The Rules of Contagion

Contents

Adam Kucharski is an associate professor at the London School of Hygiene & Tropical Medicine. A mathematician by training, his work on global outbreaks such as the Ebola epidemic and the Zika virus has taken him from villages in the Pacific Islands to hospitals in Latin America. He is a TED fellow and winner of the 2016 Rosalind Franklin Award Lecture and the 2012 Wellcome Trust Science Writing Prize. The author of *The Perfect Bet*, his writing has appeared in the *Observer*, *Financial Times*, *Scientific American* and *New Statesman*.

Also by Adam Kucharski
The Perfect Bet

For Emily

The Rules of Contagion

Why Things Spread — and Why They Stop

Adam Kucharski

P
PROFILE BOOKS

wellcome collection

Introduction

A FEW YEARS AGO, I accidentally caused a small outbreak of misinformation. On my commute to work, a friend who works in tech had sent me a stock photo of a group hunched over a table wearing balaclavas. We had a running joke about how news articles on computer hacking would often include staged pictures of people looking sinister. But this photo, below a headline about illicit online markets, had taken things much further: as well as balaclavas, there was a pile of drugs, and a man who apparently wasn't wearing any trousers. It seemed so surreal, so inexplicable.

I decided to tweet it. 'This stock photo is fascinating in so many ways,' I wrote,[1] pointing out all the quirks in the image. Twitter users seemed to agree, and within minutes dozens of people had shared and liked my post, including several journalists. Then, just as I was starting to wonder how far it might spread, some users pointed out that I'd made a mistake. It wasn't a stock photo at all; it was a still image from a documentary about drug dealing on social media. Which, in retrospect, made a lot more sense (apart from the lack of trousers).

Somewhat embarrassed, I posted a correction, and interest soon faded. But even in that short space of time, almost fifty thousand people had seen my tweet. Given that my job involves analysing disease outbreaks, I was curious about what had just happened. Why did my tweet spread so quickly at first? Did that correction really slow it down? What if people had taken longer to spot the mistake?

Questions like these crop up in a whole range of fields. When we think of contagion, we tend to think about things like infectious diseases or viral online content. But outbreaks can come in many forms. They might involve things that bring harm – like malware, violence or financial crises – or benefits, like innovations and culture. Some will start with tangible infections such as biological

pathogens and computer viruses, others with abstract ideas and beliefs. Outbreaks will sometimes rise quickly; on other occasions they will take a while to grow. Some will create unexpected patterns and, as we wait to see what happens next, these patterns will fuel excitement, curiosity, or even fear. So why do outbreaks take off – and decline – in the way they do?

THREE AND A HALF YEARS into the First World War, a new threat to life appeared. While the German army was launching its Spring Offensive in France, across the Atlantic people had started dying at Camp Funston, a busy military base in Kansas. The cause was a new type of influenza virus, which had potentially jumped from animals into humans at a nearby farm. During 1918 and 1919, the infection would become a global epidemic – otherwise known as a pandemic – and would kill over fifty million people. The final death toll was twice as many as the entire First World War.[2]

Over the following century, there would be four more flu pandemics. This raises the obvious question: what will the next one look like? Unfortunately it's difficult to say, because previous flu pandemics were all slightly different. There were different strains of the virus, and outbreaks hit some places harder than others. In fact, there's a saying in my field: 'if you've seen one pandemic, you've seen … one pandemic.'[3]

We face the same problem whether we're studying the spread of a disease, an online trend, or something else; one outbreak won't necessarily look like another. What we need is a way to separate features that are specific to a particular outbreak from the underlying principles that drive contagion. A way to look beyond simplistic explanations, and uncover what is really behind the outbreak patterns we observe.

That's the aim of this book. By exploring contagion across different areas of life, we'll find out what makes things spread and why outbreaks look like they do. Along the way, we'll see the connections that are emerging between seemingly unrelated problems: from banking crises, gun violence and fake news to

disease evolution, opioid addiction and social inequality. As well as covering the ideas that can help us to tackle outbreaks, we'll look at the unusual situations that are changing how we think about patterns of infections, beliefs, and behaviour.

Let's start with the shape of an outbreak. When disease researchers hear about a new threat, one of the first things we do is draw what we call an outbreak curve – a graph showing how many cases have appeared over time. Although the shape can vary a lot, it will typically include four main stages: the spark, growth, peak, and decline. In some cases, these stages will appear multiple times; when the 'swine flu' pandemic arrived in the UK in April 2009, it grew rapidly during early summer, peaking in July, then grew and peaked again in late October (we'll find out why later in the book).

Influenza pandemic in the UK, 2009
Data from Public Health England[4]

Despite the different stages of an outbreak, the focus will often fall on the spark. People want to know why it took off, how it started, and who was responsible. In hindsight, it's tempting to conjure up explanations and narratives, as if the outbreak was inevitable and could happen the same way again. But if we simply list the characteristics of successful infections or trends, we end up with an incomplete picture of how outbreaks actually work. Most things don't spark: for every influenza virus that jumps from

animals to humans and spreads worldwide as a pandemic, there are millions that fail to infect any people at all. For every tweet that goes viral, there are many more that don't.

Even if an outbreak does spark, it's only the start. Try and picture the shape of a particular outbreak. It might be a disease epidemic, or the spread of a new idea. How quickly does it grow? Why does it grow that quickly? When does it peak? Is there only one peak? How long does the decline phase last?

Rather than just viewing outbreaks in terms of whether they take off or not, we need to think about how to measure them and how to predict them. Take the Ebola epidemic in West Africa back in 2014. After spreading to Sierra Leone and Liberia from Guinea, cases began to rise sharply. Our team's early analysis suggested that the epidemic was doubling every two weeks in the worst affected areas.[5] It meant that if there were currently 100 cases, there could be 200 more in a fortnight and another 400 after a month. Health agencies therefore needed to respond quickly: the longer it took them to tackle the epidemic, the larger their control efforts would need to be. In essence, opening one new treatment centre immediately was equivalent to opening four in a month's time.

Some outbreaks grow on even faster timescales. In May 2017, the WannaCry computer virus hit machines around the world, including crucial NHS systems. In its early stages, the attack was doubling in size almost every hour, eventually affecting more than 200,000 computers in 150 countries.[6] Other types of technology have taken much longer to spread. When VCRs became popular in the early 1980s, the number of owners was doubling only every 480 days or so.[7]

As well as speed, there's also the question of size: contagion that spreads quickly won't necessarily cause a larger overall outbreak. So what causes an outbreak to peak? And what happens after the peak? It's an issue that's relevant to many industries, from finance and politics to technology and health. However, not everyone has the same attitude to outbreaks. My wife works in advertising;

while my research aims to stop disease transmission, she wants ideas and messages to spread. Although these outlooks seem very different, it's increasingly possible to measure and compare contagion across industries, using ideas from one area of life to help us understand another. Over the coming chapters, we will see why financial crises are similar to sexually transmitted infections, why disease researchers found it so easy to predict games like the ice bucket challenge, and how ideas used to eradicate smallpox are helping to stop gun violence. We will also look at the techniques we can use to slow down transmission or – in the case of marketing – keep it going.

Our understanding of contagion has advanced dramatically in recent years, and not just in my field of disease research. With detailed data on social interactions, researchers are discovering how information can evolve to become more persuasive and shareable, why some outbreaks keep peaking – like the 2009 flu pandemic did – and how 'small-world' connections between distant friends can help certain ideas spread widely (and yet hinder others). At the same time, we're learning more about how rumours emerge and spread, why some outbreaks are harder to explain than others, and how online algorithms are influencing our lives and infringing on our privacy.

As a result, ideas from outbreak science are now helping to tackle threats in other fields. Central banks are using these methods to prevent future financial crises, while technology firms are building new defences against harmful software. In the process, researchers are challenging long-held ideas about how outbreaks work. When it comes to contagion, history has shown that ideas about how things spread don't always match reality. Medieval communities, for example, blamed the sporadic nature of outbreaks on astrological influences; influenza means 'influence' in Italian.[8]

Popular explanations for outbreaks continue to be overturned by scientific discoveries. This research is unravelling the mysteries of contagion, showing us how to avoid simplistic anecdotes and

ineffective solutions. But despite this progress, coverage of outbreaks still tends to be vague: we simply hear that something is contagious or that it's gone viral. We rarely learn why it grew so quickly (or slowly), what made it peak, or what we should expect next time. Whether we're interested in spreading ideas and innovations, or stopping viruses and violence, we need to identify what's really driving contagion. And sometimes, that means rethinking everything we thought we knew about an infection.

1

A theory of happenings

WHEN I WAS THREE YEARS OLD, I lost the ability to walk. It happened gradually at first: a struggle to stand up here, a lack of balance there. But things soon deteriorated. Short distances became tricky, while slopes and stairs were near impossible. One Friday afternoon in April 1990, my parents took me and my failing legs to the Royal United Hospital in Bath. By the next morning I was seeing a neurological specialist. The initial suspect was a spinal tumour. Several days of tests followed; there were X-rays, blood samples, nerve stimulation, and a lumbar puncture to extract spinal fluid. As the results came in, the diagnosis shifted towards a rare condition known as Guillain-Barré syndrome (GBS). Named after French neurologists Georges Guillain and Jean Alexandre Barré, GBS is the result of a malfunctioning immune system. Rather than protecting my body, it had started attacking nerves, spreading paralysis.

Sometimes the sum of human wisdom is to be found, as writer Alexandre Dumas put it, within the words 'wait and hope'.[1] And that was to be my treatment, to wait and to hope. My parents were given a multicoloured party horn to check the strength of my breathing (there was no home equipment small enough for a toddler). If the horn failed to unroll when I blew, it meant the paralysis had reached the muscles that pumped air into my lungs.

There is a photo of me sitting on my grandfather's lap around this time. He is in a wheelchair. He'd caught polio in India aged twenty-five, and had been unable to walk since. I'd only ever known him like that, his strong arms wheeling uncooperative legs. In a way, it brought familiarity to this unfamiliar situation. Yet what linked us was also what separated us. We shared a symptom, but the mark of his polio was permanent; GBS, for all its misery, was usually a temporary condition.

So we waited and we hoped. The party horn never failed to unroll, and a lengthy recovery began. My parents told me GBS stood for 'Getting Better Slowly'. It was twelve months before I could walk, and another twelve before I could manage anything resembling a run. My balance would suffer for years to come.

As my symptoms faded, so did my memories. Events became distant, left behind to another life. I can no longer remember my parents giving me chocolate buttons before the needles. Or how I subsequently refused to eat them – even on a normal day – fearing what would come next. The memories of games of tag at primary school have faded too, with me spending all of lunchtime as 'it', my legs still too weak to catch the others. For the twenty-five years that followed my illness, I never really spoke about GBS. I left school, went to university, completed a PhD. GBS seemed too rare, too meaningless to bring up. Guillain-what? Barré who? The story, which I never told anyway, was over for me.

Except it wasn't quite. In 2015, I was in the Fijian capital Suva when I encountered GBS again, this time professionally. I'd been in the city to help investigate a recent dengue fever epidemic.[2] Transmitted by mosquitoes, the dengue virus causes sporadic outbreaks on islands like Fiji. Although symptoms are often mild, dengue can come with a severe fever, potentially leading to hospitalisation. During the first few months of 2014, over 25,000 people showed up at health centres in Fiji with a suspected dengue infection, putting a huge burden on the health system.

If you're imagining an office perched on a sunny beach, you're not picturing Suva. Unlike Fiji's resort-laden Western division, the capital is a port city in the southeast of the main island, Viti Levu. The two main roads of the city loop down into a peninsula, forming the horseshoe shape of a magnet, with the area in the middle attracting plenty of rain. Locals who were familiar with British weather told me that I'd feel right at home.

Another, much older, reminder of home was to follow soon after. During an introductory meeting, a colleague at the World Health Organization (WHO) mentioned that clusters of GBS had been

appearing on Pacific Islands. Unusual clusters. The annual par for the disease was 1 or 2 cases per 100,000 people, but in some places they'd seen double figures.[3]

Nobody ever worked out why I got GBS. Sometimes it follows an infection – GBS has been linked to flu and pneumonia, as well as other diseases[4] – but sometimes there's no clear trigger. In my case, the syndrome was just noise, a random blip in the grand scheme of human health. But in the Pacific during 2014/15, GBS represented a signal, just like birth defects would soon do in Latin America.

Behind these new signals lay the Zika virus, named after the Zika Forest in southern Uganda. A close relative of the dengue virus, Zika was first identified in the forest's mosquitoes in 1947. In the local language, Zika means 'overgrown'[5] and grow it would, from Uganda to Tahiti to Rio de Janeiro and beyond. Those signals in the Pacific and Latin America in 2014 and 2015 would gradually become clearer. Researchers found increasing evidence of a link between Zika infection and neurological conditions: as well as GBS, Zika seemed to lead to pregnancy complications. The main concern was microcephaly, where babies develop a smaller brain than usual, resulting in a smaller skull.[6] This can cause a host of serious health issues, including seizures and intellectual disabilities.

In February 2016, triggered by the possibility that Zika was causing microcephaly,[7] WHO announced that the infection was a Public Health Emergency of International Concern, or PHEIC (pronounced 'fake'). Early studies had suggested that for every 100 Zika infections during pregnancy, there could be between 1 and 20 babies with microcephaly.[8] Although microcephaly would become the primary concern about Zika, it was GBS that first brought the infection into health agencies' focus, as well as into mine. Sitting in my temporary office in Suva in 2015, I realised that this syndrome, which had shaped so much of my childhood, was one I knew almost nothing about. My ignorance was mostly self-inflicted, with some (entirely understandable) assistance from my

parents: it was years before they told me GBS could be fatal.

At the same time, the health world was facing a much deeper ignorance. Zika was generating a huge volume of questions, few of which could yet be answered. 'Rarely have scientists engaged with a new research agenda with such a sense of urgency and from such a small knowledge base,' wrote epidemiologist Laura Rodrigues in early 2016.[9] For me, the first challenge was to understand the dynamics of these Zika outbreaks. How easily did the infection spread? Were the outbreaks similar to dengue ones? How many cases should we expect?

To answer these questions, our research group started to develop mathematical models of the outbreaks. Such approaches are now commonly used in public health, as well as appearing in several other fields of research. But where do these models originally come from? And how do they actually work? It's a story that starts in 1883 with a young army surgeon, a water tank and an angry staff officer.

RONALD ROSS HAD WANTED to be a writer, but his father pushed him into medical school. His studies at St Bartholomew's in London struggled to compete with his poems, plays and music, and when Ross took his two qualifying exams in 1879, he passed only the surgery one. This meant he could not join the colonial Indian Medical Service, his father's preferred career path.[10]

Unable to practice general medicine, Ross spent the next year sailing the Atlantic as a ship's surgeon. Eventually he passed his remaining medical exam and scraped into the Indian Medical Service in 1881. After two years in Madras, Ross moved to Bangalore to take up a post as Garrison Surgeon in September 1883. From his comfortable colonial viewpoint, he claimed it was a 'picture of pleasure', a city of sun, gardens and pillared villas. The only problem, as he saw it, was the mosquitoes. His new bungalow seemed to attract far more than the other army rooms. He suspected it was something to do with the water barrel sitting outside his window, which was surrounded by the insects.

Ross's solution was to tip over the tank, destroying the mosquitoes' breeding ground. It seemed to work: without the stagnant water, the insects left him alone. Spurred on by his successful experiment, he asked his staff officer if they could remove the other water tanks too. And while they were at it, why not also get rid of the vases and tins that lay scattered around the mess? If the mosquitoes had nowhere to breed, they would have little option but to move on. The officer wasn't interested. 'He was very scornful and refused to allow men to deal with them,' Ross later wrote, 'for he said it would be upsetting to the order of nature, and as mosquitoes were created for some purpose it was our duty to bear with them.'

The experiment would turn out to be the first in a lifelong analysis of mosquitoes. The second study would come over a decade later, inspired by a conversation in London. In 1894, Ross had travelled back to England for a one-year sabbatical. The city had changed a lot since his last visit: Tower Bridge had been completed, Prime Minister William Gladstone had just resigned, and the country was about to get its first film parlour.[11] When Ross arrived, though, his mind was focused elsewhere. He wanted to catch up on the latest malaria research. In India, people regularly fell ill with the disease, which could lead to fever, vomiting, and sometimes death.

Malaria is one of the oldest diseases known to humanity. In fact, it may have been with us for our entire history as a species.[12] However, its name comes from Medieval Italy. Those who caught a fever would often blame 'mala aria': bad air.[13] The name stuck, as did the blame. Although the disease was eventually traced to a parasite called *Plasmodium*, when Ross arrived back in England the cause of its spread was still a mystery.

In London, Ross called on biologist Alfredo Kanthack at St Bartholomew's, hoping to learn about developments he may have missed while in India. Kanthack said that if Ross wanted to know more about parasites like malaria, he should go and speak to a doctor called Patrick Manson. For several years, Manson had

researched parasites in southeastern China. While there, he had discovered how people get infected with a particularly nasty family of microscopic worms called *filariae*. These parasites were small enough to get into a person's bloodstream and infect their lymph nodes, causing fluid to accumulate within the body. In severe cases, a person's limbs could swell to many times their natural size, a condition known as elephantiasis. As well as identifying how the *filariae* caused disease, Manson had shown that when mosquitoes fed on infected humans, they could also suck up the worms.[14]

Manson invited Ross into his lab, teaching him how to find parasites like malaria in infected patients. He also pointed Ross to recent academic papers he'd missed while out in India. 'I visited him often and learnt all he had to tell me,' Ross later recalled. One winter afternoon, they were walking down Oxford Street, when Manson made a comment that would transform Ross's career. 'Do you know,' he said, 'I have formed the theory that mosquitoes carry malaria just as they carry *filariae*.'

Other cultures had long speculated about a potential link between mosquitoes and malaria. British geographer Richard Burton noted that in Somalia, it was often said that mosquito bites brought on deadly fevers, though Burton himself dismissed the idea. 'The superstition probably arises from the fact that mosquitoes and fevers become formidable about the same time,' he wrote in 1856.[15] Some people had even developed treatments for malaria, despite not knowing what caused the disease. In the fourth century, Chinese scholar Ge Hong described how the qinghao plant could reduce fevers. Extracts of this plant now form the basis for modern malaria treatments.[16] (Other attempts were less successful: the word 'abracadabra' originated as a Roman spell to ward off the disease.[17])

Ross had heard the speculation linking mosquitoes and malaria, but Manson's argument was the first to really convince him. Just as mosquitoes ingested those tiny worms when they fed on human blood, Manson reckoned that they could also pick up malaria

parasites. These parasites then reproduced within the mosquito before somehow making their way back into humans. Manson suggested that drinking water might be the source of infection. When Ross returned to India, he set out to test the idea, with an experiment that would be unlikely to pass a modern ethics board. [18] He got mosquitoes to feed on an infected patient then lay eggs in a bottle of water; once the eggs had hatched, he paid three people to drink the water. To his disappointment, none of them got malaria. So how did the parasites get into people?

Ross eventually wrote to Manson with a new theory, suggesting that the infection might spread through mosquito bites. The mosquitoes injected some saliva with each bite: maybe this was enough to let the parasites in? Unable to recruit enough human volunteers for another study, Ross experimented with birds. First, he collected some mosquitoes and got them to feed on the blood of an infected bird. Then he let these mosquitoes bite healthy birds, which soon came down with the disease as well. Finally, he dissected the saliva glands of the infected mosquitoes, where he found malaria parasites. Having discovered the true route of transmission, he realised just how absurd their previous theories had been. 'Men and birds don't go about eating dead mosquitoes,' he told Manson.

In 1902, Ross received the second ever Nobel Prize for medicine for his work on malaria. Despite contributing to the discovery, Manson did not share the award. He only found out that Ross had won when he saw it in a newspaper.[19] The once close friendship between mentor and student gradually splintered into a sharp animosity. Though he was a brilliant scientist, Ross could be a divisive colleague. He got into a series of disputes with his rivals, often involving legal action. In 1912, he even threatened to sue Manson for libel.[20] The offence? Manson had written a complimentary reference letter for another researcher, who was taking up a professorship that Ross had recently vacated. Manson did not rise to the argument, choosing to apologise instead. 'It takes two fools to make a quarrel,' as he later put it.[21]

Ross would continue to work on malaria without Manson. In the process, he'd find a new outlet for his single-minded stubbornness, and a new set of opponents. Having discovered how malaria spread, he wanted to demonstrate that it could be stopped.

MALARIA ONCE HAD A MUCH BROADER reach than it does today. For centuries, the disease stretched across Europe and North America, from Oslo to Ontario. Even as temperatures dropped during the so-called Little Ice Age in the seventeenth and eighteenth centuries, the biting cold of winter would still be followed by the biting mosquitoes of summer.[22] Malaria was endemic in many temperate countries, with ongoing transmission and a regular stream of new cases from one year to the next. Eight of Shakespeare's plays include mentions of 'ague', a medieval term for malarial fever. The salt marshes of Essex, northeast of London, had been a notorious source of disease for centuries; when Ronald Ross was a student, he'd treated a woman who picked up malaria there.

Having made the link between insects and infections, Ross argued that removing mosquitoes was the key to controlling malaria. His experiences in India – like the experiment with the water tank in Bangalore – had persuaded him that mosquito numbers could be reduced. But the idea went against popular wisdom. It was impossible to get rid of every last mosquito, went the argument, which meant there would always be some insects left, and hence potential for malaria to spread. Ross acknowledged that some mosquitoes would remain, but he believed that malaria transmission could still be stopped. From Freetown to Calcutta, his suggestions were at best ignored and at worst derided. 'Everywhere, my proposal to reduce mosquitoes in towns was treated only with ridicule,' he later recalled.

In 1901, Ross had led a team to Sierra Leone to try and put his mosquito control ideas into practice. They cleared away cartloads of tins and bottles. They poisoned the standing water mosquitoes loved to breed in. And they filled potholes so 'death-dealing street-

puddles', as Ross called them, couldn't form on the roads. The results were promising: when Ross visited again a year later, there were far fewer mosquitoes. However, he had warned health authorities the effect would only last if the control measures continued. Funding for the clean up had come from a wealthy Glaswegian donor. When the money ran out, enthusiasm waned, and mosquito numbers increased once again.

Ross had more success advising the Suez Canal Company the following year. They'd been seeing around 2,000 malaria cases a year in the Egyptian city of Ismailia. After intensive mosquito reduction efforts, this number fell below a hundred. Mosquito control was also proving effective elsewhere. When the French had attempted to build a canal in Panama during the 1880s, thousands of workers had died from malaria, as well as yellow fever, another mosquito-borne infection. In 1905, with the Americans now leading the Panama project, US Army Colonel William Gorgas oversaw an intensive mosquito control campaign, making it possible to complete the canal.[23] Meanwhile further south, physicians Oswaldo Cruz and Carlos Chagas were spearheading anti-malaria programmes in Brazil, helping to reduce cases among construction workers.[24]

Despite these projects, many remained sceptical about mosquito control. Ross would need a stronger argument to persuade his peers. To make his point, he would eventually turn to mathematics. During those early years in the Indian Medical Service, he'd taught himself the subject to a fairly advanced level. The artist in him admired its elegance. 'A proved proposition was like a perfectly balanced picture,' he later suggested. 'An infinite series died away into the future like the long-drawn variations of a sonata.' Realising how much he liked the subject, he regretted not studying it properly at school. He was now too far into his career to change direction; what use was mathematics to someone working in medicine? 'It was the unfortunate passion of a married man for some beautiful but inaccessible lady,' as he put it.

Ross put the intellectual affair behind him for a while, but

returned to the subject after his mosquito discovery. This time, he found a way to make his mathematical hobby useful to his professional work. There was a vital question he needed to answer: was it really possible to control malaria without removing every mosquito? To find out, he developed a simple conceptual model of malaria transmission. He started by calculating how many new human malaria infections there might be each month, on average, in a given geographic area. This meant breaking down the process of transmission into its basic components. For transmission to occur, he reasoned, there first needs to be at least one human in the area who is infectious with malaria. As an example, he picked a scenario where there was one infectious person in a village of 1,000. For the infection to pass to another human, an *Anopheles* mosquito would have to bite this infectious human. Ross reckoned only 1 in 4 mosquitoes would manage to bite someone. So if there were 48,000 mosquitoes in an area, he'd expect only 12,000 to bite a person. And because only 1 person in 1,000 was initially infectious, on average only 12 of those 12,000 mosquitoes would bite that one infectious person and pick up the parasite.

It takes some time for the malaria parasite to reproduce within a mosquito, so these insects would also have to survive long enough to become infectious. Ross assumed only 1 in every 3 mosquitoes would make it this far, which meant that of the 12 mosquitoes with the parasite, only 4 would eventually become infectious. Finally, these mosquitoes would need to bite another human to pass on the infection. If, again, only 1 in 4 of them successfully fed off a human, this would leave a single infectious mosquito to transmit the virus. Ross's calculation showed that even if there were 48,000 mosquitoes in the area, on average they would generate only one new human infection.

If there were more mosquitoes, or more infected humans, by the above logic we'd expect more new infections per month. However, there is a second process that counteracts this effect: Ross estimated that around 20 per cent of humans infected with malaria would recover each month. For malaria to remain endemic

in the population, these two processes – infection and recovery – would need to balance each other out. If the recoveries outpaced the rate of new infections, the level of disease eventually would decline to zero.

This was his crucial insight. It wasn't necessary to get rid of every last mosquito to control malaria: there was a critical mosquito density, and once the mosquito population fell below this level, the disease would fade away by itself. As Ross put it, 'malaria cannot persist in a community unless the *Anophelines* are so numerous that the number of new infections compensates for the number of recoveries.'
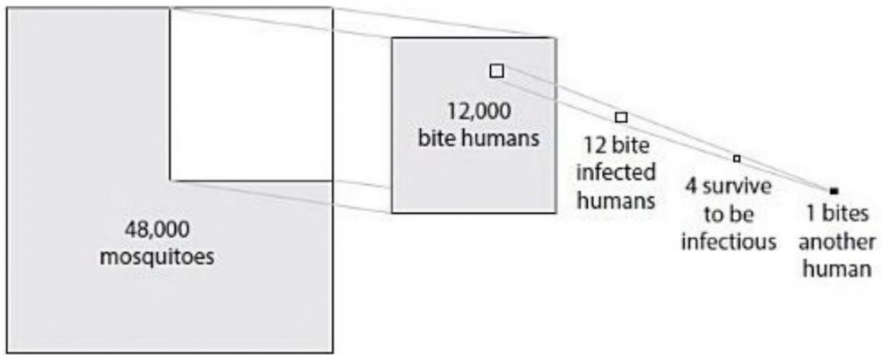
Ross calculated that even if there were 48,000 mosquitoes in a village that contained someone infected with malaria, it might only result in one additional human case

When he wrote up the analysis in his 1910 book *The Prevention of Malaria*, Ross acknowledged that his readers might not follow all of his calculations. Still, he believed that they would be able to appreciate the implications. 'The reader should make a careful study of those ideas,' he wrote, 'and will, I think, have little difficulty in understanding them, though he may have forgotten most of his mathematics'. Keeping with the mathematical theme, he called his discovery the 'mosquito theorem'.

The analysis showed how malaria could be controlled, but it also included a much deeper insight, which would revolutionise how we look at contagion. As Ross saw it, there were two ways to

approach disease analysis. Let's call them 'descriptive' and 'mechanistic' methods. In Ross's era, most studies used descriptive reasoning. This involved starting with real-life data and working backwards to identify predictable patterns. Take William Farr's analysis of a London smallpox outbreak in the late 1830s. A government statistician, Farr had noticed that the epidemic grew rapidly at first, but eventually this growth slowed until the outbreak peaked, then started to decline. This decline was almost a mirror image of the growth phase. Farr plotted a curve through case data to capture the general shape; when another outbreak started in 1840, he found it followed much the same path.[25] In his analysis, Farr didn't account for the mechanics of disease transmission. There were no rates of infection or rates of recovery. This isn't that surprising: at the time nobody knew that smallpox was a virus. Farr's method therefore focused on what shape epidemics take, not why they take that shape.[26]

In contrast, Ross adopted a mechanistic approach. Rather than taking data and finding patterns that could describe the observed trends, he started by outlining the main processes that influenced transmission. Using his knowledge of malaria, he specified how people became infected, how they infected others, and how quickly they recovered. He summarised this conceptual model of transmission using mathematical equations, which he then analysed to make conclusions about likely outbreak patterns.

Because his analysis included specific assumptions about the transmission process, Ross could tweak these assumptions to see what might happen if the situation changed. What effect might mosquito reduction have? How quickly would the disease disappear if transmission declined? Ross's approach meant he could look forward and ask 'what if?', rather than just searching for patterns in existing data. Although other researchers had made rough attempts at this type of analysis before, Ross brought the ideas together into a clear, comprehensive theory.[27] He showed how to examine epidemics in a dynamic way, treating them as a series of interacting processes rather than a set of static patterns.

McKendrick would develop a scathing view of statisticians like Karl Pearson, who relied heavily on descriptive analysis rather than adopting Ross's mechanistic methods. 'The Pearsonians have as usual made a frightful hash of the whole business,' he told Ross after reading a flawed analysis of malaria infections. 'I have no sympathy with them, or their methods.'[32] Traditional descriptive approaches were an important part of medicine – and still are – but they have limitations when it comes to understanding the process of transmission. McKendrick believed the future of outbreak analysis lay with a more dynamic way of thinking. Ross shared this view. 'We shall end by establishing a new science,' he once told McKendrick. 'But first let you and me unlock the door and then anybody can go in who likes.'[33]

ONE SUMMER EVENING IN 1924, William Kermack's experiment exploded, spraying corrosive alkali solution into his eyes. A chemist by training, Kermack had been investigating the methods commonly used to study spinal fluids. He was working alone in Edinburgh's Royal College of Physicians Laboratory that evening, and would eventually spend two months in hospital with his injuries. The accident left the 26-year-old Kermack completely blind.[34]

During his stay in hospital, Kermack asked friends and nurses to read mathematics to him. Knowing that he could no longer see, he wanted to practise getting information another way. He had an exceptional memory and would work through mathematical problems in his head. 'It was incredible to find how much he could do without being able to put anything down on paper,' remarked William McCrea, one of his colleagues.

After leaving hospital, Kermack continued to work in science but shifted his focus to other topics. He left his chemical experiments behind, and began to develop new projects. In particular, he started to work on mathematical questions with Anderson McKendrick, who had risen to become head of the Edinburgh laboratory. Having served in India for almost two decades, McKendrick had left the Indian Medical Service in 1920

and moved to Scotland with his family.

Together, the pair extended Ross's ideas to look at epidemics in general. They focused their attention on one of the most important questions in infectious disease research: what causes epidemics to end? The pair noted that there were two popular explanations at the time. Either transmission ceased because there were no susceptible people left to infect, or because the pathogen itself became less infectious as the epidemic progressed. It would turn out that, in most situations, neither explanation was correct.[35]

Like Ross, Kermack and McKendrick started by developing a mathematical model of disease transmission. For simplicity, they assumed the population mixed randomly in their model. Like marbles being shaken in a jar, everyone in the population has an equal chance of meeting everyone else. In the model, the epidemic sparks with a certain number of infectious people, and everyone else susceptible to infection. Once someone has recovered from infection, they are immune to the disease. We can therefore put the population into one of three groups, based on their disease status:

Given the names of the three groups, this is commonly known as the 'SIR model'. Say a single influenza case arrives in a population of 10,000 people. If we simulate a flu-like epidemic using the SIR model, we get the following pattern:

Simulated influenza outbreak using the SIR model

The simulated epidemic takes a while to grow because only one person is infectious at the start, but it still peaks within fifty days. And by eighty days, it's all but over. Notice that at the end of the epidemic, there are still some susceptible people left. If everyone had been infected, then all 10,000 people would have eventually ended up in the 'recovered' group. Kermack and McKendick's model suggests that this doesn't happen: outbreaks can end before everyone picks up the infection. 'An epidemic, in general, comes to an end before the susceptible population has been exhausted,' as they put it.

Why doesn't everyone get infected? It's because of a transition that happens mid-outbreak. In the early stages of an epidemic, there are lots of susceptible people. As a result, the number of people who become infected each day is larger than the number who recover, and the epidemic grows. Over time, however, the pool of susceptible people shrinks. When this pool gets small enough, the situation flips around: there are more recoveries than new infections each day, so the epidemic begins to decline. There are still susceptible people out there who could be infected, but there are so few left that an infectious person is more likely to recover than meet one.

To illustrate the effect, Kermack and McKendrick showed how the SIR model could reproduce the dynamics of a 1906 outbreak of

plague in Bombay (now Mumbai). In the model, the pathogen remains equally infectious over time; it is the shifting numbers of susceptible and infectious people that lead to the rise and fall.
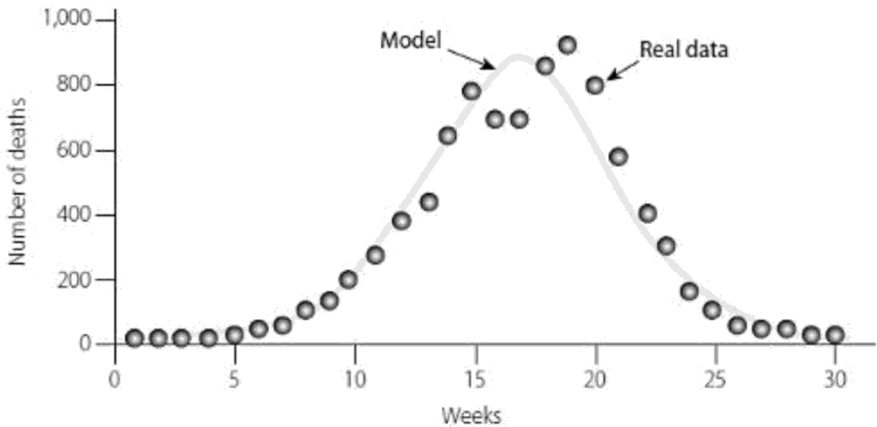
The 1906 plague outbreak in Bombay, with SIR model shown alongside real data

The crucial change happens at the peak of the epidemic. At this point, there are so many immune people – and so few susceptible – that the epidemic cannot continue to grow. The epidemic will therefore turn over and start its decline.

When there are enough immune people to prevent transmission, we say that the population has acquired 'herd immunity'. The phrase was originally coined by statistician Major Greenwood in the early twentieth century (Major was his first name, his army rank was actually captain).[36] Psychologists had previously used 'herd instinct' to describe groups that acted as a collective rather than as individuals.[37] Likewise, herd immunity meant that the population as a whole could block transmission, even if some individuals were still susceptible.

The concept of herd immunity would find popularity several decades later, when people realised it could be a powerful tool for disease control. During an epidemic, people naturally move out of the susceptible group as they become infected. But for many

infections, health agencies can move people out of this group deliberately, by vaccinating them. Just as Ross suggested malaria could be controlled without removing every last mosquito, herd immunity makes it possible to control infections without vaccinating the entire population. There are often people who cannot be vaccinated – such as newborn babies or those with compromised immune systems – but herd immunity allows vaccinated people to protect these vulnerable unvaccinated groups as well as themselves.[38] And if diseases can be controlled through vaccination, they can potentially be eliminated from a population. This is why herd immunity has found its way into the heart of epidemic theory. 'The concept has a special aura,' as epidemiologist Paul Fine once put it.[39]

As well as looking at why epidemics end, Kermack and McKendrick were also interested in the apparently random occurrence of outbreaks. Analysing their model, they found that transmission was highly sensitive to small differences in the characteristics of the pathogen or human population. This explains why large outbreaks can seemingly appear from nowhere. According to the SIR model, outbreaks need three things to take off: a sufficiently infectious pathogen, plenty of interactions between different people, and enough of the population who are susceptible. Near the critical herd immunity threshold, a small change in one of these factors can be the difference between a handful of cases and a major epidemic.

most plausible?

According to evolutionary biologist Nuno Faria and his colleagues, neither theory was particularly good.[45] Based on the genetic diversity of Zika viruses circulating in Latin America by 2016, they reckoned that the infection was introduced much earlier than previously thought. The virus probably hit the continent in mid-to-late 2013. Although too early for the canoe championship or World Cup, the time range coincided with the Confederations Cup, a regional football tournament held in June 2013. What's more, French Polynesia was one of the countries competing.

There was just one gap in the theory: the Confederations Cup occurred five months before the first Zika cases were reported in French Polynesia. But if the outbreak in French Polynesia had in reality started earlier than October 2013 – as our analysis suggested – it was just about plausible that it could have spread to Latin America during that summer. (Of course, we should be cautious about trying too hard to find a sport-shaped prologue for the Zika story: there's always a chance that it was just a random person in the Pacific taking a random flight to Brazil sometime in 2013.)

As well as analysing past outbreaks, we can use mathematical models to look at what might happen in future. This can be particularly useful for health agencies faced with difficult decisions during an outbreak. One such difficulty came in December 2015, when Zika reached the Caribbean island of Martinique. A big concern was the island's ability to handle GBS cases: if patients' lungs failed, they would need to be put on ventilators. At the time, Martinique only had eight ventilators for a population of 380,000. Would it be enough?

To find out, researchers at Institut Pasteur in Paris developed a model of Zika transmission on the island.[46] The crucial thing they wanted to know was the overall shape of the outbreak. GBS cases who required a ventilator typically stayed on it for several weeks, so a short outbreak with a large peak could overwhelm the

health system, while a longer, flatter outbreak would not. At the very start of the Martinique outbreak, there hadn't been many cases, so the team used data from French Polynesia as a starting point. Of the forty-two GBS cases reported there in 2013/14, twelve had required ventilators. According to the Pasteur model, this meant they could have a big problem. If the outbreak in Martinique followed the same pattern as French Polynesia, the island would probably need nine ventilators, one more than they had available.

Fortunately, the Martinique outbreak wouldn't be the same. As new data came in, it became clear that the virus wasn't spreading as quickly as it had in French Polynesia. At the peak of the outbreak, the researchers expected there would be around three GBS cases needing ventilators. Even in the worst-case scenario, they estimated that seven ventilators would be enough. Their conclusion about this upper limit turned out to be correct: at the peak of the outbreak, there were five GBS cases on ventilators. Overall, there were thirty cases of GBS during the outbreak, with two deaths. Without adequate medical facilities, the outcome could have been much worse.[47]

These Zika studies are just a few illustrations of how Ross's methods have influenced our understanding of infectious diseases. From predicting the shape of an outbreak to evaluating control measures, mechanistic models have become a fundamental part of how we study contagion today. Researchers are using models to help health agencies respond to a whole host of outbreaks, from malaria and Zika to HIV and Ebola, in locations ranging from remote islands to conflict zones.

Ross would no doubt be glad to see how influential his ideas have been. Despite winning a Nobel Prize for his discovery that mosquitoes transmit malaria, he did not view this as his biggest achievement. 'In my own opinion my principal work has been to establish the general laws of epidemics,' he once wrote.[48] And he didn't just mean disease epidemics.

ALTHOUGH KERMACK AND MCKENDRICK would later extend Ross's mosquito theorem to other types of infections, Ross had wider ambitions. 'As infection is only one of many kinds of events which may happen to such organisms, we shall deal with "happenings" in general,' he wrote in the second edition of *The Prevention of Malaria*. Ross proposed a 'Theory of Happenings' to describe how the number of people affected by something – whether a disease or another event – might change over time.

Ross suggested that there are two main types of happening. The first type affects people independently: if it happens to you, it generally won't increase or decrease the chances of it happening to someone else afterwards. According to Ross, this could include things like non-infectious diseases, accidents or divorce.[49] For example, suppose there is a new condition that can randomly affect anyone, but at first nobody in the population has it. If each person has a certain chance of becoming affected every year – and remains affected from that point onwards – we'd expect to see a rising pattern over time.
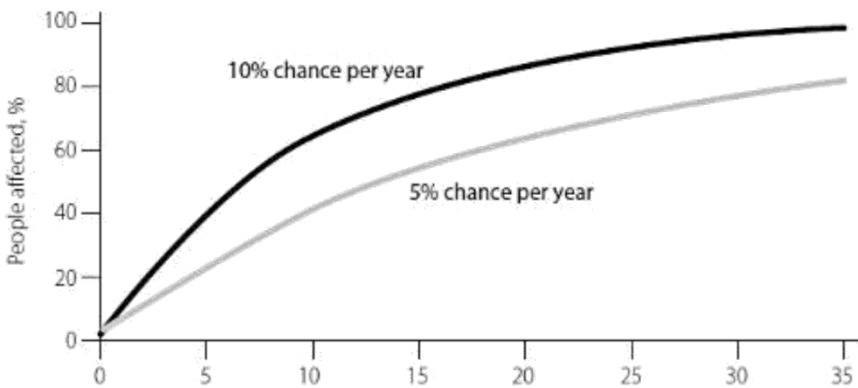

Growth of an independent happening over time. Example shows what would happen if everyone had a 5 per cent or 10 per cent chance of being affected per year

The curve gradually flattens off, though, because the size of the unaffected group shrinks over time. Each year, a proportion of

people who were previously unaffected get the condition, but because there are fewer and fewer of such people over time, the overall total doesn't grow so much later on. If the chance of being affected each year is lower, the curve will grow more slowly initially, but still eventually plateau. In reality, the curve won't necessarily level off at 100 per cent: the final amount of people affected will depend on who is initially 'susceptible' to the happening.

As an illustration, consider home ownership in the UK. Of people who were born in 1960, very few were homeowners by the age of twenty, but the majority had owned a house by the time they were thirty years old. In contrast, people who were born in 1980 or 1990 had a much lower chance of becoming a homeowner during each year of their twenties. If we plot the proportion of people who become homeowners over time, we can see how quickly ownership grows in different age groups.


Percentage of people who were homeowners by a given age, based on year of birth
Data: Council of Mortgage Lenders[50]

Of course, home ownership isn't completely random – factors such as inheritance influence people's chance of buying – but the overall pattern lines up with Ross's concept of an independent happening. On average, one twenty-year-old becoming a

homeowner won't have much effect on whether another gets on the housing ladder. As long as events occur independently of one another at a fairly consistent rate, this overall pattern won't vary much. Whether we plot the amount of people who are on the housing ladder by a certain age, or the chance your bus has arrived after a certain time waiting, we'll get a similar picture.

Independent happenings are a natural starting point, but things get more interesting when events are contagious. Ross called these types of events 'dependent happenings', because what happens to one person depends on how many others are currently affected. The simplest type of outbreak is one where affected people pass the condition on to others, and once affected, people remain so. In this situation, the happening will gradually permeate through the population. Ross noted that such epidemics would follow the shape of a 'long-drawn-out letter S'. The number of people affected grows exponentially at first, with the number of new cases rising faster and faster over time. Eventually, this growth slows down and levels off.

Illustrative example of the S-shaped growth of a dependent happening, based on Ross's model. The plot shows the growth of a more contagious and less contagious happening

The assumption that people remain affected indefinitely doesn't usually apply to infectious diseases, because people may recover, receive treatment, or die from the infection. But it can capture