you for your wit, gravity, prudence, daring, beauty, presence, experience, or innocence. How could all these aspects, existing to a greater or lesser degree in everyone, have combined so perfectly in just one—brilliant Mr. Browning? How big would London have to be before she could be sure to meet him?

This study of mixed characteristics is called *combinatorics*. It originated with a remarkable thirteenth-century Catalan missionary, Ramon Llull, who saw his vocation as converting the Muslims through logic.

He began with nine aspects of God that all three monotheistic religions agree on: Goodness, Greatness, Eternity, Power, Wisdom, Will, Virtue, Truth, and Glory. He then grouped relations (such as Concordance, Difference, and Contrariety) and divine beings and personifications (God, Angels, Hope, Charity). He went on to show that you could assemble statements from elements of these three sets, chosen at random, and always come up with a convincing result consonant with Christian doctrine.

Substituting letters for these elements of theology, Llull wrote them on three concentric, independently movable disks: a sort of doctrinal one-armed bandit. Spinning the disks at random would produce a valid statement. Moreover, the disks made every combination of elements possible, so that no awkward proposition could be suppressed by a sneaky missionary. Ideally, Llull need simply hand over his machine to a skeptical Muslim and let him convert himself.

While God's qualities may be omnipresent, uniform, and sempiternal, the disks that define secular events have intrinsic gaps or ratchets that complicate our calculations. This is the first challenge in making a model for probability: can you devise a machine that encompasses (or, at the very least, names) all that might happen? What combination of elements makes up the event that interests you? Do these elements affect one another or do they occur independently? Finally, do all of them always contribute to the event?

These are the questions we shall be examining in this book, because they crop up whenever we consider things that don't always happen or seek what turns up only every so often. These questions underline the difference between what we think we know and what we come to know—and even then, may not believe. Daniel Ellsberg ran an experiment in which he showed people two urns. One (he told them) contained 50 percent red and 50 percent black balls; the other, an unknown proportion of red to black balls. He offered $100 to any subject who drew a red ball from either urn. Which urn would they choose? Almost all chose the known proportion over the unknown. Then Ellsberg offered another $100 for a *black* ball; the same subjects still chose the known, 50-50 urn—even though their first decision suggested that they thought the "unknown" urn had fewer red balls than black ones.

The question remains "How right do you need to be?"—and there are large areas of life where we may not yet be right enough. A deeper worry, whether probability can really be truth, still looms like an avenging ghost. Einstein famously remarked that he did not believe God would play dice with the universe. The probabilistic reply is that perhaps the universe is playing dice with God.