



The Mathematics of Coincidence

Professor Sarah Hart

5th March 2024

In this lecture we are going to explore the world of coincidence. We've all heard of, and possibly experienced, occasional astounding coincidences – travelling halfway round the world only to bump into your next door neighbour, or perhaps dreaming of an event that then happens the next day. Is it more than chance when one person survives multiple lightning strikes, and another wins the lottery twice? Mathematics can give us the tools to decide when events really are coincidences, and when other factors are at play. Getting this right has implications everywhere from the science lab to the law courts.

Introduction: the curious case of Roy Sullivan

A man called Roy Sullivan has a dubious honour – during his lifetime he was struck by lightning on seven separate occasions, and survived each time, though after the third time he started carrying a water can with him at all times in case he was struck again and his hair caught fire. Perhaps understandably, he started to worry there was something more than chance behind it. Was there?

Before we start saying that he was cursed (or perhaps conversely blessed, as he survived every lightning strike) we need to try and establish whether something like this could happen by chance. To do that, we need to ask things like: how frequently does lightning strike the earth? How many people are struck by lightning each year? And what proportion of the population is that?

As a warm-up, here is something that happened to me, that threw me for a moment, where it just so happens that I can be absolutely precise about the probability of it occurring by chance. I was in a coffee shop with my daughter, and as I was paying for our drinks I asked if we could use the loo. I tapped my PIN into the card reader, they handed me a slip of paper which I thought was my receipt, but when I looked, it was blank apart from a single number – my PIN! They must have wondered why I gave them such a startled look, before I realized it was simply the keycode to the bathroom door. The odds here are easy to calculate. For a four-digit number there are ten possible first digits, ten second digits, ten third and ten fourth. So that's $10^4 = 10,000$ possible four-digit numbers. There's a one in ten-thousand chance that my PIN would match the keycode in that café.

Obviously this is quite a coincidence, but let's put it in context. Here we have to do some estimating. The population of York (where the café was) is, according to the 2021 census, 141,671. Some of those won't have bank cards, some will have several. I have three, for instance. Most adults have at least two: a bank account and a credit card. Let's conservatively estimate that the average number is 1. Let's also assume that PINs are randomly distributed. Then, on average, we would expect around 14 of the bank cards owned by residents of York to have the same PIN as that café access code. And if you widen the net to the 8.9 million tourists who visit York each year, around 890 of them would be expected to have that PIN, which is nearly 3 a day. I don't feel so special now.

This illustrates the first key point about coincidences. It may be very unlikely that a particular event will happen to me personally, or any one individual, but highly likely that it will happen to someone. The most obvious example is the lottery. Any one ticket has a minuscule chance of winning the jackpot, but nevertheless the lottery jackpot is won many times each year.

With that in mind let's think about poor Roy Sullivan, who was nicknamed the "human lightning rod". It's been estimated that around the world there are an average of one hundred lightning strikes every second. That's $366 \times 24 \times 3600 \times 100 = 3.16$ billion in 2024 alone (it's a leap year of course, so we get a few

more). Very few of these hit people, of course. In America, where Roy Sullivan was from, detailed records are kept, and several hundred people are struck by lightning each year, even though the probability of it happening to any specific individual is extremely low. The average odds of being struck by lightning at least once in your life have been estimated by the US National Weather service as 1 in 15,300. So what are the odds of someone being struck twice? If we toss a coin, the probability of it coming up heads is $\frac{1}{2}$, so the probability of two heads in a row is $\frac{1}{2} \times \frac{1}{2} = \frac{1}{4}$. So we might estimate that the probability of someone getting struck by lightning twice is $\left(\frac{1}{15,300}\right)^2 = \frac{1}{234,090,000}$. The chance of seven strikes seems infinitesimal (if you work it out, based on these estimates, it's about 1 in 2×10^{29} , or a 2 with 29 zeroes after it). However, this misrepresents things. Coins are interchangeable, but people definitely are not. Who you are, and crucially where you are, matters. Lightning is much more likely in warmer regions where there is the most upward air motion and moisture – the more southerly parts of the U.S. East Coast, for example. It's more likely in the summer than winter, and in the afternoon than during the night. We can also speculate about the reasons why 84% of the deaths by lightning strike are men, only 16% women. But the point is that different people have different risk factors. A woman working, say, as a librarian in Alaska is going to be a lot less likely to be struck by lightning than a man working as a park ranger in Virginia. Roy Sullivan, needless to say, was a park ranger in Virginia. How do you avoid being struck by lightning yourself? I read an amusing piece of advice in a book called *Life: The Odds*, based on the stark gender statistics around lightning. If you are a woman, just do what feels the sensible thing to do. If you are man, see if you can spot a woman nearby and do whatever she's doing.

Our lesson from this, anyway, is our next coincidence observation: the odds of particular events may well vary dramatically across a population. If I tell you that my daughter once found herself in a room where she was one of three people who not only had the same birthday but had been born on exactly the same day, in the same hospital, well, that sounds pretty amazing. There are probably only about 20 people on the whole planet who were born on that day in that hospital. But, of course, this group of people was not randomly selected from all the people on earth. The room was the Reception classroom at our local primary school. Most of the children in the class had been born at the local hospital, of course, and all in the same twelve months. So it's not at all unlikely that three children would share a birthday, and of course if they did, it would be the same year and probably the same hospital.

Predicting the future

Sometimes things seem so unlikely that we don't believe they can really be coincidences. If a fortune-teller predicts you'll meet a tall dark stranger, that's one thing, but what about Jeane Dixon, who predicted that the 1960 presidential election would be won by a Democrat who would be assassinated or die in office? I don't need to remind you that John F. Kennedy (a Democrat) was indeed assassinated. We've all heard of cases where people have experienced uncanny premonitory dreams, or seeming instances of telepathy. You're thinking of an old friend and then they call you out of the blue. You dream about someone, and the next day you hear that they have died. Is the universe sending us signals? Are our souls communicating beyond the vale? Surely such things are too improbable to be mere coincidences, right?

Remember, in a large population, even improbable events occur regularly. If on a given day everyone in a population has a 1 in a million chance of some event happening, then it's more likely than not to happen to at least one person as soon as that population exceeds a large enough number (if you want the exact figure, it's 693,147). How do I work that out? Well, suppose the chance of something happening is 1 in 10.

Then the chance for any given person that this event doesn't happen is $1 - \frac{1}{10} = \frac{9}{10}$. The chance that neither of two people experience the event, assuming the probabilities are independent, is $\frac{9}{10} \times \frac{9}{10} = \frac{81}{100}$. That means 19% of the time, at least one of them will experience it. For three people, we get $\left(\frac{9}{10}\right)^3 = \frac{729}{1000}$ or 72.9% of the time, none of them experience it, which means that 27.1% of the time, at least one will.

More generally, suppose an event has a 1 in n chance of happening to a given person in a given time period. Then in a population of k people, the chance it will happen to none of them is $\left(1 - \frac{1}{n}\right)^k$. There's a nice fact about expressions of this kind. It's that, as n tends to infinity, $\left(1 - \frac{1}{n}\right)^n$ gets closer and closer to $\frac{1}{e} \approx 0.367879$ (here, e is the famous mathematical constant $e \approx 2.71828$ which crops up all over mathematics). If you have a million people, this means there's only a 36.78%, or roughly 37%, chance that

none of them will experience a one-in-a-million chance on any day, meaning a 63% chance that somebody will. Of course that's the one we hear about, because nobody phones the radio station to say, guess what, nothing amazing happened to me today.

If you want to know what size population makes it more likely than not that a particular rare event will occur to someone in the population, then you have to solve $(1 - \frac{1}{n})^k = 0.5$. It turns out that the threshold is $k \approx 0.69n$. So, if the population is more than $0.69n$, and each person has a $1/n$ probability on a given day of experiencing some event, then it's more likely than not that on that day the event will happen to at least one person. All this goes to show that in a large population we get rare events often. But even if we just look at one person, you, it's highly likely that you will experience rare events in your life, because the average life is about 30,000 days, and in each one of those days there are many opportunities for coincidences. If you have on average one opportunity per hour for a 1-in-a-million coincidence each day (like the time I was reading a book about Richard III and had just learned about Jane Shore, who was one of Edward IV's mistresses, when I opened my email to find an email from someone I'd never met, called Jane Shore), then that's 720,000 shots at 1-in-a-million coincidence in your life, which is well over the magic $0.69n$, if $n = 1,000,000$. It's much more likely than not that a clairvoyant premonition or other spooky coincidence will happen to you at some point!

What the psychics and other charlatans do, of course, is to vastly increase their apparent success rate, in two main ways. Jeane Dixon, who predicted the Kennedy assassination, also predicted many things which didn't come true. These include changing her mind about the 1960 US Election and predicting that Nixon would win, predicting that Russia would be the first to put a man on the moon, and that the next World War would start in 1958. I guess we'll find out next year if she was right about China invading Russia in 2025. If you make 100 predictions for every one that comes true, it's rather less impressive. We tend only to remember, or to report on, our successes. The second trick is that, as we all know, most predictions are maddeningly vague. If you predict "something lucky will happen to you this month", or "an MP will have to resign in disgrace this year", you'll have a hard time being wrong. But perhaps we shouldn't be too hard on these psychics. As Niels Bohr famously said, prediction is very hard, especially of the future.

Incredible Luck

The kind of reasoning we've been using to consider one in a million coincidences can also explain apparently unbelievable strokes of luck like that of Evelyn Marie Adams, who won the New Jersey Lottery twice in the space of four months in 1985-6. (Why would you carry on buying a ticket if you already won? I don't know, but apparently most winners do.) The statistician David Hand calculated that the odds of two jackpots in this lottery in a four-month time period are about one in a trillion. I'm not aware of anyone winning the jackpot twice in the UK lottery, but a man called Mike McDermott matched five numbers and the bonus ball twice in a year, with the same set of numbers. To explore these impressive coincidences, we can do some rough calculations. Let's just focus on jackpot winners, and conservatively say there are 200 lotteries around the world (each US state has one, and most other countries have at least one). And there are, say, 100 winning tickets each year for each one (two draws a week; some weeks nobody wins, some weeks several people win). That means 20,000 winners each year around the world. Many lotteries are the "choose 6 numbers from 49" variety – the UK was like this for many years. That gives an approximately 1 in 14 million chance of winning the jackpot (some lotteries have better odds than others – in the UK the odds are now about 1 in 45 million but the Polish lotto the odds are only about 1 in 850,000, but it seems reasonable to use the 1 in 14 million as an average). Using the same reasoning as we did for rare events happening, the chance in any single draw of none of the 20,000 winners having a second win is $(1 - \frac{1}{14,000,000})^{20,000}$. Working this out, we find that 99.86% of the time none of them will win, but that means the chance of at least one of them winning is just over 1 in 1000, or 0.0014. And that means, using our $0.69n$ formula with $n = \frac{1}{0.014}$, that after 493 draws, less than five years, it becomes more likely than not that one of those winners will win again. And every year the ranks of the winners swell. The person I feel most sorry for though is Maureen Wilcox, who in the same week bought a ticket for the Massachusetts Lottery and a ticket for the Rhode Island lottery. She chose the winning numbers for both lotteries, but on the wrong tickets!

What can we make of the astonishing week in 2009 when on September 6th the winning numbers in the Bulgarian lottery were 4, 15, 23, 24, 35, and 42, and then on September 10th the exact same numbers came up again. Obviously, there was an outcry – it must be fraud! The machine must be broken! There

was an official enquiry into how this could happen. They should just have asked a statistician. If we are interested in the same numbers coming up in consecutive draws in a lottery somewhere in the world, then we can use the same calculation we just did – 20,000 draws each year, each with a 1 in 14 million chance of having the same numbers as the last draw, means we'd expect this to happen somewhere in the world every five years or so. But if we don't insist on the draws being consecutive, then even in a single lottery draw it's likely to happen within a few dozen years. In the Bulgarian lottery there's a 1 in 14 million chance in each draw that the same numbers as last time will come up (actually, it's 1 in 13,983,816). But if we look at the chance that a draw will match some previous draw, rather than the one immediately before, the odds improve considerably. Over the fifty years that the Bulgarian lottery had been running by September 2009, there had been more than 5000 draws from which any pair could have had the same winning numbers.

That's $\frac{1}{2}(5000 \times 4999) = 12,497,500$ possible pairs. David Hand calculated that it would be less than 43 years before it became more likely than not that the same numbers would have come up at least twice. In fact, it would start to look a bit suspicious if it hadn't happened.

Six degrees of separation

Like urban myths, where you are told "this happened to my friend's friend" and then when you tell it you also claim it happened to your friend's friend, freaky coincidences get a lot more likely if you are willing to string a few links together. You meet someone, and before long you find that their cousin lives in the town where your mother was born, or their husband once met the same celebrity that your friend's brother does the gardening for. The power of networks to offer up vast numbers of potential links is something we humans are not very good at appreciating. The famous theory of six degrees of separation suggests that we don't have to go too far along the chain to be quite likely to have something in common with everyone we meet. To use the UK (population roughly 70 million) as an example, let's suppose that on average we each know 1000 people – this is friends, relatives, work colleagues, and so on. This felt like a lot, but I checked my phone and it has 414 phone numbers stored. I know loads of people whose phone number I don't have – professional contacts, my neighbours, people in my wider social circle, my kids' friends and my friends' kids, my kids' teachers, and so on and so on. So, let's stick with the 1000 people assumption. That means, on average, the probability of you knowing a random person in the UK is $\frac{1000}{70,000,000}$, or 1 in 70,000. So, assuming (a big assumption but we'll come back to it) that our friends are randomly distributed among the UK population, the chance of you not knowing any of my 1000 friends is $\left(1 - \frac{1}{70,000}\right)^{1000} \approx 0.986$. That is, a 98.6% chance we don't have any mutual friends, which means a 1.4% chance we have a friend in common. But if we allow two intermediate links – that is, a friend of yours who knows a friend of mine, then this compounds. The chance one of my friends doesn't know any of your friends is again 98.6%. But the chance none of my friends knows any of yours is 0.986^{1000} , which is almost zero. In fact there's a better than 99.9999% chance that at least one friend of mine knows at least one friend of yours.

This all assumes a very random pattern of who knows whom. But we know human networks aren't really like that – they aren't randomly mixed. We are more likely to meet people from similar socio-economic and educational backgrounds, so you might expect in the real world if two very different people meet, they are much less likely to find common acquaintances. The mathematicians Steven Strogatz and Duncan Watts researched what they called "small world" networks, where links between nodes (here, that's people) are not necessarily random. In a truly random network, with a random probability of any two nodes being linked, the average path length, in other words the number of steps to get from one node to another, is $\frac{\ln(n)}{\ln(k)}$, where n is the number of nodes and k is the number of edges/friends per node. Fascinatingly, Strogatz and Watts found that even in a highly non-random network, with many clusters and other non-random characteristics, the presence of just a few "short-cut" edges – for example a handful of highly connected individuals with friends across many social groups, makes the path length in such graphs approximately the same as in the random graph. This is called the Watts-Strogatz model. For the UK, if we have $k = 1000$ and $n = 70,000,000$, that gives an average path length of 2.6; that is, any two people are likely linked by a chain of either one or two intermediates, just as we found when we assumed randomness. Even if we take the entire world population, of 8 billion, and make the conservative assumption that each person knows just 100 people, we find an average path length of 4.95. So the six degrees of separation idea really does hold water.

Celestial miracles?

Am I ever going to say that anything is more than coincidence? Yes. The whole of science is based, really, on spotting patterns, or apparent coincidences, and trying to explain them. John Snow produced his famous cholera map in 1854, showing that a particular outbreak of cholera in Soho in London was clustered around a water pump in Broad Street (now called Broadwick Street). This was not a coincidence. When the authorities removed the handle of the pump, making it impossible to use, case numbers dropped dramatically. That was strong evidence that cholera was caused by polluted water rather than “bad air”. Edward Jenner investigated the curious observation that an awful lot of the people immune to smallpox seemed to be dairy farmers and milk maids. Again, it wasn’t a coincidence. It turned out that having cowpox protects you, and that’s how we got the first smallpox vaccination. For thousands of years we have observed the night sky. Stars seem to move in fixed paths across the sky. The moon has the same cycle, waxing and waning in the same way each lunar month. If we notice these things, then we can start to predict, and make hypotheses. But we have to be careful. Correlation, as we all know, does not imply causation. A coincidence or pattern on its own is not evidence of some deeper cause or law. Sometimes, as we have seen, coincidences are just that. In the whole wide universe there is an awful lot of scope for them. One instance is this. Solar eclipses are only possible because the Sun is 400 times as large as the moon, but it’s also 400 times as far away from the Earth. That means they look the same size, and when the moon passes in front of the sun, it covers it almost exactly. It’s a lovely coincidence but it isn’t, as I’ve seen claimed, proof of divine intervention.

If the eclipse coincidence doesn’t convince us that “something mysterious is going on” maybe we’ll be swayed by this astonishing equation, which, so it’s claimed, “cannot be coincidence”. The famous golden number $\phi \approx 1.618$ is involved, as is π , along with the length in days of the year and of the lunar month.

$$\phi \times \frac{365.2422 \times 29.5306}{\pi} = 5555 \text{ (to the nearest whole number)}$$

Not convinced? OK, but hold on to your hats. In the Great Pyramid of Giza, if you inscribe an equilateral triangle with its top vertex at the apex of the pyramid, the side length of the triangle is 555.5 feet – this same 5555 number! Apparently, this is very significant. And there’s more! If you inscribe a circle in the pyramid, its diameter is 365.242 feet, “clearly” a representation of our planet. (Well actually, says the website claiming this, it’s 2.2 feet more than that, but the extra amount accounts for the atmosphere of the earth.) If you want to know why these measurements are in feet, which is not a unit the Ancient Egyptians used, or why you would want to multiply the lengths of the year and the lunar month together, or why the length of the year would be a good choice for the diameter in feet of a circle if it were representing the earth, I have to say I cannot enlighten you. Let’s have one final amazing thing which probably proves aliens or something. The icosahedron has 20 equilateral triangle faces each with three 60° angles, for a total of $20 \times 3 \times 60 = 3600$. Its dual solid, the dodecahedron, has 12 regular pentagon faces each with five 108° angles, for a total of $12 \times 5 \times 108 = 6480$. The sum of those is 10,080. If you add up the mean average diameter of the earth and the mean average diameter of the moon you get 10,080 to what the writer says is “over 99.6% accuracy”. Wow! I checked: the total is 10072.3. This is indeed very close to 10,080 – the accuracy is 99.9%. But somehow, I’m still not very impressed. The point is, that there are an awful lot of ways to combine numbers, and if you are allowed every interesting solid, all its side lengths and angle sizes, and all the possible shapes you can fit in a pyramid, and all the measurements you can take of the earth and the moon, and all the units of measurement you might pick to measure them in (including the Babylonian 360 degrees in a circle and the British mile and feet), then of course you will be able to find links. This is all, to use a technical term, poppycock.

Numbers are wonderful, and if you play around with them you can find all sorts of links. For example, is it fate that I’m speaking to you today? The odd numbers 1, 3, 5, 7, 9 have many mystic resonances and recur across many cultures and religions: One, unity, from which all things spring. Three for the three fates, five for the cosmic pentagram, seven for the wonders of the world, and nine for the nine Norse realms. Three sets of the numbers 1, 3, 5, 7, 9 have uncanny links to Gresham College, which as we all know was founded by the will of Sir Thomas Gresham. The Gresham coat of arms has three stars on it. What about the other odd numbers, 1, 5, 7, 9? Amazingly, 1579 is the exact year that Gresham died. Those same sacred odd numbers appear in the founding date of Gresham College: 1597, if you swap the last two digits. If you make a different swap, the very same numbers give the year I was born. In the crest and those special years we’ve seen three sets of the odd numbers 1, 3, 5, 7, 9, apart, that is from two missing threes. Can it be a coincidence that I am the 33rd Gresham Professor of Geometry? Surely not.

Coincidences in court

The misunderstanding of coincidence as it relates to the pyramids is one thing. But in a court of law such errors can have much more serious consequences. One example of this is known as the prosecutor's fallacy. Say, for example, crime scene forensic evidence finds partial DNA. The police run it through their database of known felons and they get a match. That must be the bad guy, right? Maybe at the trial an expert says the odds of someone's DNA getting that level of match with what's at the crime scene are only 1 in 10,000. With those odds it seems highly implausible that he's innocent. But this is misleading. What we should actually be interested in is the probability that he's guilty *given that he has this marker*. If 1 person in every 10,000 has it, then in a city like London with a population of 9 million, that means there are 900 people with that marker. From that pool of 900, the probability that any one of them chosen at random is guilty is a tiny 1 in 900.

Another error of probability was made, with tragic results, in the case of Sally Clark. Her young baby died in his sleep, and this was recorded as a case of Sudden Infant Death Syndrome (SIDS), sometimes called cot death. This is rare, but sadly does sometimes happen. But then a year later her second baby also died in his sleep. She was accused of murdering them. In the trial, a paediatrician called Sir Roy Meadow testified that the chance of a baby dying of SIDS in this family (the mother being a young, healthy nonsmoker) was 1 in 8,543. One could dispute the exact chances, but this wasn't the problem. The problem was that he then said that the chance of two babies dying in this way was therefore $\left(\frac{1}{8,543}\right) \times \left(\frac{1}{8,543}\right) = \frac{1}{72,982,849}$, or 1 in 73 million. He said that something like this would only occur by chance once a century. Sally Clark was convicted. But there's absolutely no reason to think that SIDS deaths within a given family are independent. It's not known exactly why, but sadly, if there has been one SIDS death in a family, then other babies born in the same family are about ten times more likely to die too. The probability asserted so confidently by Meadow was around a factor of ten off. That may seem still very unlikely, and it is, but it doesn't tell us anything unless we compare it with the probability of someone murdering their two babies. Statistics show that second SIDS deaths, which, tragically, happen four or five times a year, are nine times as likely as double homicides of infants. There was an outcry after Sally Clark's conviction. It was eventually overturned, as indeed were those of a number of other women who had been wrongly convicted after Roy Meadow gave misleading statistical evidence in their trials.

This has been a sombre topic to visit, so let's finish on a lighter note.

A royal coincidence

One of the ways we can find surprising links between people is if it turns out they are related. Programmes like BBC's *Who Do You Think You Are*, as well as websites where you can explore your family tree, even doing DNA tests to look for family members you hadn't known about, have made genealogy more popular than ever. Occasionally something truly startling happens, such as in 2016 when actor Danny Dyer was on the programme, and was shocked to discover that he was descended from Edward III, and thus also from William the Conqueror. What an incredible discovery! Well, I have wonderful news. You are descended from royalty as well! Let's look at your ancestry. You have two parents, four grandparents, and so on, doubling every generation. But that can't go on forever. If a generation is 30 years, then 1100AD is about 30 generations. But that would seem to imply that at that time you (and indeed all of us) had $2^{30} = 1,073,741,824$. But in 1100AD, the world population was estimated to be just 320 million, and the population of Europe smaller still. What's happening, of course, is that not all of these ancestors are different people. Your great-great-great-great-(etc)-grandmother might be also a great-great-great-great-(etc)-aunt, or be that greatⁿ-grandmother in multiple different ways. Family trees are not straightforward trees but interconnected networks with lots of links. They have to be – there just aren't enough ancestors to go around. By the same token, those people alive in 1100AD collectively have to be ancestors not just of you and me but of everyone alive now, in many different ways. However, not everyone contributes, because some people don't have children. If you go back a thousand years in the past, and look at everyone alive in Europe back then, then it's been shown that about 20% of those people have no descendants alive today. But a mathematician called Joseph Chang found something incredibly surprising. Every person in the remaining 80% is on the family tree of *every* European alive today. This is counterintuitive, but it's true. The mathematics behind it would take too long to go through in detail, but effectively the fact that we all "should" have well over a billion ancestors at that point but there's only a few hundred million to go around, means that everyone's family tree is forced to overlap. If you have some

European ancestry then everyone in Europe a thousand years ago, if they have any living descendant today, is your direct ancestor. If you have some European ancestry, then yes, you are descended from William the Conqueror, just like Danny Dyer, your distant cousin. You are also descended from Vikings, and from Romans, and every other European civilization. If Julius Caesar has any living descendants, then you are one of them. And if you want to find a common ancestor of everyone in the world, Chang later found that everyone alive has a common ancestor dating back a mere 3,600 years. It's amazing, but true.

Conclusion

Here, to finish with, are my seven laws of coincidence.

1. Even very improbable events are likely to happen if enough people are involved.
2. Not everyone's odds are the same.
3. You're quite likely to have a 1 in a million coincidence at some point in your life.
4. Psychic powers come from many vague predictions and forgetting your failures.
5. The more combinations, the more coincidences, to the point of meaninglessness.
6. Get your probabilities right, especially in court.
7. Don't underestimate the power of a network.

© Professor Sarah Hart 2024

Further Reading

- *Life: the odds (and how to improve them)*, by Gregory Baer, Gotham Books (2003).
- *The Improbability Principle: Why coincidences, miracles and rare events happen all the time*, by David Hand, Corgi Books (2015).
- *Innumeracy: Mathematical Illiteracy and Its Consequences*, by John Allen Paulos, Penguin Books (2000).
- *A brief history of everyone who ever lived: the stories in our genes*, by Adam Rutherford, W&N books (2017).