

SITUATION NORMAL

I recently bought an electronic kitchen scale. It has a glass platform and an Easy to Read Blue Backlit Display. My purchase was not symptomatic of a desire to bake elaborate desserts. Nor was I intending my apartment to become a haunt for local drug gangs. I was just interested in weighing stuff. As soon as the scale was out of its box I went to my local baker, Greggs, and bought a baguette. It weighed 391 grams. The following day I returned to Greggs and bought another baguette. This one was slightly heftier at 398 grams. Greggs is a chain with more than a thousand shops in the U.K. It specializes in cups of tea, sausage rolls and buns plastered in icing sugar, but I had eyes only for the baguettes. On the third day the baguette weighed 399 grams. By now I was bored with eating a whole baguette every day, but I continued with my daily weighing routine. The fourth baguette was a whopping 403 grams. I thought maybe I should hang it on the wall, like a prize fish. Surely, I thought, the weights would not rise forever, and I was correct. The fifth loaf came in at only 384 grams.

In the sixteenth and seventeenth centuries, Western Europe fell in love with collecting data. Measuring tools like the thermometer, the barometer and the perambulator—a wheel for clocking distances along a road—were all invented during this period, and using them was an exciting novelty. The fact that Arabic numerals, which provided effective notation for the results, were finally in common use among the educated classes helped. Collecting numbers became a popular pastime, and it was no passing fad; the craze marked the beginning of modern science. The ability to describe the world in quantitative, rather than qualitative, terms totally changed our relationship with our own surroundings. Numbers gave us a language

for scientific investigation and with that came a new confidence that we could have a deeper understanding of how things really are.

There is still something fun about measuring; indeed I found my daily ritual of buying and weighing bread surprisingly pleasurable. I would return from Greggs with a skip in my step, eager to see how many grams my baguette would be. The frisson of expectation was just like the feeling when you check soccer scores or financial markets.

The motivation behind my daily trip to the bakers was to chart a table of how the weights were distributed, and after ten baguettes, I could see that the lowest weight was 380 grams, the highest was 410 grams, and one of the weights, 403 grams, was repeated. The spread was quite wide, I thought. The baguettes were all from the same shop, all cost the same amount, and yet the heaviest one was almost 8 percent heavier than the lightest one.

Intrigued, I carried on with my experiment. Uneaten bread piled up in my kitchen. It was fascinating to watch how the weights spread themselves along my table. Though I could not predict how much any one baguette would weigh, when all were taken collectively it was clear that a pattern was definitely emerging. After 100 baguettes, I stopped the experiment. By the end every number between 379 grams and 422 grams had been covered at least once with only four exceptions.

Even though I had embarked on the bread project for mathematical reasons, I noticed interesting psychological side effects. Just before weighing each loaf, I would look at it and ponder the color, length, girth and texture—which varied quite considerably. I began to consider myself a connoisseur of baguettes and would say to myself, "Now, this is a heavy one," or "Definitely an average loaf today." I was wrong as often as I was right. Yet my poor forecasting record did not diminish my belief that I was indeed an expert in assessing baguettes. It was, I reasoned, the same self-delusion displayed by sports and financial pundits who are equally unable to predict random events, and yet build careers out of it.

Perhaps the most disconcerting emotional reaction I was having to Greggs's baguettes was what happened when the weights were either extremely heavy or extremely light. On the rare occasions when I weighed a record high or a record low, I was thrilled. The weight was extra-special, and it made the day seem extra-special, as if the exceptionalness of the baguette would somehow be transferred to other aspects of my life. Ratio-

nally, I knew that it was inevitable that some baguettes would be oversized and some undersize, but still, the occurrence of an extreme weight gave me a high. I don't consider myself superstitious, so I was surprised that I was unable to avoid seeing meaning in random patterns. It was a powerful reminder of how susceptible we all are to unfounded beliefs.

• • •

Despite the promise of certainty that numbers provided the scientists of the Enlightenment, they were often not totally certain. After all, sometimes when the same thing was measured twice, it gave two different results. These results were an awkward inconvenience for scientists aiming to find clear and direct explanations for natural phenomena. Galileo Galilei, for instance, noticed that when he was calculating distances of stars with his telescope, his results were prone to variation, and the variation was not due to a mistake in his calculations. Rather, it was because measuring was intrinsically fuzzy. Numbers, it seemed, were not as precise as they'd hoped.


This was exactly what I was experiencing with my baguettes. There were probably many factors that contributed to the daily variance in weight—the amount and consistency of the flour used, the length of time in the oven, the journey of the baguettes from Greggs's central bakery to my local store, the humidity of the air and so on. Likewise, there were many variables affecting the results from Galileo's telescope—such as atmospheric conditions, the temperature of the equipment and personal details, like how tired Galileo was when he recorded the readings.

Still, Galileo was able to see that the variation in his results obeyed certain rules. Despite variation, data for each measurement tended to cluster around a central value, and small errors relative to this central value were more common than large errors. He also noticed that the spread was symmetrical—a measurement was as likely to be less than the central value as it was to be more than the central value.

Likewise, my baguette data showed weights that were clustered around a value of about 400 grams, give or take 20 grams on either side. Even though none of my hundred baguettes weighed precisely 400 grams, there were a lot more baguettes weighing around 400 grams than there were ones weighing around 380 grams or 420 grams. The spread seemed pretty symmetrical, as well.

• • •

The first person to recognize the pattern produced by this kind of measurement error was the German mathematician Carl Friedrich Gauss. The pattern is described by the following curve, called the bell curve:



Gauss's graph needs some explaining. The horizontal axis describes a set of outcomes, for instance the weight of baguettes or the distance of stars. The vertical axis is the probability of those outcomes. A curve plotted on a graph with these parameters is known as a *distribution*. It shows us the spread of outcomes and how likely each is.

There are lots of different types of distributions, but the most basic type is described by the curve above. The bell curve is also known as the *normal distribution*, or the *Gaussian distribution*. Originally it was known as the *curve of error*, but because of its distinctive shape, the term *bell curve* has become more common. The bell curve has an average value, which I have marked X, called the *mean*. The mean is the most likely outcome. The further you go from the mean, the less likely the outcome will be.

When you take two measurements of the same thing and the process has been subject to random error, you tend not to get the same result. However, the more measurements you take, the more the distribution of outcomes begins to look like the bell curve; that is, the outcomes cluster symmetrically around a mean value. Of course, a graph of measurements won't give you a continuous curve—it will give you (as with my baguettes)

a jagged landscape of fixed amounts. The bell curve is a theoretical ideal of the pattern produced by random error. The more data we have, the closer the jagged landscape of outcomes will fit the curve.

In the late nineteenth century, the French mathematician Henri Poincaré knew that the distribution of an outcome that is subject to random measurement error will approximate the bell curve. Poincaré, in fact, conducted the same experiment with baguettes as I did, but for a different reason. He suspected that his local boulangerie was ripping him off by selling underweight loaves, so he decided to use mathematics in the interest of justice. Every day for a year he weighed his daily 1-kilogram loaf. Poincaré knew that if the weight was less than 1 kilogram a few times, this was not evidence of malpractice, because one would expect the weight to vary above and below the specified 1 kilogram. And he conjectured that the graph of bread weights would resemble a normal distribution—since the errors in making the bread, such as how much flour is used and how long the loaf is baked, are random.

After a year he looked at all of the data he had collected. Sure enough, the distribution of weights approximated the bell curve. The peak of the curve, however, was at 950 grams. In other words, the average weight was 0.950 kilogram, not 1 kilogram, as advertised. Poincaré's suspicions were confirmed. The eminent scientist was being cheated by an average of 50 grams per loaf. According to popular legend, Poincaré alerted the Parisian authorities and the baker was given a stern warning.

After his small victory for consumer rights, Poincaré did not let it lie. He continued to measure his daily loaf, and after the second year saw that the shape of the graph was not a proper bell curve; rather, it was skewed to the right. Since he knew that total randomness of error produces the bell curve, he deduced that some nonrandom event was affecting the loaves he was being sold. Poincaré concluded that the baker hadn't stopped his cheapskate, short-measure ways, but instead was giving Poincaré, the squeaky wheel, the largest loaf at hand, thus introducing bias in the distribution. Unfortunately for the *boulangier*, his customer was the cleverest man in France. Again, Poincaré informed the police.

Poincaré's method of baker baiting was prescient; it is now the theoretical basis of consumer protection. When shops sell products at specified weights, the product does not legally have to be that exact weight—it cannot be, since the process of manufacture will inevitably make some items a little heavier and some a little lighter. The job of trading-standards officers

is to take random samples of products on sale and draw up graphs of their weights. For any product they measure, the distribution of weights must fall within a bell curve centered on the advertised mean.

. . .

Half a century before Poincaré saw the bell curve in bread, another mathematician was seeing it wherever he looked. The Belgian Adolphe Quételet was a geometer and astronomer by training, but he soon became sidetracked by a fascination with data—more specifically, with finding patterns in figures. In one of his early projects, Quételet examined French national crime statistics, which the government started publishing in 1825. Quételet noticed that the number of murders was pretty constant every year. Even the proportion of different types of murder weapons—whether it was perpetrated by a gun, a sword, a knife, a fist, and so on—stayed roughly the same. Today this observation is unremarkable—indeed, the way we run our public institutions relies on an appreciation of, for example, crime rates, exam pass rates and accident rates, which we expect to be comparable every year—but Quételet was the first person to notice the amazing regularity of social phenomena when populations are considered as a whole. In any one year, it was possible to predict fairly accurately how many murders would occur. Quételet was troubled by the deep questions about personal responsibility this pattern raised and, by extension, about the ethics of punishment. If society was like a machine that produced a regular number of murderers, didn't this indicate that murder was the fault of society and not the individual?

Quételet's ideas transformed the use of the word *statistics*, whose original meaning had little to do with numbers. The word was used to describe general facts about the state, as in the type of information required by statesmen. Quételet turned statistics into a much wider discipline, one that was less about statecraft and more about the mathematics of collective behavior. He could not have done this without advances in probability theory, which provided techniques to analyze the randomness in data. In Brussels in 1853 Quételet hosted the first international conference on statistics.

Quételet's insights on collective behavior reverberated in other sciences. If by looking at data from human populations you could detect reliable patterns, then it was only a small leap to realize that populations of, for example, atoms also behaved with predictable regularities. James

Clerk Maxwell and Ludwig Boltzmann were indebted to Quételet's statistical thinking when they came up with the kinetic theory of gases, which explains that the pressure of a gas is determined by the collisions of its molecules traveling randomly at different velocities. Though the velocity of any individual molecule cannot be known, the molecules overall behave in a predictable way. The origin of the kinetic theory of gases is an interesting exception to the general rule that developments in the social sciences are the result of advances in the natural sciences. In this case, knowledge flowed in the other direction.

The most common pattern that Quételet found in all of his research was the bell curve. It was ubiquitous in data about human populations. Sets of data in those days were harder to come by than they are now, so Quételet scoured the world for them with the doggedness of a professional collector. For example, he came across a study in the 1814 *Edinburgh Medical Journal* containing chest measurements of 5,738 Scottish soldiers. Quételet drew up a graph of the numbers and showed that the distribution of chest sizes traced a bell curve with a mean of about 40 inches. From other sets of data he showed that the heights of men and women also plot a bell curve. To this day the retail industry relies on Quételet's discoveries. The reason why clothing shops stock more mediums than smalls and larges is that the distribution of human sizes corresponds roughly to the bell curve.

. . .

Quételet died in 1874. A decade later, across the English Channel, a 60-year-old man with a bald head and fine Victorian whiskers could frequently be seen on the streets of Britain gawking at women and rummaging around in his pocket. Francis Galton was an eminent scientist who had devised a way to measure female attractiveness. In order to discreetly register his opinion on passing women, he would prick a needle in his pocket into a cross-shaped piece of paper, to indicate whether she was "attractive," "indifferent" or "repellent." After completing his survey, he compiled a map of the country based on looks. The highest-rated city was London and the lowest-rated was Aberdeen.


Galton was probably the only man in nineteenth-century Europe who was even more obsessed with gathering data than Quételet was. As a young scientist Galton took the temperature of his daily pot of tea, together with information such as the volume of boiling water used and

how delicious it tasted. His aim was to establish how to make the perfect cuppa. (He reached no conclusions.) Galton also built an "anthropometric laboratory"—a sort of walk-in clinic in London, where members of the public could come to have their height, weight, strength of grip, swiftness of blow, eyesight and other physical attributes measured. Galton's lab compiled details on more than 10,000 people, and he achieved such fame that Prime Minister William Gladstone even dropped in to have his head measured.

Galton's research corroborated Quételet's, in that it showed that the variation in human populations was rigidly determined. Galton too saw the bell curve everywhere. In fact, the frequency of the appearance of the bell curve led him to pioneer the term "normal" as the appropriate name for the distribution. The circumference of a human head, the size of the brain and the number of brain fibers all produced bell curves, though Galton was especially interested in nonphysical attributes like intelligence. IQ tests hadn't been invented at that time, so Galton looked for other measures of intelligence. He found them in the results of the admission exams at the Royal Military Academy at Sandhurst. The exam scores, he discovered, also conformed to the bell curve. It filled him with a sense of awe. "I know of scarcely anything so apt to impress the imagination as the wonderful form of cosmic order expressed by the [bell curve]," he wrote. "The law would have been personified by the Greeks and deified, if they had known of it. It reigns with serenity and in complete self-effacement amidst the wildest confusion. The huger the mob, and the greater the apparent anarchy, the more perfect is its sway. It is the supreme law of unreason."

Galton invented a beautifully simple machine that explains the mathematics behind his cherished curve and called it the quincunx. The word's original meaning is the \therefore pattern of five dots on a die, and the contraption is a type of pinball machine in which each horizontal line of pins is offset by half a position from the line above. A ball is dropped into the quincunx at the same point, and the ball then bounces between the pins until it falls out the bottom into a rack of columns. After many balls have been dropped in, the shape they make along the bottom resembles a bell curve.

We can understand what is going on using probability. First, imagine a quincunx with just one pin and say that when a ball hits the pin, the outcome is random, with a 50 percent chance that it bounces to the left and a 50 percent chance that it bounces to the right. In other words, it has



a probability of $\frac{1}{2}$ of ending up one place to the left and a probability of $\frac{1}{2}$ of being one place to the right.

Now, let's add a second row of pins. The ball will either fall left and then left, which I will call LL, or LR or RL or RR. Since moving left and then right is equivalent to staying in the same position, the L and R together cancel each other out (as do the R and L), so there is now a $\frac{1}{4}$ chance the ball will end up one place to the left, a $\frac{3}{4}$ chance it will be in the middle and $\frac{1}{4}$ it will be to the right.

Repeating this for the third row, the equally probable options of where the ball will fall are LLL, LLR, LRL, LRR, RRR, RRL, RLR, RLL. This gives us probabilities of $\frac{1}{8}$ of landing on the far left, $\frac{3}{8}$ of landing on the near left, $\frac{3}{8}$ of landing on the near right and $\frac{1}{8}$ of landing on the far right.

In other words, if there are two rows in the quincunx and we introduce lots of balls in the machine, the law of large numbers says that the balls will fall along the bottom so as to approximate the ratio 1:2:1.



If there are three rows, they will fall in the ratio 1:3:3:1.

If there are four rows, they will fall in the ratio 1:4:6:4:1.

If I carry on working out probabilities, a ten-row quincunx will produce balls falling in the ratio 1:10:45:120:210:252:210:120:45:10:1.

Plotting these numbers gives us the first of the shapes on page 255. The shape becomes even more familiar the more rows we include. Also shown are the results for 100 and 1,000 rows as bar charts. (Note that only the middle sections of these two charts are shown since the values to the left and right are too small to see.)

So, how does this pinball game relate to what goes on in the real



world? Imagine that each row of pins in the quincunx is a random variable that will create an error in measurement. Either it will add a small amount to the correct measurement or it will subtract a small amount. In the case of Galileo and his telescope (page 248), one row of pins could represent whether there is a thermal front passing through, and another could represent the pollution in the air. Each variable contributes an error either one way or the other, just as in the quincunx the ball will bounce left or right. In any measurement there may be many millions of unobservable random errors—their combined errors, however, will give measurements that are distributed like a bell curve.

. . .

If the characteristics of a population are normally distributed, in other words are clustered around an average in the shape of a bell curve, and if

the bell curve is produced through random error, then Quételet argued, the variation in human characteristics can be seen as errors from a paradigm. He called this paradigm *l'homme moyen*, "the average man." Populations, he said, were made up of deviations from this paradigm. In Quételet's mind, being average was something to aspire to, since it was a way of keeping society in check—deviations from the average, he wrote, led to "ugliness in body as well as vice in morals." Even though the concept of *l'homme moyen* never gained acceptance in science, its use filtered down to society at large. We often talk about morality or taste in terms of what an average representative of a population may think or feel about it, such as what is seen as acceptable "in the eyes of the average man."


Whereas Quételet extolled averageness, Galton looked down on it. Galton, as I mentioned before, saw that exam results were normally distributed. Most people scored about average, while a few got very high marks and a few very low. Galton, incidentally, was himself from a very above-average family. His first cousin was Charles Darwin, and the two men corresponded regularly about their scientific ideas. About a decade after Darwin published *On the Origin of Species*, Galton started to theorize on how human evolution itself could be guided. He was interested in the heritability of smarts and wondered how it might be possible to improve the overall intelligence of a population. He wanted to shift the bell curve to the right. To this end, Galton suggested a new field of study about the "cultivation of race," or improving the intellectual stock of a population through breeding. He had thought to call his new science *viticulture*, from the Latin *vita*, life, but eventually settled on *eugenics*, from the Greek *eu*, good, and *genos*, birth. (The usual meaning of "viticulture," grape cultivation, comes from *vitis*, Latin for vine, and dates from around the same time.) Even though many liberal intellectuals of the late nineteenth century and early twentieth century supported eugenics as a way to improve society, the desire to "breed" smarter humans was an idea that was soon distorted and discredited. In the 1930s eugenics became synonymous with murderous Nazi policies to create a superior Aryan race.

In retrospect, it is easy to see how ranking traits—such as intelligence or racial purity—can lead to discrimination and bigotry. Since the bell curve appears when human features are measured, the curve has become synonymous with attempts to classify some humans as intrinsically better than others. The highest-profile example of this was the publication in 1994 of *The Bell Curve* by Richard J. Herrnstein and Charles Murray,

one of the most fiercely debated books of recent years. The book, which owes its name to the distribution of IQ scores, argues that IQ differences between racial groups are evidence of biological differences. Galton wrote that the bell curve reigned with "serenity and in complete self-effacement." Its legacy, though, has been anything but serene and self-effacing.

. . .

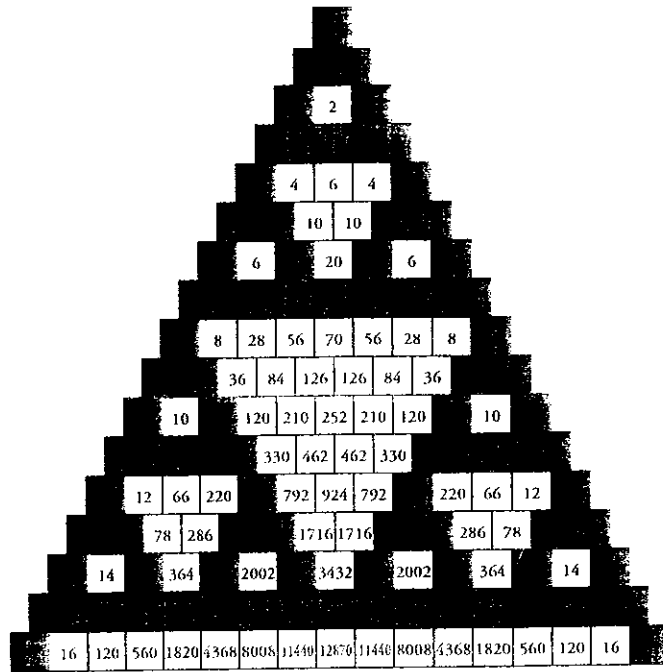
Another way to appreciate the lines of numbers produced by the quincunx is to lay them out like a pyramid. In this form, the results are better known as Pascal's triangle.




Pascal's triangle can be constructed much more simply than by working out the distributions of balls randomly falling through a quincunx. Start with a 1 in the first row, and under it place two 1s so as to make a triangle shape. Continue with subsequent rows, always placing a 1 at the beginning and end of the rows. The value of every other position is the *sum* of the two numbers above it.

The triangle is named after Blaise Pascal, even though he was a latecomer to its charms. Indian, Chinese and Persian mathematicians were all aware of the pattern centuries before he was. Unlike its prior fans, though, Pascal wrote a book about what he called *le triangle arithmétique*. He was fascinated by the mathematical richness of the patterns he discovered. "It is a strange thing how fertile it is in properties," he wrote, adding that in his book he had to leave out more than he could put in.

My favorite feature of Pascal's triangle is the following. Let each number have its own square, and color all the odd-number squares black. Keep all the even number squares white. The result is this wonderful mosaic:



The resulting pattern is reminiscent of the Sierpinski carpet, the piece of mathematical upholstery I discussed in Chapter Two, in which a square is divided into nine sub-squares and the central one is removed, with the same process being repeated with each of the sub-squares, ad infinitum. The triangular version of the Sierpinski carpet is the Sierpinski triangle, in which an equilateral triangle is divided into four identical equilateral triangles, of which the middle one is removed. The three remaining triangles are then subjected to the same operation—divide into four and remove the middle one. Here are the first three iterations:



If we extend the above method of coloring Pascal's triangle to more and more lines, the pattern looks more and more like the Sierpinski triangle. In fact, as the limit approaches infinity, Pascal's triangle *becomes* the Sierpinski triangle.

Sierpinski is not the only familiar friend we find in these black-and-white tiles. Consider the size of the white triangles down the center of the main triangle. The first is made up of one square, the second is made up of 6 squares, the third is made of 28, and the next ones 120 and 496 squares. Do these numbers ring any bells? Three of them—6, 28 and 496—are perfect numbers, which I covered in Chapter Seven. This occurrence is a remarkable visual expression of a seemingly unrelated abstract idea.

• • •

Ancient Indian interest in Pascal's triangle concerned combinations of objects. For instance, imagine we have three fruits: a mango, a lychee and a banana. There is only one combination of three items: mango, lychee, banana. If we want to select only two fruit, we can do this in three different ways: mango and lychee; mango and banana; lychee and banana. There are also only three ways of taking the fruit individually, which is each fruit on its own. The final option is to select zero fruit, and this can happen in only one way. In other words, the number of combinations of three different fruit produces the string 1, 3, 3, 1—the third line of Pascal's triangle.

If we had four objects, the number of combinations when taken none at a time, individually, two at a time, three at a time and four at a time is 1, 4, 6, 4, 1—the fourth line of Pascal's triangle. We can continue this for more and more objects and we see that Pascal's triangle is a reference table for the arrangement of things. If we had n items and wanted to know how many combinations we could make of m of them, the answer is exactly the m th position in the n th row of Pascal's triangle. (Note: by convention, the leftmost 1 of any row is taken as the zeroth position in the row.) For example, how many ways are there of grouping three fruits from a selection of seven fruits? There are 35 ways, since the third position on row seven is 35.

Now, let's move on to start combining mathematical objects. Consider the term $x + y$. What is $(x + y)^2$? It is the same as $(x + y)(x + y)$. To expand this, we need to multiply each term in the first bracket by each term in the second. So, we get $xx + xy + yx + yy$, or $x^2 + 2xy + y^2$. If we carry on, we can see the pattern more clearly. The coefficients of the individual terms are the rows of Pascal's triangle.

$$(x + y)^2 = x^2 + 2xy + y^2$$

$$(x + y)^3 = x^3 + 3x^2y + 3xy^2 + y^3$$

$$(x + y)^4 = x^4 + 4x^3y + 6x^2y^2 + 4xy^3 + y^4$$


The mathematician Abraham de Moivre, a Huguenot refugee living in London in the early eighteenth century, was the first to understand that the coefficients of these equations will approximate a curve the more times you multiply $(x + y)$ together. He didn't call it the bell curve, or the curve of error, or the normal distribution, or the Gaussian distribution, which are the names it later acquired. The curve made its first appearance in math literature in de Moivre's 1718 book on gaming, *The Doctrine of Chances*. This was the first textbook on probability theory, and another example of how scientific knowledge flourished thanks to gambling.

. . .

I've been treating the bell curve as if it is one curve, when, in fact, it is a family of curves. They all look like a bell, but some are wider than others.

Here's an explanation for why we get different widths. If Galileo, for example, measured planetary orbits with a twenty-first-century telescope, the margin of error would be less than if he were using his sixteenth-century telescope. The modern instrument would produce a much thinner bell curve than the antique one. The errors would be much smaller, but they would still be distributed normally.

The average value of a bell curve is called the mean. The width is called the *deviation*. If we know the mean and the deviation, then we know the shape of the curve. It is incredibly convenient that the normal curve can be described using only two parameters. Perhaps, though, it is too con-



Bell curves with different deviations.

venient. Often, statisticians are overly eager to find the bell curve in their data. Bill Robinson, an economist who heads KPMG's forensic accounting division in London, admits this is the case. "We love to work with normal distributions because it has mathematical properties that have been very well explored. Once we know it's a normal distribution, we can start to make all sorts of interesting statements."

Robinson's job, in basic terms, is to deduce, by looking for patterns in huge data sets, whether someone has been cooking the books. He is carrying out the same strategy that Poincaré used in weighing the loaves every day, except that Robinson is looking at gigabytes of financial data and has much more sophisticated statistical tools at his disposal.

Robinson said that his department tends to work on the assumption that for any set of data, the default distribution is the normal distribution. "I think in the financial markets it is true that we have assumed a normal distribution when perhaps it doesn't work." In recent years, in fact, there has been a backlash in both academia and finance against the historic reliance on the normal distribution.

When a distribution is less concentrated around the mean than the bell curve it is called *platykurtic*, from the Greek words *platus*, meaning flat, and *kurtos*, bulging. Conversely, when a distribution is more concentrated around the mean it is called *leptokurtic*, from the Greek *leptos*, meaning thin. William Sealy Gosset, a statistician who worked for the Guinness brewery in Dublin, drew the following aide-mémoire in 1908 to remember which was which: a duck-billed platypus was platykurtic and the kissing kangaroos were leptokurtic. He chose kangaroos because they are "noted for 'lepping,' though, perhaps, with equal reason they should be hares!" Gosset's sketches are the origin of the term *tail* for describing the far-left and far-right sections of a distribution curve.

When economists talk of distributions that are *fat-tailed* or *heavy-tailed*, they are talking of curves that stay higher than normal from the



Platykurtic and leptokurtic distributions.

axis at the extremes, as if Gosset's animals have larger than average tails. These curves describe distributions in which extreme events are more likely than if the distribution were normal. For instance, if the variation in the price of a share was fat-tailed, it would mean there was more chance of a dramatic drop, or hike, in price than if the variation was normally distributed. For this reason, it can sometimes be reckless to assume a bell curve over a fat-tailed curve.


The economist Nassim Nicholas Taleb's position in his best-selling book *The Black Swan* is that we have tended to underestimate the size and importance of the tails in distribution curves. He argues that the bell curve is a historically defective model because it cannot anticipate the occurrence of, or predict the impact of, very rare, extreme events—such as a major scientific discovery like the invention of the Internet, or a terrorist attack like 9/11. “The ubiquity of the [normal distribution] is not a property of the world,” he writes, “but a problem in our minds, stemming from the way we look at it.”

The desire to see the bell curve in data is perhaps most strongly felt in education. The awarding of grades from A to F in end-of-year exams is based on where a pupil's score falls on a bell curve, which the distribution of grades is expected to approximate. The curve is divided into sections, with A representing the top section, B the next section down, and so on. For the education system to run smoothly, it is important that the percentages of pupils getting grades A to F from year to year are comparable. If there are too many As or too many Fs in one particular year, the consequences—not enough or too many people in certain courses—would be a strain on resources. Exams are specifically designed in the hope that the distribution of results replicates the bell curve as much as possible—irrespective of whether or not this is an accurate reflection of intelligence.

It has even been argued that the reverence some scientists have for the bell curve encourages sloppy practices. We saw from the quincunx that random errors are distributed normally. So, the more random errors we can introduce into measurement, the more likely it is that we will get a bell curve from the data—even if the phenomenon being measured is not normally distributed. When the normal distribution is found in a set of data, this could simply be because the measurements have been gathered with too little care.

. . .

Which brings me back to my baguettes. Were their weights really normally distributed? Was the tail thin or fat? As you'll recall, I weighed 100 baguettes. The results showed some promising trends—there was a mean of somewhere around 400 grams and a more or less symmetrical spread between 380 and 420 grams. If I had been as indefatigable as Henri Poincaré, I would have continued the experiment for a year and had 365 (give or take days of bakery closure) weights to compare. With more data, the distribution would have been clearer. Still, my smaller sample was enough to get an idea of the pattern forming. I used a trick, compressing my results by drawing the graph with a scale that grouped baguette weights in bounds of 8 grams rather than 1 gram. This created the following graph:



Upon drawing this I felt relief, as it really looked as though my baguette experiment was producing a bell curve. My facts appeared to be fitting the theory. But when I looked closer, the graph wasn't really like the bell curve at all. Yes, the weights were clustered around a mean, but the curve was clearly not symmetrical. The left side of the curve was not as

steep as the right side. It was as if there was an invisible magnet stretching the curve a little to the left.

I could therefore conclude one of two things. Either the weights of Greggs's baguettes were not normally distributed, or they were normally distributed but some bias had crept into my experimentation process. I had an idea of what the bias might be. I had been storing the uneaten baguettes in my kitchen, and I decided to weigh one that was a few days old. To my surprise, it was only 321 grams—significantly lower than the lowest weight I had measured. It dawned on me then that baguette weight was not fixed, because bread gets lighter as it dries out. I bought another loaf and discovered that a baguette loses about 15 grams between 8 A.M. and noon.

It was now clear that my experiment was flawed. I had not taken into account the hour of the day when I took my measurements. It was almost certain that this variation was providing a bias to the distribution of weights. Most of the time I was the first person in the shop, and weighed my loaf at about 8:10 A.M.; but sometimes I got up late. This random variable was not normally distributed, since the mean would have been between 8 A.M. and 9 A.M., but there was no tail before 8 A.M. as the shop was closed. The tail on the other side went all the way to lunchtime. Then something else occurred to me. What about the ambient temperature? I had started my experiment at the beginning of spring. It had ended at the beginning of summer, when the weather was significantly hotter. I looked at the figures and saw that my baguette weights were lighter on the whole toward the end of the project. The summer heat, I assumed, was drying them out faster. Again, this variation could have had the effect of stretching the curve leftwards.

My experiment may have shown that baguette weights approximated a slightly distorted bell curve, but what I had really learned was that measurement is never so simple. The normal distribution is a theoretical ideal, and one cannot assume that all results will conform to it. I wondered about Henri Poincaré. When he measured his bread did he eliminate bias due to the Parisian weather, or the time of day of his measurements? Perhaps he had not demonstrated that he was being sold a 950-gram loaf instead of a 1-kilogram loaf at all, but had instead proved that from baking to measuring, a 1-kilogram loaf reduces in weight by 50 grams. The history of the bell curve, in fact, is a wonderful parable about the curious kinship between theoretical and applied science. Poincaré once received a letter

from the French physicist Gabriel Lippmann, who brilliantly summed up why the normal distribution was so widely exalted: "Everybody believes in the [bell curve]: the experimenters, because they think it can be proved by mathematics; and the mathematicians, because they believe it has been established by observation." In science as in so many other spheres, we often choose to see what serves our own interests.