Almost Alive

How viruses were found, fought and sometimes forgiven



Almost Alive

How viruses were found, fought and sometimes forgiven

Keith McLoughlin & Claude AI

For The Brothers and Their Better Halves
John, Peter, and Declan;
Linda, Helen, and Caroline

Preface

In January 2020 a string of letters was posted to a public database, and within about twenty-four hours a team in Maryland had worked out how to turn it into a vaccine. Two days after that the design was finished. Proving that it worked took the rest of the year, which was the fastest anything of the kind had ever been done, by a factor of roughly ten.

Six years later, measles has re-established itself in six countries of the World Health Organization's European region, and Irish coverage of a vaccine that has been free and routine for decades fell between one quarter and the next.

This book is about the distance between those two paragraphs.

It is a history of viruses and the things people have done about them, and it is arranged in the order the knowledge actually arrived rather than the order that would be tidiest. That means vaccination comes first and viruses come third, because Edward Jenner was vaccinating against smallpox for a

century before anyone could say what smallpox was. It means germ theory has to be fought for in the middle of the book rather than assumed at the start. And it means the thing we are best at — designing a vaccine — turns out to be the most recent and the least difficult part of the story.

Who this is for

One reader, imagined in some detail, and every decision in the book answers to them.

They sat a school science paper some years ago and have let the details go. They are intelligent, curious and out of practice. They have no specialist vocabulary and they are not going to look anything up, because a book that requires its reader to look things up has failed.

So nothing here is assumed to be remembered. Terms are explained where they first appear, including ones a schoolchild has met. Ideas are built in the order a person needs them rather than the order a textbook would use. There is no algebra anywhere.

How it is arranged

Four parts.

Part I runs from 1721 to 1980, from a scratch on a child's arm in Constantinople to the last case of smallpox on earth. Nobody in the first half of it knows what a virus is, and the reader does not need to either.

Part II covers 1981 to 2019: HIV, cancers that turn out to be infectious, a fraudulent paper that is still doing damage, and the two decades of unglamorous work that made Part III possible before anyone knew it would be needed.

Part III is Covid-19, in six chapters. Three of them tell what happened. The other three take questions that are still argued about — where the virus came from, whether the restrictions were worth their cost, and what the vaccine safety findings actually showed — and try to set out the real state of each.

Part IV looks forward from where this was written, in November 2026: at the virus most likely to do this again, at the vaccines now arriving for diseases this book describes as unbeatable, and at the one problem that outlasts every technical success in these pages.

The conventions

Three kinds of box interrupt the text. **KEY CONCEPT** explains an idea the chapter depends on. **HOW WE KNOW** gives the evidence behind a claim — who did the experi-

ment, what they measured, why it convinced people. **THE NUMBER** takes a single figure and works it through, and says what it does not mean as well as what it does.

A reader who skips every box can still follow every chapter. That was a rule during the writing, not a hope.

Chapters refer to each other by number, which lets each one stay short by pointing elsewhere instead of repeating itself.

The later chapters say in their text when they were written. This is not modesty. A chapter about an inquiry that has not reported, or a treaty that cannot yet be signed, or a trial whose result is due in two years, is a photograph of a moving thing, and the reader is entitled to know the shutter speed.

On getting it right

Every date, figure and attribution in this book was checked against a source at the time of writing. Where something could not be checked, it was either cut or marked in the text as uncertain. Several charts that were planned could not be drawn, because the data they needed does not exist in a form anyone has published; in each case the text says so rather than quietly substituting something that looks similar.

This matters more than usual here. A book about vaccines that got its numbers from memory would be exactly the thing it complains about.

Where a question is open, the book says it is open. Where the evidence points hard in one direction, it says that too. Those are different sentences and the difference is the subject.

A note on how this was made

This book was written in collaboration with an artificial intelligence, which is why there are two names on it.

The arrangement was roughly this. The structure, the subject, the judgements about what belonged and what did not, and the final say on every page were mine. The drafting, the searching, the chart-making and a great deal of the argument were Claude's. Each chapter came back with a list of everything in it that rested on a weak source, which is a kind of collaborator I had not had before.

I have not hidden this and I have not made a feature of it. It seemed better to say plainly what happened than to let a reader wonder. Whether the result reads as one voice is for you to judge; it was intended to.

What this book is not

It is not medical advice, and it does not tell anyone what to do about their own health.

The book explains how evidence is gathered and how strong it turns out to be. It does not give doses or schedules as guidance, and where a passage could be read as a personal recommendation it says what it means instead: that decisions of that kind belong with you and the clinician who knows you.

That is not a disclaimer bolted to the front. It is the book's actual position, and the last chapter argues for it at length.

Where this leaves us

The argument of what follows is never stated outright in the chapters, so it may as well be stated here.

Humanity learned to defeat viruses before it knew what they were. Then it learned what they were. Then it built the means to make a vaccine against a new one in under a year. And then it discovered that the hardest part had never been the science.

Contents

Part One — History 11

1 The Milkmaid's Secret 13
2 The Boy from Alsace 33
3 Through the Filter 53
4 The Year the World Caught Cold 73
5 The Summer Disease 95
6 A Daughter's Sore Throat 117
7 The Last Case 139

Part Two — The modern era 159

8 The Virus That Hid 161
9 A Cancer You Can Catch 185
10 Twelve Children 205
11 The Pig, the Bird and the Needle 225
12 Out of the Forest 249
13 The Molecule Nobody Wanted 271

Part Three — Covid-19 291

14 December in Wuhan 293
15 Where It Came From 315
16 The Year Indoors 337
17 Three Hundred Days 361
18 Rare, Real and Rumoured 383
19 The Long Tail 405

Part Four — The future 425

20 The Bird in the Dairy 427
21 Vaccines Against Everything 449
22 The Trust Problem 471
23 The Other Enemy 493

Colophon 509

PART ONE

History

The Milkmaid's Secret

> *"The annihilation of the Small Pox, the most dreadful scourge of the human species, must be the final result of this practice." — Edward Jenner, On the Origin of the Vaccine Inoculation, 1801*

A letter from Adrianople

On the first of April 1717, in the city the Ottomans called Edirne and the English called Adrianople, a woman sat down to write to a friend at home. Lady Mary Wortley Montagu was twenty-seven, the wife of the new British ambassador to the Sultan, and she had spent the winter travelling across Europe toward Constantinople with a small son and a household in tow. She was clever, restless and unusually observant, and the letter she wrote that day to Sarah Chiswell has become the most quoted thing she ever put on paper.

She had a reason to notice what she described. Four years earlier her brother had died of smallpox. Two years earlier she had caught it herself, and though she lived, it had taken

her eyelashes and left her face pitted. Smallpox was not a rare misfortune in the England she had left. It was the ordinary way a family lost a child, a wife lost her looks, a village lost a tenth of its people in a bad year. Nobody knew what caused it. Everybody knew what it looked like: the fever, the rash that rose into hard pustules, the stench, and then either death or a lifetime of scars.

What she had seen in Turkey was a thing the local women did every autumn. They called it engrafting. An old woman would arrive with a nutshell full of matter taken from the sores of someone with a mild case, scratch the skin of a healthy child with a needle, and press a little of the matter into the wound. The child would sicken for a few days, come out in a few spots, and recover. After that, Lady Mary was assured, they never caught the smallpox at all. She wrote that she intended to try it on her own son, and that she would happily fight the English doctors over it when she got home.

She did both. In March 1718, in Constantinople, the embassy surgeon Charles Maitland and a Turkish woman engrafted five-year-old Edward Montagu, one arm each. He had his few days of fever and his few spots and was well. In 1721, with smallpox raging through London, she had Maitland do the same to her small daughter Mary, and this time she made sure there were physicians in the room to watch.



Figure F01-01. *The whole technology of early vaccination. The same lancet that inoculators had used for smallpox served Jenner for cowpox; what changed was what was on the blade.*

Boston, 1721

The same year, on the other side of the Atlantic, a ship arrived in Boston harbour. HMS *Seahorse* came in from Barbados on 22 April 1721 with a crew that had recently been through smallpox, and within weeks the disease was loose in a town of about 11,000 people, most of whom had never met it.

Boston had a man who thought he knew what to do, and he was not a doctor. Cotton Mather was a Puritan minister with

a long memory and a reputation that had not fully recovered from the Salem witch trials. Years before, he had asked an enslaved West African man in his household, whom he had named Onesimus, whether he had ever had smallpox. The answer was yes and no. Onesimus explained that in his own country people cut the skin and put in a little of the matter from a sick person, and so had a mild case on purpose and never a bad one. He showed Mather the scar. Mather later read accounts of the same practice from Turkey, sent to the Royal Society in London by two physicians, Emanuel Timoni and Jacob Pylarini, and when the epidemic came he wrote to every doctor in Boston urging them to try it.

One agreed. Zabdiel Boylston began in June 1721 with his own six-year-old son and two enslaved people in his household, a man and a boy. All three recovered. He went on inoculating through the summer and autumn while the town argued, preached and in one case threw a bomb through Mather's window. By the time the epidemic burned out in February 1722, Boylston had inoculated 282 people and six of them had died. Of the 5,759 Bostonians who caught smallpox the ordinary way, 844 were dead.

Those two numbers, set side by side, are the whole argument for inoculation, and Boylston and Mather understood that they were an argument. They counted. The practice was dan-

gerous, and it was also, by the only test that mattered, seven times less dangerous than taking your chances.

THE NUMBER — One in seven against one in fifty

In Boston in 1721, 5,759 people caught smallpox naturally and 844 died: 844 divided by 5,759 is 0.147, or about one death in every seven cases. Boylston inoculated 282 people and 6 died: 6 divided by 282 is 0.021, about one death in every 47.

So a person who chose inoculation in Boston that year took roughly a one-in-fifty risk of dying from it, against a one-in-seven risk if they caught the disease in the ordinary way. Later and larger counts in England put the inoculation risk nearer one in a hundred, and the natural risk for smallpox as a whole nearer one in three or four, because Boston's 1721 strain was milder than many.

What the number does not say is just as important. It compares the inoculated with everyone who fell ill, not with everyone exposed; some of the uninoculated never caught it at all. It says nothing about whether an inoculated person could pass smallpox to others, which they could. And it is a comparison, not a trial: nobody was choosing at random who got the needle. The strength of the figure is that the gap is enormous. A difference of

> seven to one survives a great deal of sloppiness in the counting.

Spreading the needle

Back in London, Lady Mary's demonstration on her daughter did what she intended. The physicians who watched went away and talked. The Princess of Wales took an interest. In August 1721 a handful of condemned prisoners in Newgate were offered their lives in exchange for being inoculated under observation; they survived and were freed. Two royal daughters were inoculated the following year. Within a generation, inoculation was an established trade in England, practised by surgeons who travelled from village to village and charged by the arm.

It was never universal and never safe. A person who had been inoculated was infectious for a fortnight, and careless practitioners started epidemics as often as they prevented them. In the 1760s a Suffolk family named Sutton made a fortune by refining the method, using a shallower scratch and a smaller dose, and claiming far lower death rates. Opponents pointed out that an inoculated child in a village without smallpox could kill a dozen neighbours. Both sides were right. For the better part of the eighteenth century, that

was where things stood: a crude, risky, effective practice, borrowed from Africa and the Ottoman Empire, administered by people who had no idea why it worked.

KEY CONCEPT — Why you cannot catch it twice

Everyone in the eighteenth century knew one fact about smallpox that mattered more than any other: a person who had survived it never had it again. Nurses who had been through it could tend the sick without fear. Families looked for a servant with a pitted face.

Nobody could explain this, and this book will not be able to explain it properly until later chapters. For now, the idea is simple enough to state. The body keeps a memory of what has attacked it. The first time a particular disease arrives, the body is slow to answer and the person is ill for weeks, if they live. If the same disease arrives again, the body recognises it and answers so fast that the person often never notices. This is immunity, and it is specific: having had smallpox protected you from smallpox and from nothing else.

Inoculation worked by exploiting this memory on purpose. Give someone a small, deliberate dose through the skin, where the disease seems to run milder than when it is breathed in, and if they survive the mild case they

> have the memory without having risked the full disease. The trick is old and the explanation is modern, and for seventy years the trick was all anyone had.

A country doctor’s notebook

Edward Jenner was inoculated as a schoolboy in Gloucestershire and remembered the experience as a misery: weeks of purging, bleeding and starvation beforehand, as the practice of the time required, and then the disease. He trained as a surgeon in London under John Hunter, the most original anatomist of the age, who taught him to distrust books and test things. Then he went home to Berkeley, a small town between the Severn and the Cotswolds, and practised medicine among farmers for the rest of his life.

In dairy country there was a piece of folk knowledge that every farmhand knew and no physician took seriously. Cows got a disease on their udders called the cowpox, a crop of sores that the milkers caught on their hands. It was unpleasant and it passed. And people who had had the cowpox, the farmers said, did not take the smallpox. Milkmaids were famous for clear complexions in an age when most faces were marked.

Jenner's own account, written in 1801, says the idea came to him from the inoculators' trade rather than from any milkmaid. Practitioners in the county had found that some patients simply could not be given smallpox by inoculation, however many times they tried, and that these people had usually had the cowpox. A farmer told him so in as many words. For two decades Jenner collected such cases, asked questions, examined hands and udders, and wrote it all down. He also worked out something the farmers had not: that there was more than one disease on the cows, and that only the true cowpox protected. Many of his failures to find protection turned out to be cases of the wrong pox.

He was not the first to act on the idea. In 1774 a Dorset farmer named Benjamin Jesty, during a smallpox outbreak, had taken matter from a cow's udder on a stocking needle and scratched it into the arms of his wife and two sons. They were never known to catch smallpox, and the neighbours thought him a lunatic. Jesty did not publish, did not test, and did not pursue it. Jenner did all three, and the difference between the two men is the difference between knowing a thing and proving it.

The boy and the cow

In May 1796 a dairymaid named Sarah Nelmes came to Jenner with cowpox sores on her hand, caught from her master's cows. Local tradition names the cow as Blossom, whose hide now hangs in a London medical school, and whether or not the name is right the sores were real. On 14 May, Jenner took matter from one of them and scratched it into two small cuts on the arm of James Phipps, the eight-year-old son of his gardener.

The boy ran a slight fever, felt out of sorts for a day or two, and recovered. Then came the part that made the experiment an experiment. On 1 July, Jenner inoculated James Phipps with smallpox itself, in the usual way, from a fresh case. Nothing happened. No fever, no spots. He did it again, later, with the same result. The cowpox had done what seventy years of inoculation had done, without the smallpox.

It is hard to read this now without flinching. Jenner deliberately exposed a child to the most lethal disease of the age on the strength of farm gossip and a notebook. By the standards of his time it was not reckless: the boy would have been inoculated anyway, and the danger from the test was no greater than from the standard practice. By the standards of ours it

would not be permitted. Both things are true, and the book will meet the tension between them more than once.

He sent an account to the Royal Society. They declined to publish it, and a friend warned him that he should not risk his reputation on a single case. So Jenner went looking for more cows.

There was no cowpox in the district in 1797, and Jenner could do nothing. In 1798 it returned, and he worked quickly: matter from a cow to a child, from that child to another, and on down a chain of arms, testing as he went. In the summer of 1798 he published, at his own expense, a slim book with a long title: *An Inquiry into the Causes and Effects of the Variolae Vaccinae*. The Latin means "the smallpox of the cow", and from *vaccinae*, the cow's, we have the word vaccine. Jenner coined it. He did not yet know how large a word it would become.

> **HOW WE KNOW — The twenty-three cases of 1798**
>
> The *Inquiry* is not a theory. It is a list. Jenner set out 23 numbered cases, and the argument is in their arrangement.
>
> The first group were people who had caught cowpox naturally from cattle, some of them decades earlier, and

whom Jenner or other practitioners had since tried to inoculate with smallpox. In case after case the inoculation did not take: no fever, no pustule, the arm barely reacting. These cases showed that the folk belief was true, but they could not rule out that these people were simply resistant for some other reason.

Case XVI is Sarah Nelmes and her infected hand. Case XVII is James Phipps, deliberately given cowpox from that hand and then deliberately challenged with smallpox. This is the step that turned observation into proof: Jenner chose the subject, controlled the exposure and chose the test. The later cases carried cowpox from one child's arm to the next, so that each new child was protected by the one before, and then challenged them too.

What convinced his readers was not any single case but the combination. Natural cowpox protected; deliberate cowpox protected; cowpox passed from arm to arm still protected; and in every instance the test was the same one every inoculator in England used daily. A sceptic could repeat it, and within two years hundreds had.



Figure F01-02. *The source of everything in this chapter. A mild disease of cattle, caught by hand, turned out to carry protection against the deadliest disease of people, and the reason would not be understood for more than a century.*

Spreading the arm-to-arm chain

The *Inquiry* was attacked, mocked and, more dangerously, tried by people who got it wrong. Matter taken too late from a sore did not protect. Matter contaminated with smallpox gave smallpox. A London surgeon vaccinated a group of children from a sample that turned out to be mixed, and the resulting cases were held against Jenner for years. He spent much of the rest of his life answering letters, sending threads

soaked in cowpox lymph to anyone who asked, and defending the difference between true and spurious cowpox.

But the practice spread faster than the arguments against it, because the alternative was so plainly worse. There was no need for a cow in every parish. Once a child had been vaccinated, the fluid from the sore on that child's arm, taken on the eighth day or so, would vaccinate the next. A single successful vaccination could be carried across a country from arm to arm, a week at a time, and it was. Within a few years it had reached India, the Americas and Russia. In 1803 Spain despatched an expedition to its colonies carrying the vaccine in the only container that would keep it alive on a long voyage: a succession of orphan boys, two vaccinated at a time, each pair passing the cowpox to the next before their own sores healed.

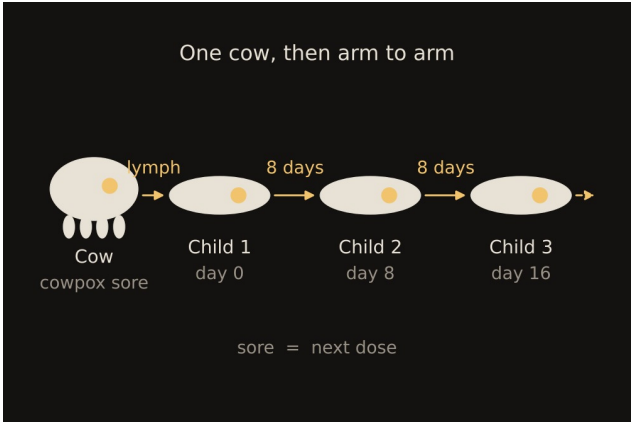


Figure F01-03. *Why vaccination could outrun smallpox without a laboratory. Cowpox needed no cow after the first case; the sore on one vaccinated arm supplied the next, so a single chain could cross a continent one week at a time.*

Dublin’s Cow-Pock Institution

Ireland took up vaccination almost as soon as England did, and in one respect ahead of it. Within a few years of the *Inquiry*, vaccination was being offered to children at the Dispensary for the Infant Poor in Dublin, and physicians in Cork had vaccinated on a scale large enough to be talked about. In January 1804, a group of Dublin subscribers opened an institution with no other purpose than to vaccin-

ate: the Cow-Pock Institution, first in North Cope Street and later in a townhouse on Sackville Street, where a subscription book opened in March of that year survives in the archive of the Royal College of Physicians of Ireland.

The Institution did two things. It vaccinated the poor of the city without charge, which in a city as crowded and as poor as Dublin was no small undertaking. And it kept the chain alive. Because the vaccine could only be had from a fresh human sore, the Institution asked each vaccinated child to return a week later so that lymph could be taken from the arm, dried on threads or sealed between glass, and sent out in packets to practitioners across the country. A physician in Mayo or Kerry who wanted to vaccinate wrote to Sackville Street. In 1805 the Institution's physician, Samuel Labatt, published an address to the medical practitioners of Ireland setting out how to do it properly, with a plate showing the stages of the sore and an appendix of letters from doctors around the country reporting their results.

By 1811 the directors were able to report that they had seen no case of smallpox after what they called a perfect vaccination. The qualification mattered, and it had been Jenner's own: a sore that did not develop properly, or lymph taken at the wrong time, protected no one. But the Institution survived, on Sackville Street, into the second half of the cen-

tury, and for decades it was the point from which the vaccine reached the rest of the island. When Parliament extended vaccination by statute in Ireland in 1840, it was building on a network that a committee of Dublin subscribers had kept going, arm to arm, since the year Jenner's book was six years old.



Figure F01-04. *Where the chain was kept alive. An institution that existed only to vaccinate also had to be a source of vaccine, taking lymph from each week's children to supply the next week's and the rest of the country.*

What the bills showed

London had been counting its dead for a century and a half. Every week the parish clerks published the Bills of Mortality, listing burials by cause, and smallpox had its own column. Through the eighteenth century the column held steady at something between six and ten of every hundred burials in the city, year in and year out, with a worse epidemic every few years. Inoculation, for all its spread, had barely dented the London figure; the practice was patchy, and in a great city the inoculated may have spread as much as they prevented.

After 1800 the column began to fall. In the first decade of the new century smallpox was about seven in a hundred London burials. In the 1810s it was under five. In the 1820s it was between three and four. The city had grown, the counting had not improved, and nothing else about London had changed for the better; it had simply been vaccinating for a generation. The fall was not steady and it was not complete. There were bad years, and a serious epidemic in 1838 when the practice had been allowed to lapse. But the direction was plain to anyone who read the bills, and people did.

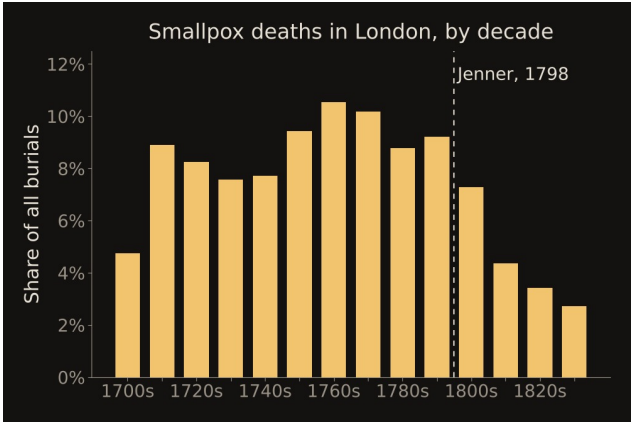


Figure F01-05. *The first evidence that a vaccine could change a city. Smallpox's share of London burials held near a tenth for a hundred years of inoculation, then fell by more than half in the thirty years after Jenner's Inquiry. Drawn from Charles Creighton's tabulation of the Bills of Mortality.*

Where this leaves us

In 1801 Jenner allowed himself a prediction. The annihilation of smallpox, he wrote, must be the final result of this practice. He was right, and it would take 179 years.

Consider what he and his contemporaries did not know. They did not know what smallpox was. They did not know what cowpox was, or why one should protect against the

other. They had no idea what happened inside an arm that made it refuse a second infection. They could not see the thing they were fighting, could not grow it, could not measure it, and would not be able to do any of those things for another hundred years. They worked entirely by result. A thing was tried, people were counted, and the practice that left fewer of them dead was kept.

That is a respectable way to proceed, and it is worth remembering, because much of what follows in this book is the story of understanding catching up with practice. The Turkish women and Onesimus knew that a mild case prevented a bad one. The dairymaids knew that cowpox kept off smallpox. Jenner knew how to turn that knowledge into a test, and a committee in Dublin knew how to turn the test into a service. None of them knew why.

The next chapter stays in the nineteenth century and watches a chemist in Paris, with no better idea of what a germ was than Jenner had of a virus, work out that Jenner's trick was not a trick at all. It was a method, and it could be made to work for other diseases, if only one could find the cow.

The Boy from Alsace

> *“In the fields of observation, chance favours only the prepared mind.” — Louis Pasteur, address at Douai on the opening of the Faculty of Sciences of Lille, 7 December 1854*

A knock at the Rue d’Ulm

On the morning of Monday 6 July 1885, a woman from Alsace arrived at a laboratory in the Rue d’Ulm in Paris with her nine-year-old son. Two days earlier, on his way to school, the boy had been attacked by a dog and bitten fourteen times, on the hands, the legs and the thighs. The dog’s owner had shot it, and when its stomach was opened it was found to be full of hay, straw and splinters of wood, which is what a rabid dog eats. A local doctor had cauterised the worst of the wounds with carbolic acid, twelve hours after the attack, which was too late to do any good, and told the mother that there was a man in Paris who was said to be working on rabies.

The boy's name was Joseph Meister. The man was Louis Pasteur, who was sixty-two, a chemist by training, half paralysed down his left side since a stroke seventeen years earlier, and not a doctor. He had never treated a human being for anything.

Rabies was then, as it is now once symptoms appear, a disease with no survivors. The bite of a rabid animal did not always transmit it; perhaps one bitten person in six developed the disease, and more if the bites were deep or to the face. But those who did develop it, after weeks or months in which they felt perfectly well, died within days, raving, convulsing at the sight of water, and conscious for much of it. There was no treatment except cautery, and the fear of the disease was out of all proportion to its frequency. It is worth holding that fear in mind. It explains why a mother would take a child across France to a chemist, and why the chemist would agree to see her.

Pasteur did not decide alone. That evening he put the case to two physicians, Alfred Vulpian and Jacques-Joseph Grancher, and they examined the boy. Fourteen bites, some deep. In their judgement he would almost certainly die. Pasteur had a preparation that had protected dogs; he had never given it to a person, and he was not sure it was safe. Vulpian and Grancher told him that the only alternative was to do noth-

ing, and that doing nothing was a sentence. At eight o'clock that night Grancher, because Pasteur was not licensed to do it himself, injected Joseph Meister under the skin of his right side with half a syringe of a liquid made from the dried spinal cord of a rabbit that had died of rabies fifteen days before.

What a chemist knew about disease

To understand what was in the syringe, the reader has to go back thirty years and meet the idea that made it possible, because Jenner had not had it.

In Jenner's lifetime nobody knew what caused any infectious disease. The leading view blamed bad air, miasma, rising from marshes and filth. Alongside it ran an older notion, never fully worked out, that some diseases passed from person to person by a seed or ferment of some kind. Pasteur came to the question sideways. As a young chemist he had studied crystals and then fermentation, the process by which sugar becomes alcohol in a vat or milk turns sour in a jug. Chemists of the time thought fermentation was a purely chemical event, a kind of decay. Pasteur showed, through the 1850s, that it was done by living things: tiny organisms, visible under a microscope, each kind producing its own characteristic change. Yeast made alcohol. A different, smaller

organism made lactic acid and soured the milk. Kill the organisms with heat and the change stopped.

This raised an obvious question: where did the organisms come from? Many scientists believed that they arose spontaneously from the liquid itself, that life could begin afresh in any suitable broth. In a series of experiments around 1860, Pasteur showed that they did not. He boiled broth in flasks whose necks he had drawn out into long, downward-curving tubes, open to the air but shaped so that dust settled in the bend. The broth stayed clear for months. Snap the neck off, or tip the flask so the broth touched the dust, and within days it was cloudy with life. The organisms came from outside, carried on particles in the air. They did not generate themselves.

From there it was a short step, though it took the best part of two decades and a great many arguments, to the idea now called germ theory. If each fermentation is the work of a specific microscopic organism, perhaps each infectious disease is too. Pasteur spent the late 1860s on diseases of silkworms, which were ruining the French silk industry, and found microscopic parasites in the sick moths. In Germany, Robert Koch showed in 1876 that anthrax, a disease of sheep and cattle that also killed people, was caused by a rod-shaped organism he could see, grow and transfer. The word for such

an organism was microbe, or germ, and the smaller rod-shaped kinds came to be called bacteria.

KEY CONCEPT — Germ theory

Germ theory is the claim that infectious diseases are caused by living things too small to see, that each disease has its own kind, and that the disease spreads when the organism passes from one body to another.

That sounds obvious now. In 1870 it was a radical simplification. It said that cholera was not a quality of the air or a punishment or an imbalance in the body but a specific thing, which could in principle be found, grown, studied and stopped. It said that the same disease was the same everywhere, because the same organism caused it. And it said that whatever protected a person against a disease was protecting them against that organism in particular, which is why having had smallpox was no help against measles.

The theory also ruled things out. If a disease required its germ, then a body that had somehow learned to defeat the germ could not catch the disease, however foul the air. That is the connection to Jenner. Vaccination now had a possible explanation: the cowpox germ, whatever it was, so resembled the smallpox germ that a body

> which had beaten one could beat the other. Nobody could yet see either germ. The theory said they must exist, and said what to look for.

By the end of the 1870s Pasteur, who had never studied medicine, was running a laboratory devoted to diseases, and had come to a conviction that most of his medical contemporaries thought absurd. He believed that what Jenner had done for smallpox could be done for other diseases: that for each germ there might be a weakened form that would teach the body to resist the full one. The difficulty was that Jenner had been handed his weakened form by the cows. Nobody knew how to make one.



Figure F02-01. *The shape that settled an argument. Open to the air but bent so that dust could not reach the broth, the flask stayed clear for months, and showed that the organisms of decay came from outside and did not arise on their own.*

Forgotten chickens

The usual story is that chance made it for him. In 1879 Pasteur's laboratory was studying chicken cholera, a disease that swept through farmyards and killed birds within a day or two, and had managed to grow its germ in flasks of broth. Before the summer holiday, the story goes, Pasteur told an assistant to inject a batch of chickens with a fresh culture. The assistant did not, and went on holiday himself. When the

laboratory reassembled in the autumn, the old cultures, which had stood in their flasks for weeks, were injected into chickens anyway. The chickens sickened slightly and recovered. And when the same chickens were then given a fresh, lethal culture, they did not die.

How much of that is true is uncertain; Pasteur's notebooks do not record the holiday incident in the form the story gives it, and his own published account is more guarded and more gradual. What is not in doubt is the conclusion he drew and tested through the following months. A culture of the chicken cholera germ that was left standing, exposed to the air, slowly lost its power to kill. It still produced a mild illness, and birds that had survived the mild illness were protected against the deadly one. He called the loss of killing power attenuation, a weakening, and the word has been used for vaccines of this kind ever since. Crucially, the weakening could be controlled. Cultures sealed from the air kept their full force; cultures left open for different lengths of time gave a graded series, from lethal to harmless.

This was the thing Jenner had never had. A method. Not one lucky cow for one disease, but a procedure that might, in principle, be applied to any germ that could be grown: grow it, weaken it, and give the weakened form on purpose. Pasteur, with a sense of history that his rivals found insufferable,

proposed that the word vaccine, which had meant only Jenner's cowpox, should now mean any such preparation, in Jenner's honour. The proposal stuck.

HOW WE KNOW — Pouilly-le-Fort, 1881

Pasteur's next target was anthrax, and his claim to have made a vaccine against it was met with open disbelief from veterinarians. So he agreed to a public trial, on a farm at Pouilly-le-Fort near Melun, with the terms set in advance and the press invited.

On 5 May 1881 and again on 17 May, his team vaccinated 24 sheep, a goat and 6 cows. Another 24 sheep, a goat and 4 cows were left untreated. On 31 May all of the animals were injected with a fresh, fully virulent culture of anthrax. On 2 June the visitors returned. Every vaccinated animal was healthy. Every unvaccinated sheep and the unvaccinated goat were dead or dying by the end of the day, and the unvaccinated cows were visibly ill.

As a demonstration it was close to perfect: two matched groups, one difference between them, the same challenge for both, and an outcome that needed no statistics. What it did not reveal, and what Pasteur did not say, was how the vaccine had been made. Everyone assumed it

had been weakened by the air, like the chicken cholera germ. It had not. The laboratory notebooks, which Pasteur ordered his family never to release and which were made public only in 1971, show that the vaccine used at Pouilly-le-Fort had been weakened with a chemical, potassium dichromate, by a method close to one his rival Henri Toussaint had published the year before. Pasteur's own method was not yet reliable, and he knew it. The sheep lived. The account of why was not the whole truth.



Figure F02-02. *Where vaccination became a method. The tools were simple, but the bench could grow a germ, weaken it to order and test the result, which is everything a vaccine maker has needed since.*

Thirteen injections

Anthrax and chicken cholera were caused by bacteria, which could be grown in broth and seen under a microscope. Rabies was different, and the difference is the hinge of this book. Pasteur and his assistants, Émile Roux and Charles Chamberland among them, hunted for the rabies germ for years and never found it. Nothing grew in their flasks. Nothing showed under the lens. The disease was plainly caused by something that passed from a rabid animal’s saliva into a

wound and travelled to the brain, and whatever that something was, it was too small to see with any microscope then in existence. Rabies, as chapter 3 will explain, is caused by a virus, and Pasteur was trying to make a vaccine against a thing that nobody would see for another half century.

He did it anyway, by treating the invisible agent as if it were a bacterium he could not culture. If the germ lived in the nervous system, then nervous tissue from a rabid animal must contain it. His team injected material from rabid dogs directly into the brains of rabbits, and found that after many passages from rabbit to rabbit the disease became faster and more regular, arriving reliably on the sixth or seventh day. They then took the spinal cords of these rabbits and hung them in sterile jars with a drying agent at the bottom. Each day of drying weakened whatever was in the cord. A cord dried for a fortnight no longer caused rabies at all. A cord dried for one day was nearly as dangerous as the fresh disease.

That gave them a graded series, like the open flasks of chicken cholera, and they used it on dogs. A dog injected with the fourteen-day cord one day, the thirteen-day cord the next, and so on down to fresh cord, did not get rabies from the fresh cord. Nor did it get rabies when it was afterwards bitten by a rabid animal. By the summer of 1885 Pasteur had

put fifty dogs through the series, and all fifty had resisted. He had not tried it on a person, and he was afraid to.

Joseph Meister forced the decision. Over ten days, from the evening of 6 July, Grancher gave the boy thirteen injections, each from a cord that had been drying for a shorter time than the last, so that the final doses were from cords only a day or two old, which would have killed an untreated rabbit. Pasteur, who slept badly throughout, wrote to his family that the boy was eating and sleeping and that he himself could not. On the last day, as a test, the final dose was from a cord virulent enough to have caused the disease on its own. Nothing happened. Meister went home to Alsace at the end of July, and he never had rabies.

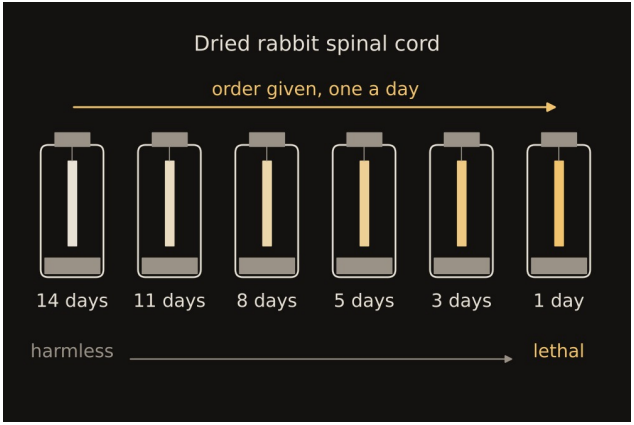


Figure F02-03. *How to weaken a germ you cannot see. Drying the infected cord for a known number of days gave Pasteur a series of doses from harmless to lethal, and giving them in order taught the body to resist the last before it arrived.*



Figure F02-04. *The rabies vaccine of 1885 in its entirety. A dried spinal cord in a jar held a germ nobody could see, weakened by a number of days that could be counted.*

A hundred in six weeks

Three months later a second patient arrived: Jean-Baptiste Jupille, a fifteen-year-old shepherd from the Jura who had been savaged while fighting off a rabid dog that had attacked a group of younger children. Six days had already passed since the bite. Pasteur treated him, and he too lived. On 26 October 1885, Pasteur told the Académie des sciences what he had done. Within six weeks a hundred bitten people had come to the Rue d’Ulm. By the end of February 1886 there

had been 350. By the end of October 1886, when Pasteur reported again, nearly 2,500 people had made the journey to Paris from across Europe, including a party of Russian peasants mauled by a rabid wolf, most of whom survived.

THE NUMBER — One in six, and then one in 170

Not everyone bitten by a rabid dog gets rabies. Figures compiled in the decades after Pasteur, when many bites still went untreated, put the risk at about 18 per cent overall, roughly one person in six, rising to nearly one in two for bites to the head or face and falling to about one in twenty for a bite on a hand or foot. Joseph Meister, with fourteen deep bites, was at the dangerous end of that range.

Pasteur's first full year gives the other number. Of about 1,700 French patients treated in Paris between November 1885 and October 1886, 10 died: 10 divided by 1,700 is 1 in 170, well under 1 per cent. Two more died who had arrived more than a month after being bitten, too late for the treatment to work, and Pasteur counted them separately, which his critics said he should not have.

Here is what the comparison cannot do. Nobody knew how many of the 1,700 had been bitten by dogs that

were actually rabid, or how many would have fallen ill untreated. If the true untreated risk was one in six, the treatment had saved something like 270 lives. If many of the dogs were not rabid at all, it had saved far fewer. Pasteur's own best evidence was indirect: in the year before his method, 60 people had died of rabies in the hospitals of Paris over five years, and in the year after, 3, of whom 2 had not been treated. It was persuasive. It was not proof, and the proper design for proof, which chapter 5 describes, was seventy years away.

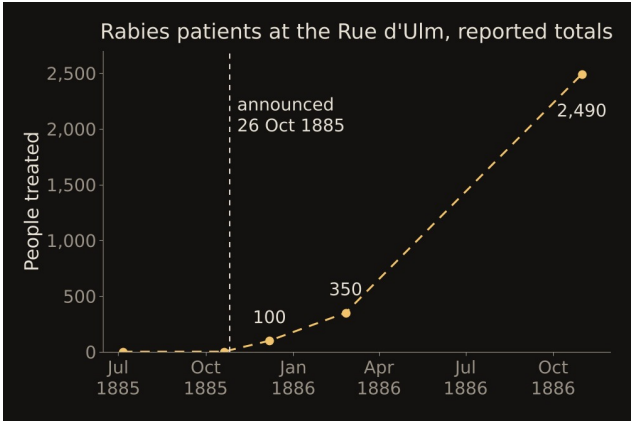


Figure F02-05. *How fast a cure travels. One child in July 1885, a hundred within six weeks of the announcement, and nearly 2,500 within a*

year, from across Europe, before anyone could say with certainty how well the treatment worked. Drawn from Pasteur's reports to the Académie des sciences as printed in 1886.

The notebooks reopened

Pasteur died in 1895, a national hero with an institute named after him, and for most of the next century his story was told the way he had told it. Then in 1971 his grandson gave his laboratory notebooks, more than a hundred of them, to the Bibliothèque nationale, against Pasteur's own explicit instruction that they never be shown. The historian Gerald Geison spent twenty years reading them, and in 1995 published what he had found.

The notebooks do not show a fraud. They show a man in a hurry, in a race, who presented his results in the best light and did not always say what he had actually done. The Pouilly-le-Fort vaccine was one case. The rabies treatment was another. When Meister arrived, Pasteur had not in fact completed the series of dog experiments that his public account implied. He had tested a slightly different method on dogs, and the exact course of injections given to the boy had never been tried in full on any animal. He was improvising, under pressure, with a child, and he knew it. Two earlier patients, treated quietly in the spring of 1885 by a different

method, do not appear in the public account at all, and neither case settled anything.

None of this changes what happened to Joseph Meister or to the 2,500 who followed him. The treatment worked, and versions of it, grown in cells rather than rabbits, are given to millions of people every year. But it is right to know how close to the edge the first step was taken, because the question it raises does not go away. A desperate patient, an untested treatment, a scientist sure of himself and short of proof: the book will meet that scene again in 2020, and the arguments on both sides will be the same.

Meister himself grew up, served in the army, and in later life was employed as a caretaker at the Pasteur Institute. He died on 24 June 1940, ten days after the German army entered Paris, and a story grew up that he had killed himself rather than open Pasteur's tomb for the soldiers. The contemporary record and his own family tell it differently: he had sent his wife and children away to safety, stayed to look after the Institute, and took his own life believing, wrongly, that they had been caught in the invasion. They came home hours later. He was sixty-four.

Where this leaves us

Between Jenner's cow and Pasteur's rabbit, vaccination changed from an accident into a method, and the method rested on an idea that neither man could prove. Germ theory said that each disease had its organism. Pasteur showed that an organism could be weakened in a controlled way and given deliberately, and that the body which had met the weak form would defeat the strong one. For bacteria, which he could grow and see, the method was complete by 1881. For rabies, which he could neither grow nor see, it worked anyway, and the fact that it worked was the first evidence that there were germs of a different kind, smaller than any bacterium, that obeyed the same rules.

Pasteur never saw one. Nobody did. By the time he died, the thing in the dried spinal cord had still not been found, and the best that could be said of it was that it passed through places a bacterium could not. In the year of his death, a Russian botanist and a Dutch microbiologist, working on a disease of tobacco plants, were about to say exactly that, and to begin the chapter in which the subject of this book finally appears.

Through the Filter

> *“Concerning a contagium vivum fluidum as cause of the spot disease of tobacco leaves.” — Martinus Beijerinck, title of his paper to the Royal Academy of Sciences, Amsterdam, 1898 (trans. James Johnson, 1942)*

Spotted leaves

The subject of this book was found in a tobacco field, by people who were trying to save a crop and not to found a science.

In the 1880s, growers across Europe were losing tobacco to a disease that mottled the leaves with light and dark patches, stunted the plants and ruined the harvest. At the agricultural station at Wageningen in the Netherlands, a German chemist named Adolf Mayer studied it, gave it the name it still carries, mosaic disease, and showed in 1886 that it was infectious: sap pressed from a sick plant and rubbed into a healthy one gave the healthy one the disease. He looked for the germ, as germ theory told him to, and found nothing he

could name. He supposed it must be a bacterium too small or too delicate for his methods.

In 1892, in St Petersburg, a young Russian botanist named Dmitri Ivanovsky, who had been sent to the tobacco plantations of the south to investigate the same disease, did the experiment that matters. He took sap from a diseased plant and passed it through a filter of unglazed porcelain. These filters had been designed in Pasteur's laboratory by Charles Chamberland eight years earlier for a practical purpose: their pores were so fine that no bacterium could pass, and water forced through them came out sterile. Pasteur's whole approach to disease rested on the assumption that the filter caught the germ. Ivanovsky forced his tobacco sap through one, took the clear liquid that came out the other side, and rubbed it on a healthy plant.

The plant got mosaic disease.

Ivanovsky reported the result to the Academy of Sciences and drew the cautious conclusion: perhaps the bacterium made a poison that passed through, or perhaps a few very small bacteria had slipped through a flaw in the filter. He did not pursue it. He had found the first evidence of something new in biology and had not recognised it, which is the usual way such things go.



Figure F03-01. *Where virology began. The disease that cost growers their crop in the 1880s was the first to be traced to an agent smaller than any germ, and its virus remains the most studied of all.*

A living fluid

Six years later, Martinus Beijerinck repeated the experiment in Delft without knowing it had been done. Beijerinck was a microbiologist of great originality and difficult temper, a man who never married, never studied a human disease and spent his career on the organisms of soil and water. He filtered the sap, as Ivanovsky had, and got the same result. Then he kept going, and the extra steps are what made 1898 the year viruses were discovered rather than 1892.

He showed that the infectious agent in the filtered sap was not a poison, because it multiplied. A drop of filtrate rubbed on one plant produced a diseased plant whose sap could infect another, and another, indefinitely, with no loss of strength. A poison would have been diluted away. He showed that it would not grow in any broth or on any jelly that grew bacteria; it multiplied only inside living plant tissue, and only, he noticed, in the parts of the plant that were themselves growing and dividing. And he showed that it would soak slowly through a block of jelly, as a dissolved substance does and a bacterium does not.

From this he concluded that the agent was not a tiny organism at all but something of a different order, which he called a *contagium vivum fluidum*: a contagious living fluid. The phrase is wrong, as it turned out, since the agent is made of particles and not fluid, but the thought behind it was right and new. Here was something that reproduced, and so was in some sense alive, but that had no cell, could not be grown alone, and passed through a filter like a dissolved salt. For want of a better word Beijerinck reached for an old one that physicians had used loosely for centuries to mean a poison or a morbid principle. He called it a virus. The word has meant this, and only this, ever since.

HOW WE KNOW — The filter experiment

The Chamberland filter was a hollow candle of unglazed porcelain, closed at one end. Liquid was forced through its wall under pressure, and the pores in the porcelain, a fraction of a thousandth of a millimetre across, held back anything larger. The test that it worked was simple: pass a broth teeming with bacteria through it, and the liquid that emerged, put in fresh broth and kept warm, stayed clear. Nothing grew. Every bacterium then known was stopped.

So when infected tobacco sap passed through the same filter and still caused disease, there were only three possibilities. The disease was caused by a poison dissolved in the sap; or the filter had a crack; or the cause was something alive that was smaller than the pores.

Beijerinck's additional experiments eliminated the first two. A poison does not multiply, and this agent did, through an unlimited series of plants. A cracked filter lets bacteria through, and no bacteria could be found or grown from the filtrate however hard anyone tried. That left the third, and the third was a thing that nobody had imagined: an agent of disease that was alive by the test of reproduction and invisible by every instrument in ex-

> istence. The filter did not show what a virus was. It showed, by subtraction, that there had to be one.



Figure F03-02. *The instrument that found what it could not hold. Designed to catch every bacterium, the porcelain candle let the cause of mosaic disease pass, and the gap between what it stopped and what it did not was the first measurement of a virus.*

The cattle at Greifswald

In the same year, 1898, two of Robert Koch's former students in Germany did for animals what Beijerinck had done for plants. Friedrich Loeffler and Paul Frosch, working at Greifswald on the Baltic coast, were investigating foot-and-

mouth disease, which was costing German farmers their cattle. They took fluid from the blisters of sick animals, passed it through a porcelain filter, and found that the clear filtrate still gave calves the disease. They, too, showed that it multiplied, since fluid from the newly infected animals could be filtered and passed on again. Unlike Beijerinck they held to the more conservative view that it was a very small microbe, too small to see; but the result was the same, and it was the first disease of animals shown to be caused by one of these filter-passing agents.

Within a few years there was a list of them. In 1901 a United States Army commission in Cuba, building on the Cuban physician Carlos Finlay's long-mocked insistence that yellow fever was carried by mosquitoes, showed that the disease passed through a filter: the first virus of human beings. Others followed, found one by one by the same method. Rabies was on the list, which closed a question Pasteur had left open: the thing in his dried spinal cords that would not grow and could not be seen was one of these. So was smallpox, and so was cowpox, which meant that Jenner's whole practice had been the deliberate use of one virus to protect against another. For the first two decades of the twentieth century the agents were simply called filterable viruses, after the one thing anyone knew about them, and the list was a

catalogue of absences: no growth in broth, no body under the microscope, no stain that caught them.

KEY CONCEPT — What a virus is

Every living thing the reader has heard of, from a bacterium to a whale, is made of cells. A cell is a tiny bag of watery chemistry, bounded by a skin, that takes in food, burns it for energy, builds its own parts, and divides to make two. A bacterium is one such cell living alone. A human being is tens of trillions of them living together. The instructions for building and running a cell are written in long molecules inside it, its genes, and every cell has the machinery to read those instructions and carry them out.

A virus is not a cell. It is a set of genes wrapped in a protective coat of protein, and sometimes an outer envelope of fat borrowed from the last cell it left. It has no machinery of its own. It cannot take in food, cannot burn anything for energy, cannot build a part or divide. Outside a cell it does nothing at all, and can sit for hours or years as inert as a grain of salt.

What it can do is get into a cell. Its coat is shaped to fit a particular kind of cell, in a particular kind of host, and when it finds one it attaches, and its genes go inside.

From that moment the cell's own machinery, which cannot tell the difference, reads the virus's instructions instead of its own and builds copies of the virus: hundreds or thousands of them, until the cell bursts or is used up. The copies spread to the next cell, and the next host, and the process repeats.

That is why a virus cannot be grown in broth, why it needs living tissue, and why it passes through a filter: it is not an organism in the ordinary sense but a message that hijacks organisms, and it is small because a message can be. It is also why, in the end, a vaccine works against it. The body can learn to recognise the coat, and a virus that is stopped at the door never gets to deliver its instructions.

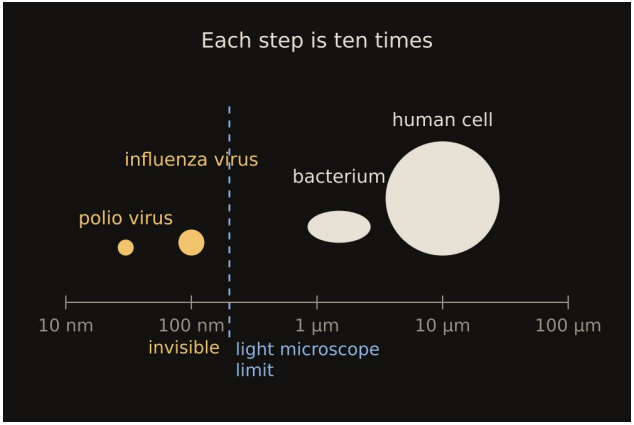


Figure F03-03. *Why the virus was found by subtraction. Each step down the ladder is roughly tenfold, and the limit of what light can show falls between the bacterium and the virus, so that one could be seen and the other only inferred.*

Something that eats bacteria

The next virus to be found was not a disease of anything the reader would care about. It was a disease of bacteria, and it was discovered twice.

In 1915 Frederick Twort, a London bacteriologist, noticed that some of his bacterial colonies, growing on plates of jelly, developed glassy transparent patches where the bacteria had dissolved, and that a trace of material from such a

patch, touched onto a fresh colony, made it dissolve too. He published a short paper, speculated that it might be a virus of bacteria, and was taken off to the war. In 1917 Félix d'Hérelle, a self-taught French-Canadian working without pay at the Pasteur Institute in Paris, found the same thing in the stools of soldiers recovering from dysentery: a filterable agent that destroyed the dysentery bacterium in a tube of broth, turning a cloudy culture clear overnight, and that multiplied as it did so. He called it a bacteriophage, an eater of bacteria, and unlike Twort he spent the rest of his life on it.

Phages, as they came to be called, mattered for two reasons. The first was practical. D'Hérelle saw at once that a virus which killed bacteria might be used to treat bacterial infections, and for twenty years before antibiotics, phage therapy was tried across the world, with results that were sometimes dramatic and more often impossible to reproduce. In the Soviet republic of Georgia, where d'Hérelle's colleague Giorgi Eliava founded an institute for the purpose, it is practised still. The second reason was scientific. Phages were cheap, fast and safe to work with. A phage made a clear spot on a lawn of bacteria within hours, and the spots could be counted. That made it possible for the first time to count viruses, to measure how many there were in a drop and how fast they multiplied, and most of the questions about what a virus ac-

tually does inside a cell were first answered, decades later, with phages.

A virus in a crystal

By the 1930s, the filterable viruses had been studied for forty years without anyone seeing one, and a reasonable biologist could still hold that they were a kind of very small bacterium. What settled the matter was a chemist's trick.

Wendell Stanley, a young biochemist at the Rockefeller Institute's laboratory in Princeton, set out to purify tobacco mosaic virus, Ivanovsky's and Beijerinck's agent, by the methods used for purifying proteins. He worked through vast quantities of infected leaves, and in 1935 announced that he had obtained from them a crystalline protein that carried the full infectivity of the virus. The crystals could be dissolved and recrystallised, and a vanishingly small quantity, rubbed on a leaf, still gave the disease. He said, in the title of his paper, that he had isolated a crystalline protein possessing the properties of tobacco mosaic virus.

Nothing that is alive forms crystals. Crystals are what salt and sugar do, and the idea that an agent which reproduced in a plant could be bottled as a crystalline powder, like a chemical, caused something close to a scandal. It also settled Beijerinck's old question: the virus was not a fluid but a

particle, regular enough to stack in a lattice. Stanley thought his crystals were pure protein and that the protein somehow copied itself. He was wrong about that. Within two years, Frederick Bawden and Norman Pirie in England showed that the crystals contained a second substance besides protein, a nucleic acid of the kind called RNA, making up a few per cent of the weight. It was in that few per cent, as the whole of molecular biology would later show, that the instructions lay. The protein was the coat. The acid was the message. Stanley shared a Nobel Prize in 1946 for the crystals all the same, and the question his work raised, whether a thing that can be crystallised is alive, has never been given an answer everyone accepts, which is why this book is called what it is.

THE NUMBER — How small is small

A human cell, the kind that lines the inside of the nose, is about 10 thousandths of a millimetre across. Written as a number, that is 0.01 mm. A single bacterium, say the common gut bacterium E. coli, is about a tenth of that: 0.001 mm, or one micrometre. The viruses that cause influenza are about a tenth of that again, 100 nanometres, where a nanometre is a millionth of a millimetre. The virus that causes polio is 30 nanometres. Tobacco mosa-

ic virus is a thin rod about 300 nanometres long and only 18 across.

Those numbers are hard to picture, so scale them up. Make the human cell the size of a football pitch, 100 metres long. The bacterium is then a family car parked on the centre spot. The influenza virus is a football beside the car. The polio virus is a golf ball. And the cell's own machinery, the parts the virus hijacks, are scattered across the pitch like gravel.

The reason the numbers mattered in 1898 is that light has a size too. A ray of visible light is about 500 nanometres across, and nothing much smaller than half that can be seen with it, however good the lenses. A bacterium, at 1,000 nanometres, is comfortably visible in a good microscope. A virus, at 30 or 100, is not: it is smaller than the light being used to look for it. The largest viruses, the family that includes smallpox, at about 200 nanometres, sit right at the limit, and an Edinburgh physician named John Buist probably glimpsed clumps of them in vaccine lymph in 1886 without knowing what he saw. Everything else was invisible by the physics of light, not by any failure of skill.

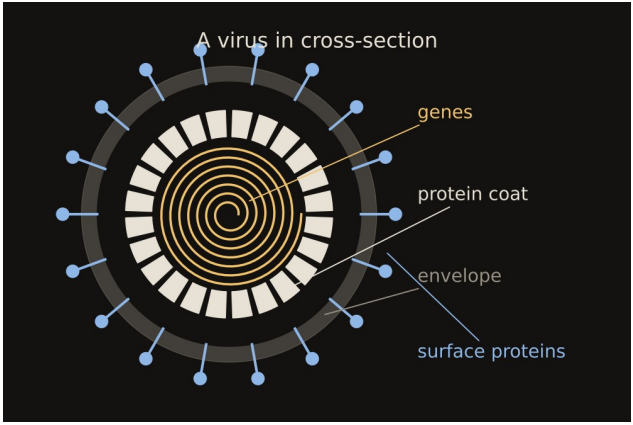


Figure F03-04. *The whole of a virus. Genes at the centre carry the instructions; the coat protects them and fits the door of a particular cell; the envelope, where there is one, is stolen from the last cell the virus left. There is no machinery because there is no need for any.*

First photographs

The instrument that could see them used not light but electrons, which behave as waves far shorter than light. The physicist Ernst Ruska built the first electron microscope in Berlin in 1931 while still a student. His brother Helmut was a physician, and it was Helmut who grasped what the machine was for. In 1938 he showed a meeting of the Medical Society of Berlin the first photographs ever taken of virus

particles, and in 1939, with Gustav Kausche and Edgar Pfankuch, published images of tobacco mosaic virus and the first direct measurements of its size. The rods were there, lying across the picture like scattered matchsticks, exactly as the crystals had implied. By the end of 1939 the Berlin group had taken nearly 2,000 photographs of viruses, among them the smallpox virus. The first picture of a phage followed in 1940, and showed it to be a particle with a head and a tail, which silenced the remaining doubters who had thought phages were an enzyme.

It is worth pausing on what these pictures changed, and what they did not. They confirmed what the filters, the crystals and the chemistry had already established: that viruses were particles, of definite size and shape, much smaller than bacteria. They made it possible to tell viruses apart by appearance, which the filters never could. But the understanding of what a virus is, the genes-in-a-coat account in the box above, came from chemistry and from the work on phages, not from the pictures. One can look at a photograph of a matchstick for a long time without learning what fire is.



Figure F03-05. *Seeing with something smaller than light. The electron microscope of the late 1930s was the first instrument whose probe was finer than a virus, and within two years of its first biological use it had photographed smallpox.*

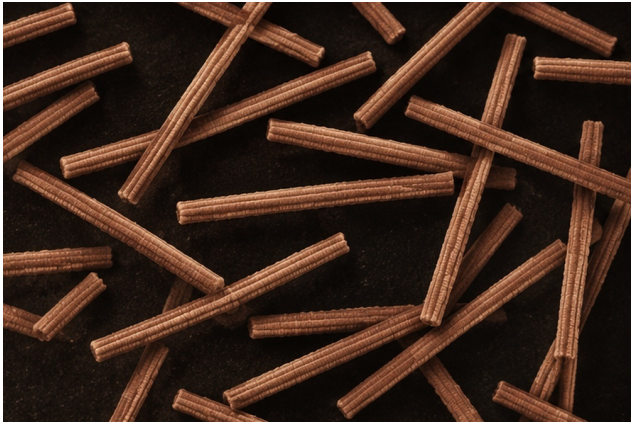


Figure F03-06. *What forty years of inference had been about. The rods are 300 nanometres long and 18 across, smaller than the wavelength of light, and were first seen in 1939 only because the microscope used electrons instead.*

Where this leaves us

In 1892 nobody knew that viruses existed. By 1940 it was known that they were particles made of a nucleic acid wrapped in protein, that they could be crystallised like a chemical and photographed with electrons, that they multiplied only inside living cells by taking over the cell's machinery, and that they caused diseases of plants, animals, people and bacteria alike. Smallpox, cowpox, rabies, yellow

fever, foot-and-mouth and influenza were all on the list, and the practice of vaccination, which had been working for a century and a half, finally had a subject.

What was still missing was almost everything that would be needed to make a vaccine against one of them on purpose. Pasteur had made his by passage through animals, because he had no other way to grow the agent, and that was still the only way. No virus had yet been grown in a dish. Nobody knew how the coat recognised a cell, or how the body recognised the coat, or what the nucleic acid did. Those were questions for the second half of the century, and the second half of this part of the book.

Before that, the reader needs to meet one of these particles at work, because the virus that is about to arrive did more harm in a year than any other event of the twentieth century, and it did it in the same decade in which Twort and d'Hérelle were watching their bacteria dissolve. The next chapter is 1918.

The Year the World Caught Cold

> *"If grippe condemns, the secondary infections execute."*
> *— Louis Cruveilhier, 1919*

A cook reports sick

On the morning of 4 March 1918, an army cook named Albert Gitchell walked into the infirmary at Camp Funston, a training camp on the Fort Riley reservation in eastern Kansas, with a sore throat, a headache and a fever. He is usually called the first case of the 1918 pandemic, and he almost certainly was not; other men in his unit reported the same symptoms within minutes of him, and the camp had 56,000 soldiers in it. By the end of the day more than a hundred men were in the infirmary. Within three weeks, eleven hundred were ill enough to be hospitalised and thirty-eight were dead.

Something had been moving through that part of Kansas already. In the sparsely populated county of Haskell, two hundred miles to the south-west, a country doctor named Loring Miner had spent January and February treating an in-

fluenza he did not recognise: violent, fast and killing people who should not have died of influenza. He was sufficiently alarmed to report it to the United States Public Health Service, whose weekly bulletin printed a short notice on 5 April 1918, by which time it no longer mattered. Draftees from Haskell County had been travelling back and forth to Camp Funston all winter.

From Funston it went where the soldiers went: to other camps across the United States, to the ports, and onto the troopships crossing the Atlantic at the rate of tens of thousands of men a month. By 11 March it was in Queens. By April it was in France, and from the trenches it spread into every army and every civilian population they touched. The spring wave was unpleasant and not especially deadly. Soldiers called it three-day fever. Most who caught it were ill for a week and went back to duty.

It acquired its name in May. Spain was neutral, and so Spanish newspapers reported freely on an epidemic that censors in the belligerent countries were suppressing as bad for morale. When the Spanish king fell ill the story ran everywhere, and the disease became the Spanish influenza, a name that was wrong, stuck, and is a useful reminder that the public name of a disease records who was willing to talk about it rather than where it began.

The autumn wave

In the last week of August 1918 something changed. The same virus, or a descendant of it, appeared almost simultaneously in Brest, in Freetown and in Boston, and this time it killed.

The second wave is the event that makes 1918 different from every other influenza year. It arrived in a city, burned through it in about six weeks, and left. The characteristic story is the same everywhere: a healthy adult feeling unwell in the morning, feverish by the afternoon, and by evening labouring to breathe, with a cough that brought up blood-stained foam, and a dusky blue colour spreading from the lips and ears across the face as the lungs stopped delivering oxygen. Some died within a day or two, their lungs filled with fluid in a way doctors had not seen before and could do nothing about. Many more lasted a week or ten days and died of a bacterial pneumonia that settled into lungs the virus had stripped bare, which in the era before antibiotics was itself close to a death sentence. That is what the epigraph to this chapter means.

London counted its dead weekly, as it had since the seventeenth century, and the record shows the shape of the thing plainly. Through the summer of 1918 the city buried

between 750 and 1,000 people a week. In the week ending 2 November it buried 4,290.

Hospitals ran out of beds, then out of staff. Public buildings became wards. Gravediggers fell ill and the dead waited. In Philadelphia, which held a war-bond parade at the end of September with two hundred thousand people in the streets and saw its hospitals overwhelmed within a fortnight, bodies were collected from houses by horse-drawn carts. Almost every account from almost every country contains the same details: the suddenness, the blue faces, the young adults, and the silence of the streets.

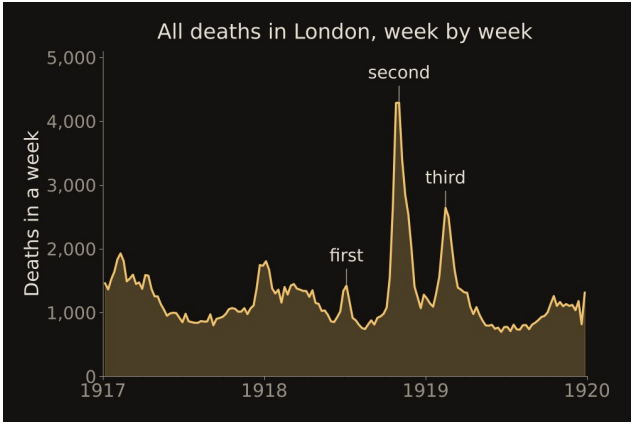


Figure F04-01. *Three waves in one city. The ordinary winter rise of 1917 is dwarfed by the autumn of 1918, which did its work in six weeks, and by the third wave of February 1919 that caught those the second had missed. Drawn from the weekly London returns.*

Ireland, October 1918

Ireland had the epidemic in three waves like everywhere else, and had it while in the middle of something else entirely.

The first wave arrived in the early summer of 1918, mild and widely ignored. The second came in October and November and was neither. The official count of deaths certified to influenza in Ireland across 1918 and 1919 was 20,057. Doctors

certifying deaths were themselves overwhelmed and many deaths went uncertified or were recorded as bronchitis or pneumonia; adding the excess pneumonia deaths of those two years gives a toll of at least 23,288. Working from the surviving statistics, the historian Ida Milne and the medical statistician Anthony Kinsella estimated that about 800,000 people on the island caught it, roughly a fifth of the population.

The geography was uneven in ways nobody could explain at the time. Kildare had the highest county death rate in 1918, at 3.95 per thousand people, which is 263 deaths in a small county. Clare had the lowest, 0.46 per thousand. Dublin city and county lost 1,767 people to influenza in 1918 and another 1,099 in 1919. The age pattern matched the rest of the world: in 1918, nearly 23 per cent of all Irish influenza deaths were of people between twenty-five and thirty-five.

Milne's book takes its title, *Stacking the Coffins*, from the mortuary at the Dublin workhouse hospital, where coffins were piled many deep because burials could not keep pace. She interviewed survivors in the 1990s and 2000s, people who had been children in 1918 and were by then in their nineties, and found that the epidemic had been remembered privately and intensely and talked about publicly almost not

at all. It had been swallowed by the war and by what came after it.

What came after it was a revolution, and the influenza got into that too. Since the arrests of May 1918, several hundred republicans had been interned without trial in English and Welsh prisons. In October, reports that more than a hundred Sinn Féin prisoners in Belfast jail had caught influenza produced questions in Parliament about their treatment, and the authorities, acutely aware of what a dead prisoner would cost them after the hunger-strike death of Thomas Ashe the previous year, improved conditions sharply. At Usk in Monmouthshire no such calculation was made. Six prisoners fell ill in late November and were left in cold cells without regular access to a doctor until 1 December. Richard Coleman, a Volunteer from Swords who had fought in 1916, died there in December 1918, days before the general election.

The election was held on 14 December. Sinn Féin won seventy-three seats, and among those elected, from a cell in Gloucester prison, was Pierce McCan of Tipperary East. The flu reached Gloucester in February. McCan died on 6 March 1919, aged thirty-six, without ever sitting in the Dáil that had assembled in Dublin six weeks earlier. His body was brought home on the mail boat along with eleven released prisoners too ill with influenza to travel under their own power, and

was buried at Dualla with a guard of Volunteers. Two of the prisoners taken in May 1918 had died in British custody, both of influenza, and the funerals of both were political events of the first order.

It is worth being careful here about cause and effect. The influenza did not make the Irish revolution; the election result was not a verdict on two deaths in English jails. But a pandemic is never only a medical event. It arrives into whatever a society is already doing, and in Ireland in the winter of 1918 it supplied funerals, grievances and absences to a movement that knew what to do with them.



Figure F04-02. *What a health service looks like when it has run out of beds. Halls, schools and barracks became wards in the autumn of 1918, and the nursing they could offer was warmth, fluids and company.*

Why the young died

Influenza kills the very young and the very old. That is the ordinary pattern, and if you plot deaths against age in a normal year you get a U: high at the left-hand end, low across the middle of life, rising again at the right.

In 1918 the middle of the U filled in. Death rates among people in their twenties and thirties rose to levels normally seen only in infants and the very elderly, and the curve ac-

quired a third peak in between. The shape is called the W, and it is the single strangest fact about the 1918 pandemic.

Two things in that chart need pointing out. The first is the size of the middle peak: for men aged twenty-five to thirty-four the death rate from influenza and pneumonia went from 79 per 100,000 in 1917 to 1,217 in 1918, a fifteenfold rise. The second is easy to miss. At the far right of the chart, the 1918 line runs *below* the 1917 line. People over seventy-five were less likely to die of influenza in the worst influenza year in history than in an ordinary one.

Nobody has a complete explanation, and it is worth saying so plainly rather than tidying it up. Two partial explanations are taken seriously. One is about the old: a person who was seventy in 1918 had been born in the 1840s and had lived through earlier influenza pandemics, and may have carried an immunity, acquired in childhood, to something sufficiently like the 1918 virus to help. The reverse argument applies to the young adults, whose first childhood encounter with influenza had been with a virus quite unlike this one. The other explanation is about the immune system itself: that the catastrophic early deaths, the ones that filled healthy lungs with fluid in two days, were caused not by the virus directly but by the body's own violent response to it, which a vigorous immune system in its twenties is better equipped to

mount. Neither account is settled, and the question matters beyond history, because the age pattern of the next pandemic will determine who has to be protected first.

KEY CONCEPT — Why influenza keeps coming back

Chapter 3 described a virus as genes in a coat. Influenza's coat is studded with two kinds of protein spike, and these are what the body learns to recognise. Make antibodies that grip those spikes and the virus is stopped at the door.

The trouble is that influenza copies its genes carelessly. Every time it multiplies it makes small mistakes, and over a season or two these accumulate until the spikes have changed shape enough that last year's antibodies no longer fit properly. This slow change is why there is a new influenza vaccine every year and why having had influenza once protects you for a few years rather than a lifetime.

There is also a fast version. Influenza's genome is not one strand but eight separate pieces, and the virus infects birds and pigs as well as people. If two different influenza viruses get into the same animal at the same time, the eight pieces can be dealt out between them like two

hands of cards shuffled together, and a virus can emerge carrying a spike from a bird virus that no human being alive has ever met. That is not a slow drift but a jump, and it is how pandemics start. Everyone is vulnerable at once, because nobody's immune memory has anything to match it against.

Chapter 11 returns to this, because the jump has happened three more times since 1918, and chapter 20 asks what would happen if it happened again.

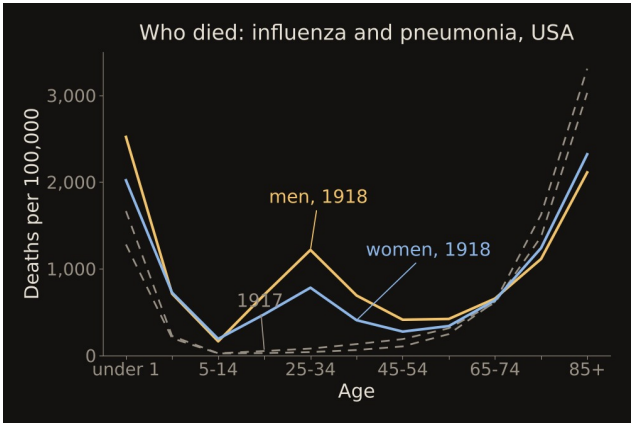


Figure F04-03. *The year the curve grew a third peak. Deaths among people in their late twenties and early thirties rose roughly twentyfold, while the rate among the very old was, remarkably, lower than in an*

ordinary year. Drawn from United States vital statistics for the death-registration states.

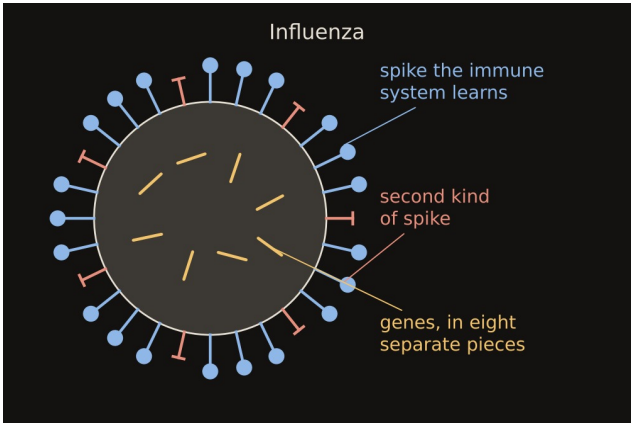


Figure F04-04. *Why influenza is the virus that keeps returning. Two kinds of surface spike are what immunity recognises, and a genome in eight separate pieces means two viruses meeting in one animal can exchange parts and produce something nobody is immune to.*

What could be done

Almost nothing, and the almost is interesting.

There was no vaccine, because nobody knew what caused influenza. The leading theory blamed a bacterium, Pfeiffer’s bacillus, found in the lungs of many victims, and several of the vaccines rushed into production in 1918 were made

against it. They did not work, for the reason chapter 3 would have predicted: the agent was a virus, it would not be seen for another twenty years, and it would not be grown in a laboratory until 1931. There were no antibiotics for the bacterial pneumonias that killed most of the dead; penicillin was a decade from discovery and two decades from use. Treatment was nursing: warmth, fluids, propping patients up so they could breathe, and someone to stay with them.

What remained was keeping people apart. Cities closed schools, theatres, cinemas and churches, banned public gatherings, staggered shop hours to thin the crowds on trams, and in some places required masks in public. The measures were adopted unevenly, abandoned early under commercial pressure, and often imposed after the local epidemic had already peaked. Even so, a century of looking at the records has produced a reasonably clear finding: the cities that acted sooner and kept the measures longer had lower peak death rates than those that did not. The gauze masks of 1918 were of little use in themselves, being loose and porous, but the acts they accompanied were not useless at all.

It is a modest conclusion and worth stating carefully, because it is the only tool the chapter contains and because the same measures, and the same arguments about them, return in chapter 16 in a form the reader will remember personally.

THE NUMBER — Fifty million, or seventeen, or a hundred

The figure usually given for the 1918 pandemic is fifty million dead. The honest version is that nobody knows, and that the published estimates differ by a factor of six.

The first serious attempt, by Edwin Jordan in 1927, arrived at about 21.6 million. In 1991 Patterson and Pyle revised it upward to 30 million, with a range from 24.7 to 39.3 million. In 2002 Johnson and Mueller proposed 50 million and said the true figure might be as high as 100 million. In 2018 Spreeuwenberg, Kroneman and Paget worked back from national mortality records, controlled for the war and for the underlying trend in deaths, and arrived at 17.4 million, concluding that totals above 25 million were not realistic.

Why the spread? Because most of the world in 1918 did not register deaths. Reliable vital statistics existed for western Europe, North America and a few other places, covering perhaps a quarter of humanity. Everywhere else, and above all in India, where the death toll was enormous and the records were not, the figures are reconstructions built on assumptions, and a small change

in an assumption applied to hundreds of millions of people moves the total by millions.

What the number does not tell you is more useful than the number. It does not tell you how many caught it, which was perhaps a third of everyone alive. It does not say where the dying was concentrated, which was overwhelmingly in the poorest places. And it is not a measure of how dangerous the virus was to an individual; on the best estimates somewhere between one and three in every hundred people who caught it died, which is a frightening figure precisely because so many people caught it.



Figure F04-05. *The entire medical armoury of 1918. Loose gauze stopped little by itself; what helped, where it was tried early and kept up, was the closing and distancing that went with it.*

Frozen ground

The virus itself was recovered seventy-nine years later, and the story of how is one of the better ones in this book.

In November 1918 the influenza reached Brevig Mission, a small Inuit village on the Seward Peninsula in Alaska, probably carried by men bringing mail and supplies overland. In five days, between 15 and 20 November, it killed 72 of the village's 80 adults. The territorial authorities hired miners from Nome to dig a mass grave in ground that had to be

thawed with steam hoses, and the 72 were buried in permafrost, where they stayed frozen.

In 1951 a Swedish graduate student at the University of Iowa named Johan Hultin obtained the village elders' permission to open the grave, hoping to recover live virus. He got tissue, but the techniques of 1951 could do nothing with it. He went on to a career as a pathologist, and he did not forget.

In March 1997 Jeffery Taubenberger and Ann Reid, working at the Armed Forces Institute of Pathology in Washington, published in *Science* a partial genetic sequence of the 1918 virus, recovered from a scrap of preserved lung tissue belonging to a soldier who had died in 1918 and whose sample had sat in a military warehouse ever since. The fragment was small, and they needed more material. Hultin read the paper, wrote to Taubenberger offering to go back to Alaska, and in August 1997, aged seventy-two and paying his own way, he returned to Brevig Mission. The elders gave permission a second time. Seven feet down he found the body of a young woman, obese in life, whose fat had insulated her lungs and kept them intact. He called her Lucy, took the lung tissue, reburied her, and sent the samples to Washington in four separate packages in case one went astray.

HOW WE KNOW — Reading a virus that no longer exists

There was no virus left alive in Lucy's lungs. What survived was fragments of its genetic material, chemically degraded and broken into short pieces mixed in with the dead woman's own.

The technique that made the work possible multiplies a chosen stretch of genetic material millions of times over, so that a quantity far too small to study becomes a quantity large enough to read. To use it you must know roughly what you are looking for, because you have to supply short matching sequences to mark the start and end of the stretch you want copied. Taubenberger's group knew the general shape of influenza genes from modern strains, and so could fish for the 1918 versions a short piece at a time and assemble the overlapping pieces into a continuous sequence.

What made the result trustworthy was that it came from three independent sources: the soldier's preserved sample, a second soldier's, and Lucy's lung. Three bodies, two continents, one sequence. Contamination of a laboratory sample is always the first objection to ancient genetic material, and three matching recoveries from

> separately stored specimens is the answer to it. By 2005 the whole genome, all eight pieces, had been published.

What the sequence showed was that the 1918 virus was an influenza of the type labelled H1N1, that it was not a mixture of human and bird strains thrown together in a pig, as had been widely assumed, but something that looked as though it had come into people from birds more or less whole, and that no single feature of it explained its lethality. The explanation, in so far as there is one, is distributed across the whole genome, and the question of why that particular combination of genes killed healthy adults is still open in 2026.



Figure F04-06. *Where the evidence was kept. Burial in permafrost preserved lung tissue well enough that, seventy-nine years later, the genetic material of the virus could be recovered from it and read.*

Where this leaves us

The 1918 pandemic killed more people in a year than the war it accompanied, and then it was largely forgotten. There are war memorials in every town in Ireland and Britain and almost no memorials to the influenza. Ida Milne's explanation, drawn from the survivors she interviewed, is that it was remembered as a private grief rather than a public one: it killed in houses rather than in battles, it killed people in their prime in a way that felt shameful and arbitrary, and the soci-

eties it hit had already used up their capacity for public mourning.

For this book the chapter marks the low point. Here is a virus doing what viruses do, in a world that had germ theory, had Pasteur's method, had vaccines against smallpox and rabies and a handful of bacterial diseases, and could do nothing whatever against this one. It could not see the virus, grow it, test for it or vaccinate against it. It could close the schools.

Everything needed to change that was about to arrive. Within fifteen years influenza had been grown in a laboratory, the first flu vaccines were in trials, and the electron microscope was photographing virus particles. Within thirty, a method for growing viruses in dishes would make vaccines possible for diseases nobody had been able to touch.

The next chapter is about the first of them, and about a disease that frightened the twentieth century out of proportion to the number it killed, because of who it took and what it left behind.

The Summer Disease

> *"Well, the people, I would say. There is no patent. Could you patent the sun?" — Jonas Salk, asked who owned the patent on the polio vaccine, interviewed by Edward R. Murrow, April 1955*

A summer in Cork

On 13 June 1956 a single case of poliomyelitis was notified to the medical authorities in Cork city. Six days later there were two more. By 3 July there were six, which was as many as the city had seen in the whole of the previous year, and the chief medical officer, Dr Saunders, understood what he was looking at.

What followed is familiar now in a way it was not then. The open-air swimming baths at Victoria Cross were closed. Parents were told to get their children to bed before dark, on the theory, which was wrong, that tiredness and chill brought the disease on. By the middle of July, St Finbarr's Hospital, which housed one of the country's three regional poliomyel-

itis centres, was admitting four new cases a day. On 15 July a five-year-old girl died there. By 9 August there had been 90 cases in the city and 34 in the county, and of the city's 90, seventy-one were children under the age of six.

Cork emptied. People from outside stopped coming in; people inside stopped going out. Cinemas and shops were deserted, visitors cancelled, and a city of eighty thousand acquired the atmosphere of a place under quarantine, which in every practical sense it was. The Gaelic Athletic Association discussed whether Cork teams should be allowed to travel to Dublin for the All-Ireland finals, and the finals were postponed. Schools, due back at the start of September, stayed shut. The Minister for Health, T. F. O'Higgins, came down to inspect St Finbarr's and St Mary's Orthopaedic Hospital and pronounced himself satisfied with the arrangements.

Primary schools reopened on Monday 8 October. By then the epidemic was over. It had produced 220 cases across the city and county and had killed five children. In the country as a whole, 1956 brought 499 notified cases and 20 deaths, a quarter of them in Cork.

The particular cruelty of the timing is that a vaccine already existed. It had been declared safe and effective in Michigan fifteen months earlier. It had not reached Cork, and would not be offered in Ireland until April 1957.



Figure F05-01. *The place parents dreaded in the summer of 1956. Cork had one of Ireland's three regional polio centres, and for three months it was the destination for a city's frightened families.*

The disease that took the strong

Polio is an old disease that became a modern epidemic, and the reason is one of the stranger facts in this book.

The poliovirus lives in the gut and spreads by the oral and faecal route: unwashed hands, shared water, contaminated food. In most people it causes nothing at all, or a few days of fever and sore throat. In perhaps one infection in two hundred it leaves the gut, reaches the nervous system and destroys the motor nerve cells of the spinal cord, and the

muscles those nerves served are paralysed, permanently. If the nerves serving the diaphragm are destroyed, the patient cannot breathe.

In a society with poor sanitation, nearly every infant meets the virus in the first months of life, while still partly protected by antibodies acquired from the mother, and nearly every infant comes through it with lifelong immunity and no paralysis. Clean water, flush lavatories and careful mothers break that pattern. Children meet the virus later, at four or seven or fifteen, with no maternal protection left, and the proportion who are paralysed rises. Polio became an epidemic disease of prosperous countries in the twentieth century because those countries had cleaned themselves up. It arrived in the summer, every summer, and nobody could say why.

What it left behind was visible, which is much of why it terrified. A town could point to the children who had had it. And for those whose breathing muscles failed there was the iron lung, a sealed steel cylinder invented in 1928 in which a patient lay with only their head protruding while a pump cycled the air pressure and made their chest rise and fall. Some spent weeks in one and walked away. A few spent the rest of their lives in one.

Franklin Roosevelt was paralysed in 1921, at thirty-nine, and concealed the extent of it for the rest of his career. In 1938

he founded the organisation that became the March of Dimes, which asked Americans to post single coins to the White House and received them by the sackload, and which went on to fund, almost single-handed, the research and the trial described below. It is the one great instance in this book of a vaccine paid for by small donations from the people who wanted it.

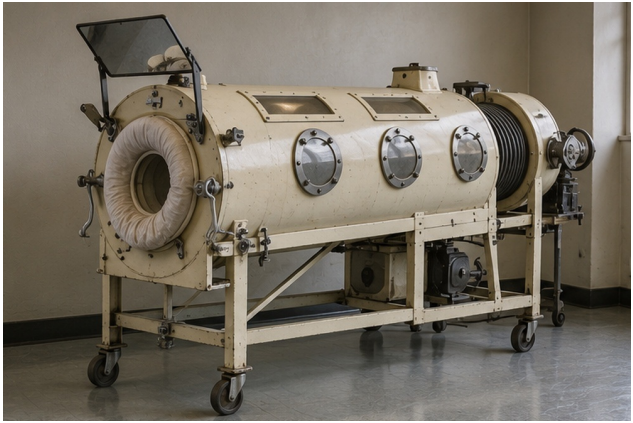


Figure F05-02. *What medicine could offer before the vaccine. The iron lung breathed for patients whose own muscles could not, and kept alive people for whom there was no treatment and no recovery of the nerves that had been destroyed.*

Growing a virus in a dish

Everything in this chapter after 1949 rests on a technical advance that sounds dull and was not.

To make a vaccine you need the virus, in quantity. Pasteur had got his rabies agent by passing it through rabbits, because there was no other way, and for polio the position in the 1940s was the same but worse: poliovirus had only ever been grown in nervous tissue, in the brains of monkeys. Vaccine made from monkey brain carried fragments of monkey nerve tissue, and injecting that into people provoked the immune system to attack the recipient's own nerves. Several early attempts at a polio vaccine, in the 1930s, did exactly that, with terrible results.

In March 1948, in a laboratory at Children's Hospital in Boston, a young researcher named Thomas Weller finished an experiment on chickenpox virus and found he had four flasks of human embryonic tissue left over. On the suggestion of the laboratory's head, John Enders, he fetched a sample of poliovirus from the freezer and put it into them, on the off chance. The flasks contained skin, muscle and connective tissue: no nerve cells at all. The virus grew.

Enders, Weller and their colleague Frederick Robbins spent the next year confirming it and working out how to keep

such cultures alive for weeks by replacing the nutrient fluid every few days. They published in 1949, and the consequences were immediate. Poliovirus could now be grown in bulk, in tissue that no vaccine maker needed to worry about. It could be identified in a dish by the damage it did to the cells, instead of by injecting a monkey and waiting. It could be counted. The three shared the Nobel Prize in 1954, by which time a vaccine built on their method was already in the arms of nearly two million children.



Figure F05-03. *The technique the rest of the century was built on. Once a virus could be grown in ordinary tissue rather than in a living animal's nervous system, vaccines could be made in quantity and tested in a dish.*

Salk’s gamble

Jonas Salk was thirty-three when the University of Pittsburgh recruited him in 1947, and he had learned his trade on influenza, where he had worked under Thomas Francis making killed-virus vaccines for the army. He proposed to do the same for polio: grow the virus, kill it with formaldehyde, and inject the dead virus, so that the body would learn the shape of the coat without ever meeting a live one.

Much of the field thought this was the wrong approach. Albert Sabin, the leading figure in American polio research, argued that only a live weakened virus, in the tradition running back to Pasteur, would produce immunity that lasted, and that a killed vaccine was at best a temporary measure. The argument between the two men was technical, personal and public, and it ran for twenty years.

Salk went ahead. In 1952 he began testing on children at a home for the disabled and at a state institution, which was then accepted practice and would not be now. In 1953 he inoculated himself, his wife and his three sons. By March 1954 he had given his vaccine to more than five thousand people and nobody had been harmed by it. That established that it was safe, in the limited sense that it did not obviously injure

anyone. It did not establish that it worked, and there was no agreement about how to find out.

The largest experiment ever made

The trial that settled it was run not by Salk but by his old teacher Thomas Francis, who agreed to take it on only if he could design it himself and if nobody, Salk included, saw the results until he had finished. It began in April 1954 and involved about 1.8 million schoolchildren across forty-four states, paid for by the March of Dimes at a cost of around seven and a half million dollars, which was raised from the public a coin at a time. The children were called Polio Pioneers.

> **HOW WE KNOW — The 1954 field trial**
>
> Francis used two designs at once, which is the key to the whole thing.
>
> In most of the participating areas, second-graders were vaccinated and first- and third-graders were watched as controls. This is an observed-control design. It is easy to run, and it has a flaw everyone knew about: the children who were vaccinated were those whose parents had consented, and parents who consent differ from parents who do not, in income, in education, in how carefully they

watch for symptoms. Any difference in outcome might be a difference between the families rather than a difference made by the vaccine.

In the remaining areas Francis insisted on something better. Every child whose parents consented was enrolled, and then each child was allocated at random to receive either the vaccine or an injection of salt water that looked identical. Neither the child, nor the parents, nor the doctor giving the injection, nor the physician who later diagnosed the illness knew which a given child had received; the vials carried code numbers that were not deciphered until the analysis was complete. That is a randomised, placebo-controlled, double-blind trial, and in 1954 it was close to unprecedented on this scale.

Each of those words does one job. *Randomised*: the two groups differ only by chance, so anything that makes families unalike is dealt out evenly between them. *Placebo-controlled*: both groups go through the same experience, so the comparison is between vaccine and no vaccine rather than between being injected and not being injected. *Double-blind*: nobody who could influence a diagnosis knows which group a sick child was in,

which matters enormously for a disease whose milder forms are a judgement call.

Francis's report ran to 164 pages. Both designs gave the same answer, and the randomised arm is the one that proved it.

On 12 April 1955, ten years to the day after Roosevelt's death, Francis presented the results to five hundred scientists and reporters in an auditorium at the University of Michigan. The vaccine, he said, was safe, effective and potent. Church bells were rung. Factories stopped. It was, by any measure, one of the few unambiguously good days of the twentieth century.

THE NUMBER — Thirty-three against a hundred and fifteen

The headlines said the vaccine was 80 to 90 per cent effective. Francis's own tables say something more modest, and the difference is worth understanding.

In the randomised placebo-controlled areas, 200,745 children received the vaccine and 33 of them developed paralytic polio. In the same areas, 201,229 children re-

ceived the placebo and 115 of them developed paralytic polio.

Work it out. Among the vaccinated, 33 in 200,745 is about 16 cases per 100,000 children. Among the placebo group, 115 in 201,229 is about 57 per 100,000. The vaccine cut the rate from 57 to 16, which is a reduction of 41 in every 57, or about 72 per cent.

So where does 90 per cent come from? From separating out the severity. The vaccine performed better against the worst forms of the disease than against paralysis in general, and better still against death: there were no deaths among the vaccinated children in the placebo-controlled areas. The honest summary is that it prevented roughly seven paralytic cases in ten and nearly all the fatal ones.

Two things the figures do not say. They do not mean that 72 per cent of vaccinated children were protected and 28 per cent were not; they describe a reduction in risk across a population, and nobody can say which individual was saved. And they say nothing about how long protection lasts, because a one-year trial cannot answer that question. It was answered later, and the answer was good.

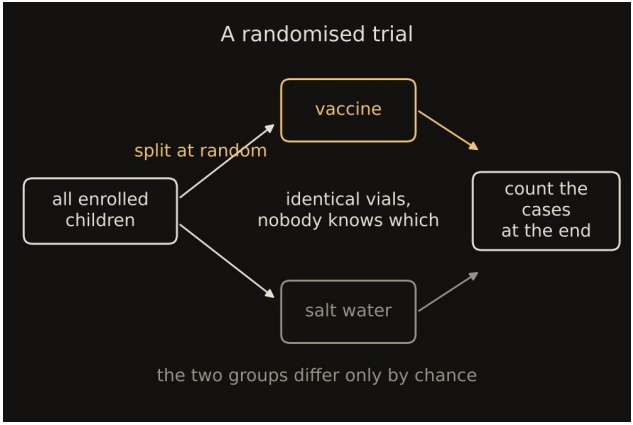


Figure F05-04. *The design that makes a result believable. Allocating at random makes the two groups alike in everything but the vaccine; making the injections and the diagnoses blind stops anyone's expectations from shaping the count.*

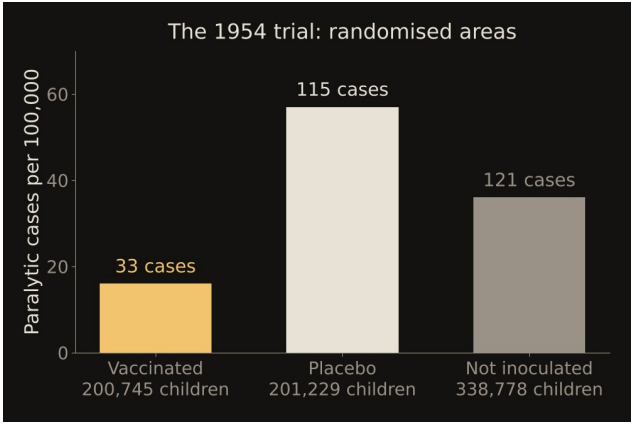


Figure F05-05. *What the trial actually showed. Children who received the vaccine developed paralytic polio at about a quarter the rate of those who received salt water, and the two groups were comparable because the allocation between them was made at random.*

Cutter

Six companies were licensed to make the vaccine, on the day the results were announced, under enormous pressure to produce it fast.

Within two weeks, children who had been vaccinated began to develop polio, and the paralysis began in the arm that had received the injection. The vaccine from one manufacturer, Cutter Laboratories of Berkeley, California, contained live

poliovirus. The formaldehyde treatment that was supposed to kill every particle had not killed every particle, and the safety test then in use was not sensitive enough to detect what was left.

About 165,000 doses of Cutter's vaccine had gone out. By the end, roughly 40,000 children had developed polio from it, 200 were left with some degree of permanent paralysis and 10 died. Some of those paralysed had never been vaccinated at all; they had caught the virus from a vaccinated child, because a child given live poliovirus sheds it, and the disease spread into families and neighbourhoods. The programme was suspended within a month.

The investigations afterwards reached a conclusion that is still worth repeating. Three of the larger manufacturers had followed Salk's inactivation protocol carefully and their vaccine was safe. Cutter had departed from it. But the deeper fault lay with the regulator: there was no adequate requirement that each lot be tested rigorously, no proper inspection of the manufacturing records, and nobody whose job it was to check. The federal government, as the historian of the episode concluded, was in the best position to prevent it and did not.

Two things followed. The safety requirements for vaccine manufacture in the United States were rewritten, and have

been tightened repeatedly since; the system of lot-by-lot release testing that chapter 18 will rely on dates from this. And a lawsuit against Cutter, which found the company liable without finding it negligent, established a principle of liability that over the following decades drove most manufacturers out of the vaccine business altogether. The world is still living with that, as chapter 20 explains when it asks who would make the vaccines for the next pandemic.

Sabin’s sugar lump

Albert Sabin had spent the Salk years developing the thing he had always argued for: a live virus, weakened by repeated passage through culture until it could no longer reach the nervous system. Live vaccine had advantages that were real. It was given by mouth, on a sugar lump or a few drops, needing no needle and no trained injector, which made mass campaigns possible in places with almost no health service. Because it multiplied in the gut, it produced immunity in the gut itself, so that a vaccinated child not only did not get polio but stopped passing the virus on. And it spread: an immunised child could pass the weakened virus to siblings and neighbours, immunising them too.

Unable to run a large trial in a United States committed to Salk’s vaccine, Sabin arranged one in the Soviet Union,

where millions of children were vaccinated in 1959. It worked. His vaccine was licensed in the United States in 1961 and became the world's standard, and it was the oral vaccine, not the injected one, that drove the global eradication campaign described in chapter 7.

KEY CONCEPT — Killed against live

These are the two ways to make a vaccine from a virus, and almost every vaccine before the 1980s was one or the other.

A **killed** vaccine, like Salk's, uses virus that has been grown and then destroyed with heat or a chemical so that it cannot multiply at all. The body sees the coat, learns the shape, and makes antibodies. Because nothing in the vial is alive, it cannot cause the disease, and it is safe for people whose immune systems are weak. The costs are that it usually needs several doses and later boosters, because the body is never shown anything that behaves like a real infection, and that it protects the person without necessarily stopping them carrying the virus.

A **live attenuated** vaccine, like Sabin's, uses virus that is alive but has been tamed by growing it in unnatural conditions until it has lost the ability to cause disease. It multiplies in the body, which means the immune system

> meets it as it would meet a real infection and often remembers it for life after one or two doses. The costs are that it must be kept cold and alive, that it cannot safely be given to people with weakened immunity, and that something alive can change. Sabin's weakened poliovirus can, rarely, mutate back towards virulence in the gut, and when it does it can paralyse. That is why wealthy countries, having eliminated wild polio, went back to the killed vaccine, and why the world's remaining polio cases in 2026 include some caused by the vaccine itself.
>
> Neither type is better. They trade different things, and which trade is right depends on the disease and on where you are standing.

What Cork was left with

Vaccination reached Ireland in April 1957, as the epidemic that had closed Cork's schools was already subsiding. Polio notifications fell away over the following years and the last case acquired in Ireland was recorded in 1984.

But the disease did not finish with the people it had taken. Nerve cells destroyed in childhood do not grow back; what recovery happens comes from the surviving nerves taking over the work of the lost ones, and those overworked cells

can fail decades later. The result, muscle weakness, pain and fatigue appearing thirty or forty years after an apparent recovery, is called post-polio syndrome. In 2026, seventy years after the Cork epidemic, Polio Survivors Ireland put the number of Irish people living with it at more than four thousand. Some of them were among the 220 children of 1956.

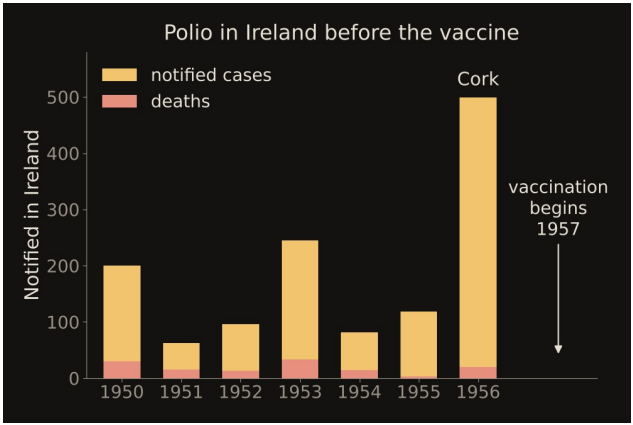


Figure F05-06. *The last epidemic. Ireland's polio notifications had fluctuated for fifteen years before 1956 produced 499 cases and 20 deaths, a quarter of them in Cork, in the last summer before a vaccine was available. Drawn from the Report of the Department of Health, 1956-57.*

Where this leaves us

Three things in this chapter outlast polio.

The first is tissue culture. Once a virus could be grown outside a living animal, the limiting step in vaccine making stopped being biology and started being engineering, and chapter 6 is the story of what a single determined laboratory could do with that freedom in twenty years.

The second is the trial. Jenner persuaded people with twenty-three cases and Pasteur with a tally of the treated and the dead, and both were right; but neither could have proved a vaccine that was 70 per cent effective rather than 100 per cent, and almost every vaccine since has been in that range. Randomisation, placebo and blinding are how a modest, real benefit is distinguished from no benefit at all, and after 1954 no vaccine was accepted without them. Every efficacy figure in the rest of this book, including the ones in chapter 17 that the reader lived through, comes from a trial built on the design Francis used on Cork-sized populations of American schoolchildren.

The third is Cutter, and it is the uncomfortable one. The polio vaccine is among the most successful things medicine has ever done, and the first mass campaign gave ten children a fatal disease it was meant to prevent. Both of those are

true, and a book that tells you only the first is not telling you about vaccines; it is advertising them. Chapter 11 returns to this at length, because the question of how a rare harm is found, acknowledged and paid for is one of the hardest in the subject, and because the answers given in the 1950s shaped the ones given in 2021.

The next chapter stays with the people who built the childhood vaccination schedule, and with a father who took a throat swab from his sleeping daughter.

A Daughter's Sore Throat

> *"If I had to name a person who has done more for the benefit of human health, with less recognition than anyone else, it would be Maurice Hilleman." — Robert Gallo, 2005*

A father with a swab

At one o'clock in the morning on 21 March 1963, a five-year-old girl in a house outside Philadelphia woke her father to tell him her throat was sore. He felt the side of her jaw, found it swollen, and recognised mumps.

Most fathers would have given her a glass of water and gone back to bed. This one drove to his laboratory, fetched cotton swabs and a vial of nutrient broth, came home, woke her again, swabbed the back of her throat, drove back to the laboratory, and put the swab in the freezer. His name was Maurice Hilleman and he ran virus research at Merck. The girl was Jeryl Lynn. The virus from her throat was grown, weakened by passage through hen's eggs and cell culture,

tested over the following four years, and licensed as a vaccine in 1967. It is still the mumps virus in the vaccine given to children today, and it still carries her name.

Hilleman is the most consequential figure in this book whom almost nobody has heard of. Over a career running from the 1940s to the 1990s he had a hand in some forty vaccines, including those against measles, mumps, rubella, chickenpox, hepatitis A, hepatitis B, meningitis and pneumonia. He was abrasive, swore constantly, kept a row of shrunken-head models on his desk representing colleagues who had crossed him, and was interested in results rather than credit. The Jeryl Lynn strain was his one sentimental gesture.

Mumps deserves a word before it disappears from the story. It is a swelling of the salivary glands that most children get through in a week, and for most of them that is all it is. For a minority it is not: mumps was, before the vaccine, one of the commonest causes of acquired deafness in childhood, it inflames the covering of the brain in a few cases in a hundred, and in boys past puberty it inflames the testicles, occasionally with lasting effect. It also used to close schools and barracks: during the First World War it put more American soldiers in hospital than any disease except influenza and gonorrhoea. A week in bed for nearly everyone and permanent damage for a few is the characteristic profile of the diseases

in this chapter, and it is precisely the profile that makes a vaccine look unnecessary until you count.

This chapter is about the two decades in which the childhood vaccination schedule was assembled, mostly by a handful of people, using the technique chapter 5 described. Once viruses could be grown in cell culture, the question stopped being whether a vaccine was possible and became which disease to do next.

Measles, the one everyone caught

Before the vaccine, essentially every child caught measles. It was so universal that it was not thought of as a danger at all, which is why the figures come as a shock.

In England and Wales, where measles had been notifiable since 1940, notifications ran between 160,000 and 800,000 a year, rising and falling in a two-year rhythm as each crop of susceptible infants accumulated and was burned through. In the peak year of 1961 there were 763,531 notified cases. Around a hundred children died of measles each year, and many more were left deaf or brain-damaged by it, and because the notifications counted only the cases a doctor saw, the true number of infections was higher still.

What measles does to a child is worth setting out, because the generation that caught it has mostly forgotten and the generation that did not has never known. It is the most contagious virus known: one case in a susceptible household infects almost everyone in it. It begins as a heavy cold with sore, streaming, light-hating eyes, builds to a high fever and a rash that spreads from the hairline downwards, and leaves the child flattened for a fortnight. Of those who die, most die of the pneumonia that follows. About one child in a thousand develops inflammation of the brain, which can leave permanent damage, and a far rarer complication, appearing years later, destroys the brain entirely and is always fatal.

And there is something else, which was not understood until long after the vaccine was in use. Doctors had noticed for a century that children seemed unusually prone to other infections in the months and years after measles, and had assumed this was general debility. In 2019 a team working with blood samples from 77 unvaccinated Dutch children, taken before and two months after they caught measles in an outbreak in 2013, showed what was happening. The virus destroys the cells that hold the immune system's memory of past infections. The children lost between 11 and 73 per cent of their existing antibodies to other viruses and bacteria: protection built up over a lifetime of ordinary exposures, erased.

Children who had been given the vaccine instead lost nothing. Measles does not merely make a child ill for a fortnight. It takes away part of what the child had already learned, and the learning has to be done again.

John Enders, whose tissue-culture work had made the polio vaccine possible, turned to measles next, isolated the virus in 1954 from a boy named David Edmonston, and by 1961 had a working vaccine. It caused a high fever and a rash in many children, and had to be given alongside an injection of antibodies to damp this down. Hilleman took the Enders strain, passed it further to weaken it, and in 1968 produced the version that needed no antibody injection and is used still.

The effect is one of the clearest things in this book. Measles vaccination began in England and Wales in 1968. The two-year epidemic cycle, which had run without interruption for as long as anyone had been counting, stopped.

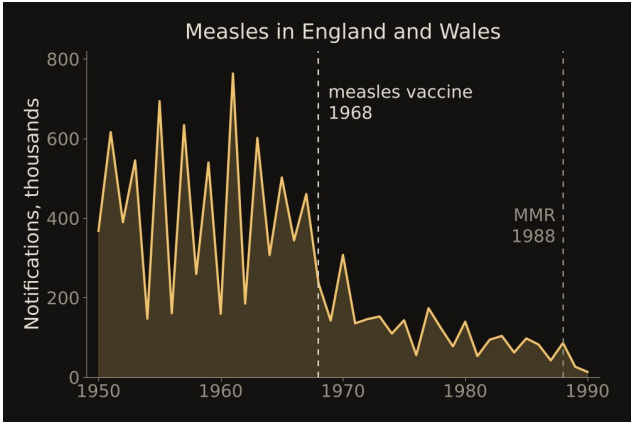


Figure F06-01. *What a vaccine does to a disease every child used to catch. The sawtooth before 1968 is the natural rhythm of measles exhausting and renewing its supply of susceptible children; vaccination flattened it, and the arrival of MMR in 1988 pushed notifications to a few thousand. Drawn from UK Health Security Agency notification records.*

The rubella spring

Rubella, or German measles, is a mild illness. A child who catches it has a few days of rash and slight fever and is generally thought to be lucky, because it is over with.

In 1941 an Australian ophthalmologist, Norman Gregg, noticed that he was seeing an unusual number of babies with

cataracts, and that their mothers had something in common: they had all had rubella early in pregnancy. The mild childhood illness, caught by a woman in the first weeks of carrying a child, attacks the developing eye, ear, heart and brain. The result is called congenital rubella syndrome.

In 1964 and 1965 a rubella epidemic ran through the United States, and because rubella is mild in the people who catch it, it ran almost unopposed. There were an estimated 12.5 million cases. Among the consequences were 11,250 pregnancies lost, 2,100 newborn deaths, and 20,000 infants born with congenital rubella syndrome: about 11,600 of them deaf, 3,580 blind, and 1,800 with lasting intellectual disability. Schools for deaf children in America were still absorbing that cohort fifteen years later.

Live rubella vaccines were licensed in the United States in 1969, and the one that proved best, developed by Stanley Plotkin in Philadelphia, followed shortly after. The point of vaccinating children against rubella was never really to protect the children. It was to stop them passing it to pregnant women, and it is the first vaccine in this book whose purpose was to protect somebody other than the person receiving it. Chapter 7 gives that idea a name.

Three in one

In 1971 Hilleman combined his measles, mumps and rubella vaccines into a single injection. It is the same combination, with the same Jeryl Lynn mumps strain, that chapter 10 will find at the centre of the worst vaccine controversy of the twentieth century.

The case for combining was practical and powerful. Three diseases, one needle, one visit, one set of records. Every separate appointment is a chance for a child to be missed, and uptake of a combined vaccine is reliably higher than the uptake of its parts given separately. The case against was that a combination is harder to test, harder to make, and that if anything goes wrong with any component the whole schedule is affected. Both of those turned out to be true as well.



Figure F06-02. *How several of these vaccines were grown. Passing a virus repeatedly through hen's eggs weakened it for human beings, and the method that tamed the Jeryl Lynn mumps strain is still used to make influenza vaccine today.*

A virus found in an Australian's blood

The last of the great mid-century vaccines came from a piece of research that had nothing to do with vaccines.

Baruch Blumberg was studying why different populations differ in their susceptibility to disease, which meant collecting blood from all over the world and looking for variations in its proteins. In 1965 he reported a protein he had found in the blood of an Australian Aboriginal man that reacted with

antibodies in the blood of an American haemophiliac who had received many transfusions. He called it the Australia antigen and had no idea what it was. Over the next two years he and his colleagues worked out that it was not an inherited variation at all: it was part of the virus that causes hepatitis B, floating free in the blood of people who carried the infection.

That mattered immediately, because it gave the world a blood test. From 1971 donated blood could be screened for the Australia antigen, and a major route of transmission closed. It mattered more in the longer run, because Blumberg realised that the antigen itself, purified from carriers' blood, might serve as a vaccine: here was the coat of the virus, in quantity, with no virus attached.

> **HOW WE KNOW — The hepatitis B vaccine trial, 1980**
>
> Making a vaccine out of protein harvested from infected human plasma was an alarming proposition, and proving that it was both safe and effective required a trial in people at high risk of catching the disease soon.
>
> Wolf Szmuness, an epidemiologist at the New York Blood Center who had survived the war in Poland and the Soviet Union and rebuilt his career from scratch in

America, designed it. He enrolled 1,083 homosexual men in New York, a group then known to have a very high rate of hepatitis B infection, and allocated them at random to receive either the vaccine or a placebo, with neither the men nor the investigators knowing which. It is the same design as the Francis trial in chapter 5, on a far smaller population, made possible by the fact that the disease was common in that group.

The result, published in 1980, was that the vaccine prevented about 92 per cent of hepatitis B infections, and that antibody levels rose in nearly everyone who received it.

Two things are worth noting about this trial beyond its result. It was the first demonstration that a vaccine could prevent a chronic blood-borne infection, and so the first step towards preventing the liver cancer that hepatitis B causes decades later, which chapter 9 takes up. And it was conducted in a community that, within two years, would be at the centre of a far worse epidemic: many of the men in Szmuness's trial would be dead of AIDS before the decade ended. Chapter 8 begins in those same clinics.

The other half of the schedule

A reader looking at a child's vaccination record will notice that it is not all viruses. Alongside measles, mumps, rubella and polio sit diphtheria, tetanus, whooping cough, and in many countries tuberculosis and meningitis. Those are bacteria, and this book is about viruses. It is worth a few hundred words to say what the difference is and why this volume leaves the other half to another.

> **KEY CONCEPT — Why antibiotics do nothing to a virus**
>
> A bacterium is a cell. It is a complete living thing in a bag, with its own genes, its own machinery for reading them, its own way of making energy and building proteins, and its own method of dividing in two. It can live in a puddle, on a doorknob or in a wound, needing nothing from anybody.
>
> A virus, as chapter 3 explained, is not a cell. It has genes and a coat and nothing else, and it can do nothing at all until it gets inside one of our cells and borrows the machinery there.
>
> That difference is why antibiotics work on one and not on the other. An antibiotic is a chemical that jams some

piece of bacterial machinery — the wall-building, the protein-making, the gene-copying — choosing a piece different enough from ours that we are unharmed. Penicillin breaks the construction of the bacterial cell wall; our cells have no wall, so penicillin does nothing to us.

A virus offers no such target. It has no wall to break, no protein factory to jam, no metabolism to poison. The machinery it is using is ours, and a drug that stopped it would stop us too. This is why there are thousands of antibiotics and only a few dozen antiviral drugs, why most antivirals had to wait for the detailed knowledge of particular viruses that chapter 8 describes, and why vaccination has mattered so much more against viruses than against bacteria. Against a bacterium there is usually a cure. Against a virus, prevention is often all there is.

Taking an antibiotic for a cold or for influenza, therefore, does nothing for the illness, and it does do something for the bacteria living in and around the person taking it: it selects the ones that can survive it. Chapter 23 returns to where that has led.

The bacterial vaccines came first, in fact. Emil von Behring showed in the 1890s that serum from an animal that had survived diphtheria could save a child dying of it, work for

which he received the first Nobel Prize in Medicine in 1901; a true diphtheria vaccine followed in the 1920s, along with vaccines against tetanus and whooping cough. And in 1921, at the Pasteur Institute, Albert Calmette and Camille Guérin produced a living, weakened strain of the cattle tuberculosis bacterium, after thirteen years of passing it through culture, which became the vaccine known by their initials: BCG.

Britain and Ireland were slow to accept it. The medical establishment in both countries doubted its effectiveness, and in Ireland, which had the highest tuberculosis mortality among young adults of any country in western Europe or North America, that scepticism cost lives for twenty years.

The person who broke it was Dorothy Stopford Price. Born in Dublin in 1890, she had begun medicine at Trinity in 1916, worked as a dispensary doctor in Kilbrittain in west Cork during the War of Independence, treating the wounded on the republican side, and then joined St Ultan's, an infants' hospital in Dublin founded by Kathleen Lynn and Madeleine ffrench-Mullen for the children of the slums. There she met tuberculosis in its most ordinary form: children in overcrowded rooms, infected by adults, dying slowly.

She taught herself German in order to read the Continental literature, travelled to Vienna and to Scandinavia to see the tuberculin test and BCG in use, and in January 1937 vaccin-

ated two infants at St Ultan's with BCG obtained from Sweden. It was the first use of the vaccine in Ireland or Britain. She was arguing, against her own profession, both that Irish children were not already protected by early exposure, as was widely assumed, and that a vaccine the British establishment distrusted was worth using.

She was vindicated politically rather than medically. In 1949 the Minister for Health, Noël Browne, himself a survivor of a family destroyed by tuberculosis, made her chair of the National BCG Committee and put the national centre at St Ultan's. Mass BCG vaccination, mass radiography and a building programme of sanatoria followed, and Irish tuberculosis mortality, which had been a national disgrace, fell steeply through the 1950s. Antibiotics arrived at the same time and take much of the credit, which is a fair division, though it should be noted that the vaccination campaign had begun before streptomycin was generally available. Stopford Price had a stroke in the late 1940s and died in 1954, as the campaign she had argued for across two decades was taking effect.

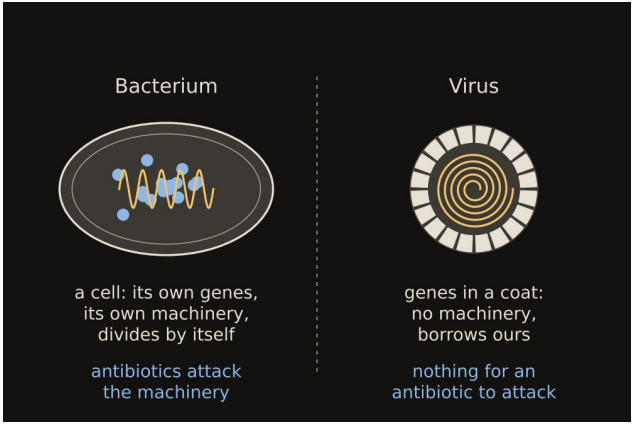


Figure F06-03. *The reason one half of the schedule is in this book and the other is not. A bacterium is a living cell with machinery a drug can jam; a virus carries none of its own, which is what makes vaccines so much more important against it.*



Figure F06-04. *Where the Irish argument about BCG was settled. St Ultan's, founded for the infants of the Dublin slums, was the first hospital in Ireland or Britain to use the tuberculosis vaccine, in January 1937, and became the national centre for it twelve years later.*

Building a schedule

Having a vaccine and having a vaccinated population are different things, and the gap between them is organisation.

The schedule that emerged in the rich countries during the 1960s and 1970s rests on a few simple ideas. Give the vaccines when children are most vulnerable and most reachable, which in practice means in the first two years of life, when parents are already bringing infants to be weighed and

checked. Combine what can be combined, to reduce the number of visits. Keep a record, so that a missed dose can be found and given late. Make it free, because any charge falls hardest on the families whose children are at greatest risk. And tie it loosely to school entry, which is the one moment when the state sees every child.

None of this is medicine. It is clerical work, and it is the reason a vaccine invented in a laboratory in Pennsylvania reaches a child in Mayo.

It also carries a cost that only became visible later. A schedule is a routine, and a routine is something people stop thinking about. Through the 1960s and 1970s parents brought their children because they remembered the diseases: they had seen a cousin in callipers or a classmate who had gone deaf, and nobody needed to be persuaded. The schedule worked so well that within a generation no parent had seen any of it. What remained was a list of injections given to a healthy infant for reasons the parents had to take on trust, against diseases that had become words rather than memories. That is not an argument against the schedule. It is the explanation for most of Part II.

THE NUMBER — Before and after

Take measles in England and Wales, where the counting was consistent across the change.

In the ten years before vaccination, 1958 to 1967, notified cases averaged about 412,000 a year. In the ten years after the vaccine was introduced in 1968, 1969 to 1978, they averaged about 149,000. That is a fall of roughly two-thirds, achieved with vaccine coverage that never exceeded about half of children in that period.

Deaths fell further. Measles was killing about a hundred children a year in England and Wales in the early 1960s. By the late 1970s it was killing about twenty.

Three cautions about those numbers. First, notification is incomplete and the proportion of cases reported may itself have changed. Second, nothing in a before-and-after comparison proves cause; the fall is attributed to the vaccine because it begins in the year the vaccine does and because the two-year cycle breaks at the same moment, not because the arithmetic alone could show it. Third, and most important for chapter 22: two-thirds was not enough. Measles kept circulating in England and Wales for another twenty years, because the proportion of children vaccinated stayed too low to stop it. The

> threshold required is much higher than people expect, and the next chapter explains why.



Figure F06-05. *The room where the schedule actually happens. Routine vaccination worked because it was attached to visits parents were already making, and because somebody kept a list of who had not come.*



Figure F06-06. *The least glamorous object in this book and one of the most useful. A record is what turns a vaccine into a programme, by making it possible to know who has been missed.*

Where this leaves us

By 1980 a child in a wealthy country was routinely protected against measles, mumps, rubella, polio, diphtheria, tetanus and whooping cough, and in many places tuberculosis, and a vaccine against hepatitis B had just been proved. Most of this had been done in about twenty years, by a small number of laboratories, using a technique that did not exist in 1948.

It was also the high point of a particular kind of confidence. The diseases being prevented were the familiar ones, the

ones every adult remembered having; the vaccines were given by family doctors and nurses to children whose parents had seen what measles and polio did; and the question of whether to vaccinate was barely a question at all.

Everything after this chapter is harder. The viruses in Part II are ones that defeated this method: a virus that hides inside the genes of the cells it infects, viruses that cause cancer decades after infection, viruses that jump out of forests. And the confidence does not survive either, for reasons that are partly the fault of the people who had it.

The next chapter finishes Part I with the one unambiguous victory: the only time in history that a human disease has been deliberately driven out of existence.

The Last Case

> *"Declares solemnly that the world and all its peoples have won freedom from smallpox." — Resolution WHA33.3, Thirty-third World Health Assembly, 8 May 1980*

A disease with nowhere to hide

Chapter 1 left smallpox spreading arm to arm across Europe while nobody knew what it was. This chapter is about the last of it, and it begins with the question of why this disease and no other.

Humanity has tried to eradicate several diseases and has succeeded, so far, with one. The reasons are properties of the virus rather than of the people fighting it, and they are worth listing because they explain both the victory here and the difficulty everywhere else in this book.

Smallpox infected human beings and nothing else. There was no animal that harboured it, which meant that if it could

be driven out of people it had nowhere left to go. It produced no chronic carriers: a person either had it and recovered with lifelong immunity, or died; nobody walked around shedding virus for years. It was unmistakable. The rash was so characteristic that a villager who had seen one case could recognise another, which made it possible to find cases without a laboratory. It spread by close contact rather than drifting across a room, so its chains of transmission were short and traceable. And the vaccine was unusually good: one dose, lasting protection, and after freeze-drying was perfected in the 1950s, a preparation that survived tropical heat without refrigeration.

It is worth remembering what was being eradicated. Smallpox began like influenza, with fever, headache and backache so severe that patients remembered it decades later. Then the rash came, starting on the face and forearms and spreading, and unlike the rash of chickenpox its spots all developed together, rising over a week into hard, deep, round pustules packed closest where the skin was thinnest. In the severe form about three people in ten died, usually in the second week. Those who lived were left pitted, often deeply, and a proportion were blinded by scarring of the cornea. In the eighteenth century it was among the commonest causes of blindness in Europe. The milder form, variola minor, which

is the one that survived longest in the Horn of Africa, killed about one person in a hundred, and it was that comparative mildness which let it circulate unnoticed when the deadlier form was already gone.

Take away any one of those and the story ends differently. Chapter 20 asks what it would take to do the same to influenza, which has an animal reservoir the size of the world's bird population, and the answer is that it cannot be done.

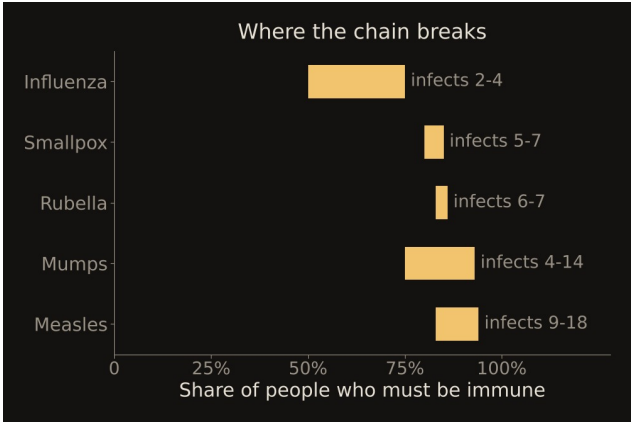


Figure F07-01. *Why one disease could be eradicated and another cannot be held down. The more people each case infects, the higher the share of the population that must be immune before the chain breaks, and measles sets a threshold that few countries have ever sustained.*

Henderson’s unpopular programme

The World Health Assembly had resolved to eradicate smallpox in 1959 and then funded the effort at about a hundred thousand dollars a year, which is a statement of intent rather than a programme. For most of the next decade little happened.

In 1966 the Assembly voted, narrowly and against the private scepticism of much of the organisation, to try properly. The Intensified Smallpox Eradication Programme began on 1 January 1967, with a target of ten years. Donald Henderson, an American epidemiologist from the Centers for Disease Control, was put in charge of it a few months before it started, somewhat against his will.

The scale of what they were taking on is easy to understate. In 1967 smallpox was endemic in more than thirty countries containing well over a billion people. The WHO’s own estimate was that there were ten to fifteen million cases a year and around two million deaths. Reporting was so poor as to be close to meaningless: in some countries fewer than one case in a hundred was being notified, which meant that nobody knew where the disease actually was.

Almost everyone expected it to fail. Malaria eradication, attempted with great confidence in the 1950s, had just col-

lapsed, and the memory of that failure hung over everything. The programme's annual budget never exceeded a small fraction of what a single rich country spends on health in a day.

Rings, not blankets

The strategy that everyone assumes was used, and the one Henderson spent the rest of his life correcting people about, was mass vaccination: vaccinate as near to everybody as possible, and the disease runs out of people to infect.

That is how smallpox had been driven out of Europe and North America, and it works, but it is enormously expensive, and in countries with poor roads, scattered populations and no reliable registers it never quite finishes. The last ten per cent of a population is always the hardest to reach, and smallpox lives in exactly that ten per cent.

In late 1966 an American doctor named William Foege was working as a missionary physician in eastern Nigeria when smallpox appeared in the area and he found himself short of vaccine. Rather than spread it thinly, he did something counter-intuitive. Using a radio network of mission stations, he mapped where the cases actually were, and then concentrated every dose he had on the people around each case: the household, the neighbours, the people who had visited, and

then their contacts in turn. He was not trying to make the population immune. He was trying to build a wall around each fire.

It worked far faster than the vaccine supply should have allowed, and when the Nigerian civil war interrupted the programme the following year, the outbreaks that had been contained this way stayed contained. The approach became the centre of the global strategy: find every case, isolate it, vaccinate the ring of people around it, and move on. It is called surveillance and containment, and it is the reason the ten-year deadline was nearly met.

Henderson was careful, in later years, about the way this story is told. Surveillance had been written into the programme's field manual from the beginning, in 1966, and mass vaccination was never the whole of the plan; the shift between 1967 and 1969 was one of emphasis rather than a conversion. What is not in dispute is that the emphasis shifted, that Foege's outbreak work in West Africa is what shifted it, and that it transformed a programme which was running out of money into one that was running ahead of schedule.

KEY CONCEPT — Herd immunity

A disease needs a supply of people to infect. If every person an infected person meets is already immune, the chain stops, and it stops whether or not the immune people notice anything.

This is why a disease dies out long before everybody is immune. Picture a hundred people in a village and one case of a disease that, in a village where nobody is immune, passes to five others. If nobody is immune, those five infect twenty-five, and the outbreak grows. Now suppose eighty of the hundred are immune. Of the five people the first case would have infected, four are immune and nothing happens to them; one is not, and catches it. One case has produced one case, and the outbreak is holding steady rather than growing. Make it eighty-five immune, and one case produces less than one, and the chain dies out on its own.

The people protected by this are not the immune ones. They are the twenty, or the fifteen, who are not immune and who never meet the disease because it ran out of road before it reached them: the newborn too young to be vaccinated, the child on chemotherapy whose immune system cannot be trusted with a live vaccine, the

small number in whom the vaccine did not take. Vaccination is the only medical intervention that routinely protects people other than the one receiving it, and herd immunity is how.

Two cautions. The protection is not spread evenly: it depends on who is immune and who mixes with whom, so a town with high overall coverage can still contain a school or a street where almost nobody is vaccinated, and an outbreak will find it. And it is not permanent. Each year brings a new crop of babies who are not immune to anything, which is why a vaccination programme is not a campaign that can be completed but a thing that must be done forever.

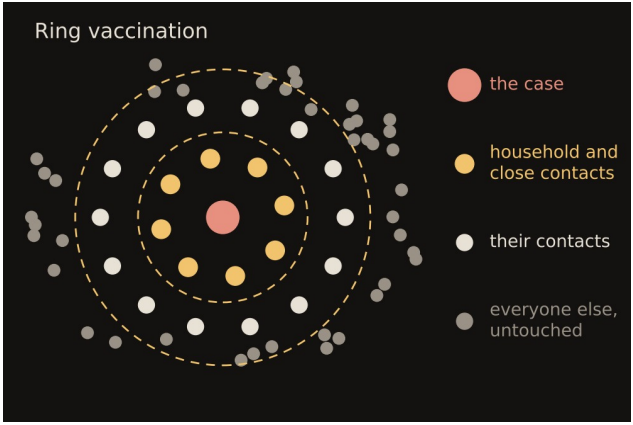


Figure F07-02. *Why a programme short of vaccine beat one that was not. Instead of trying to immunise a whole population, surveillance and containment found each case and put a wall of immunity immediately around it, so the chain of transmission ended there.*

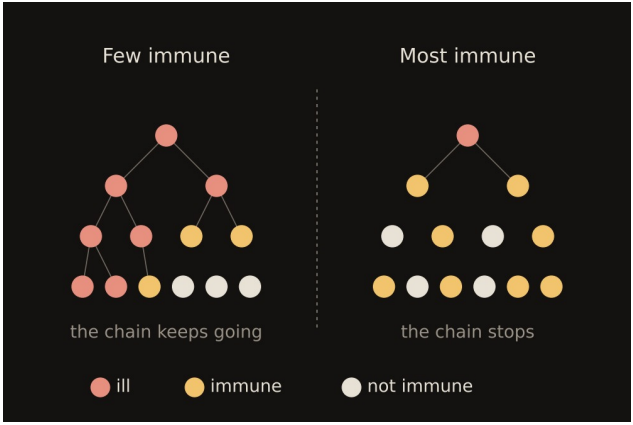


Figure F07-03. *The same disease, the same first case, two populations. Where most people are already immune, the chain of transmission runs into immune people and stops, and the few who are not immune are never reached.*

A needle with two prongs

The other change was a piece of ironmongery.

Vaccination against smallpox had always been awkward. The vaccine had to be scratched into the skin, and the traditional methods, a lancet or a rotary instrument, wasted most of the dose, were hard to teach and were easy to do badly. Jet injectors, introduced in the 1960s, were fast but heavy, needed maintenance and broke down in the field.

In 1965 a scientist at Wyeth named Benjamin Rubin took an ordinary sewing machine needle, ground the eye down to leave two prongs, and found that the gap between them held exactly the right amount of vaccine by capillary action. A vaccinator dipped the needle in the vial, lifted it out with a droplet suspended between the prongs, and jabbed the skin a dozen times in a few seconds. It used a quarter of the vaccine the old methods needed. It cost almost nothing. It could be boiled and reused hundreds of times. And it could be taught in a quarter of an hour to someone with no medical training, which in villages with no clinic was the whole point.

The programme's internal honour, given to those who had worked on it, was a lapel pin in the shape of the bifurcated needle bent into a circle. It is a fair emblem: the eradication of smallpox was accomplished by epidemiology, political persistence and a modified sewing needle.

> **THE NUMBER — How many must be immune**
>
> Epidemiologists describe how contagious a disease is with a single figure: the average number of people one infected person passes it to, in a population where nobody is immune and nobody is taking precautions. For smallpox the usual estimate is between five and seven.

From that figure you can work out how much of a population has to be immune to stop the disease spreading, and the arithmetic is the village example above. If one case would infect six people, you need enough immunity that on average fewer than one of those six is susceptible. One in six is about 17 per cent, so you need roughly 83 per cent of people immune, and a little more to be safe. Published estimates for smallpox put the threshold at 80 to 85 per cent.

Now do it for measles, which is the most contagious virus known and passes to somewhere between twelve and eighteen people. One in fifteen is under 7 per cent, so more than 93 per cent of the population must be immune. Allow for a vaccine that is very good but not perfect, and for the fact that immunity is never spread evenly, and you arrive at the figure chapter 22 is built around: 95 per cent.

This is the single most useful number in the book, and the most misunderstood. It explains why smallpox could be eradicated with coverage that would be nowhere near sufficient for measles. It explains why a measles programme that reaches nine children in ten is not nearly as good as it sounds. And it explains why a few per cent of

a population declining a vaccine can matter enormously for one disease and hardly at all for another.

What the number does not do is tell you about any individual. It is a property of a population, and it assumes people mix at random, which they do not.



Figure F07-04. *The instrument that made the arithmetic work. Holding one dose between two prongs by capillary action, it cut the vaccine needed per person to a quarter and could be taught to anyone in fifteen minutes.*

Merca, October 1977

By the mid-1970s the disease was cornered. South America had been clear since 1971, Asia since 1975. The last naturally occurring case of the severe form of smallpox was a three-year-old girl in Bangladesh, Rahima Banu, in October 1975; she was found because an eight-year-old neighbour reported her to the eradication team and collected the reward.

What remained was the milder form, variola minor, in the Horn of Africa, in a region with a war going on in it.

On 12 October 1977, a twenty-three-year-old hospital cook in the Somali port town of Merca was asked to ride along in a vehicle taking two children with smallpox to an isolation camp. His name was Ali Maow Maalin. He also worked part-time as a vaccinator for the eradication programme, which required him to be vaccinated, and he was not; he admitted many years later that he had avoided it because he disliked the needle.

He was with the children for perhaps ten minutes. On 22 October he developed a fever. He was diagnosed first with malaria and then with chickenpox, and continued to move around the hospital for several days before an eradication worker recognised the rash on 30 October. He was isolated, he recovered, and as a result of the investigation that fol-

lowed, more than fifty thousand people in and around Merca were vaccinated within two weeks and every one of his contacts was traced.

None of them developed smallpox. That is the end of it: Ali Maow Maalin was the last person in the history of the world to catch smallpox in the ordinary way, and the disease that had killed perhaps three hundred million people in the twentieth century alone went extinct in the body of a hospital cook who had skipped his own vaccination.

He spent the rest of his working life as a vaccinator, campaigning against polio in Somalia and telling people in villages that he had had smallpox and that they should not take the risk he had taken. He died of malaria in 2013, in Merca, while working on a polio campaign.



Figure F07-05. *The last front. Eradication ended not in laboratories but in places like this, where teams found each case, traced its contacts and vaccinated outward until there was nobody left to infect.*

Birmingham, 1978

There is a coda, and it is worth telling because it is the part most accounts leave out.

In August 1978 a medical photographer named Janet Parker, who worked at the University of Birmingham Medical School, became ill. Her darkroom was on the floor above a laboratory where smallpox virus was still being studied. She developed a rash on 15 August, was diagnosed with variola major on 20 August, and died on 11 September. Her mother,

who nursed her, caught smallpox too and survived. Her father died of a heart attack during her illness. Henry Bedson, the professor who ran the laboratory, took his own life on 6 September.

The subsequent inquiry could not establish with certainty how the virus had travelled, and the question is still argued about. What was not in doubt was that laboratories holding the virus were a danger in proportion to their number, and the programme's response was to reduce that number. Stocks were consolidated, destroyed or transferred, until by the early 1980s the declared stocks of smallpox virus in the world were held in just two places, one in the United States and one in Russia, where they remain. Whether they should be destroyed has been debated at the World Health Assembly for forty years without resolution.

Janet Parker was the last person to die of smallpox. She died eleven months after the disease had stopped circulating anywhere on earth.

> **HOW WE KNOW — Proving an absence**
>
> Declaring a disease eradicated is logically awkward. You cannot observe the absence of something everywhere;

you can only fail to find it, and failing to find a disease is exactly what a weak surveillance system does best.

The programme's answer had three parts. First, build the reporting system before you need it: by the 1970s the eradication teams had turned the weakness of national reporting into a strength by going out and looking, visiting markets, schools and clinics rather than waiting for returns to come in. Second, pay for information. Rewards were offered, and publicised with pictures of the rash, to anyone who reported a case of smallpox that turned out to be genuine; the reward was raised repeatedly in the final years, which had the useful effect of generating a flood of reports of chickenpox and measles, each of which had to be checked, and each of which demonstrated that the system was working. Third, wait. The rule adopted was that a country could be certified only after two years with no case, a period chosen to be several times the longest plausible interval between one case and the next.

Between Maalin's recovery in November 1977 and the certification of eradication, teams searched the Horn of Africa and every other recently endemic region, and the rewards went unclaimed. On 9 December 1979 the Global Commission signed its conclusion that smallpox

had been eradicated, and on 8 May 1980 the World Health Assembly formally accepted it.

The reasoning is worth noticing, because it is the same reasoning applied to any claim that something does not happen: an absence is only evidence when you can show that you were looking hard enough to have found it.

Where this leaves us

Smallpox is the only human disease ever deliberately eradicated. It took 184 years from the publication of Jenner's *Inquiry* and about twelve from the decision to do it properly, and it has saved, on any reasonable estimate, more lives than any other single act of public health in history. Nobody has been vaccinated against it routinely since the early 1980s, which means that almost everybody alive under the age of forty-five is fully susceptible to a disease that no longer exists.

Part I ends here, and it ends on the only unambiguous victory the book contains. It is worth summarising what has been established in it, because everything in Part II assumes it.

A vaccine works by showing the immune system the shape of a pathogen without the disease, and the body's memory

does the rest. The trick was stumbled on with cowpox, turned into a method by Pasteur, and given a subject when viruses were found at the end of the nineteenth century. Growing viruses in cell culture, from 1949, turned vaccine-making from an art into an industry. The randomised placebo-controlled trial, from 1954, made it possible to know whether a vaccine worked rather than to believe it. Herd immunity explains why vaccinating a population protects people who are not vaccinated, and the arithmetic of the threshold explains why some diseases are far easier to contain than others.

Part I also contains most of this book's good news. The chapters ahead are about viruses that broke the method: one that writes itself into the genes of the cells it infects and has defeated forty years of vaccine research; ones that cause cancers decades after the infection has been forgotten; ones that come out of forests without warning. They are also about what happened to public confidence in the whole enterprise, which in 1980 was close to total and is not now.

The next chapter opens in June 1981, in Los Angeles, with a report of five cases of pneumonia.

Part Two

The modern era

The Virus That Hid

> *"The fact that these patients were all homosexuals suggests an association between some aspect of homosexual lifestyle or disease acquired through sexual contact." — Editorial note, Morbidity and Mortality Weekly Report, 5 June 1981*

Five young men in Los Angeles

The United States Centers for Disease Control publishes a bulletin every Friday called the *Morbidity and Mortality Weekly Report*. It is a plain document, read by public health officials, and most of its contents are routine. On 5 June 1981 it carried a short piece under the heading "Pneumocystis Pneumonia — Los Angeles".

Between October 1980 and May 1981, it reported, five young men had been treated at three different hospitals in Los Angeles for pneumonia caused by an organism called *Pneumocystis carinii*. Two of them had already died. The rest died soon after it was printed.

What made this worth a page was the organism. *Pneumocystis* is everywhere, and it does nothing to anybody whose immune system is working; it causes pneumonia only in people whose defences have collapsed, which in 1981 meant cancer patients on chemotherapy, transplant recipients on drugs to suppress rejection, and the severely malnourished. Five cases in previously healthy young men in one city in eight months was not a statistical fluctuation. It meant that something was destroying the immune systems of healthy adults.

The report was written by Michael Gottlieb, an immunologist at UCLA who had noticed the first cases, with Wayne Shandera of the county health department and others. Its editorial note observed that all five patients were homosexual men and suggested that something about that might be connected. That observation was correct as an association and disastrous as a frame. It attached the disease to a group rather than to a mechanism, it produced the first name anyone used for it, gay-related immune deficiency, and it meant that when the same collapse began appearing in people who had received blood transfusions, in haemophiliacs, in injecting drug users and in infants, it took longer than it should have to be recognised as the same thing. The name was changed in 1982 to acquired immune deficiency syndrome,

AIDS, which described what was happening and implied nothing about who it happened to.

It is worth holding on to how little was known. No cause, no test, no treatment, and no way to tell an infected person from an uninfected one: only a pattern of infections that healthy people do not get, appearing in people who should not have been vulnerable, and killing almost all of them within two years.



Figure F08-01. *Where the first years were spent. Before there was a test, a cause or a treatment, what medicine could offer was nursing, the treatment of each opportunistic infection as it arrived, and company.*

A virus that writes itself into you

The agent turned out to be a virus, and a particular kind of virus that explains almost everything difficult about this chapter.

> **KEY CONCEPT — A retrovirus**
>
> Chapter 3 described the ordinary arrangement: a virus carries genes, gets into a cell, and the cell's machinery reads those genes and builds copies of the virus. Chapter 13 will set out the usual direction of the traffic in a cell, which is that information flows from the cell's permanent genetic library, written in DNA, out to temporary working copies, written in RNA, which are what the protein factories read.
>
> A retrovirus runs that traffic backwards, which is where the name comes from. Its genes arrive written in RNA, and it carries with it an enzyme that copies RNA back into DNA. That DNA copy is then inserted into the cell's own genetic library, in among the host's genes, and it stays there.
>
> Consider what that means. The virus is no longer a thing in the cell that might be cleared out. It is part of the cell, copied every time the cell divides and passed to every

descendant, and there is no process in the body for finding and removing it, because there is no way to tell it apart from the cell's own material.

HIV's particular cruelty is its choice of cell. It gets in by latching onto a protein called CD4, which sits on the surface of the white blood cells whose job is to co-ordinate the rest of the immune response: the cells that recognise an invader and tell everything else what to do. HIV infects the commanders, hides inside them, and over years destroys them, and the immune system does not so much lose a battle as lose its ability to organise one. The *Pneumocystis* and the rest are the consequence.

This also explains why the disease was invisible for so long. Infection produces a brief flu-like illness and then nothing: years, often a decade, during which the person is well, is infectious, and has no idea. By the time the first case appeared in Los Angeles the virus had been spreading quietly for a generation.

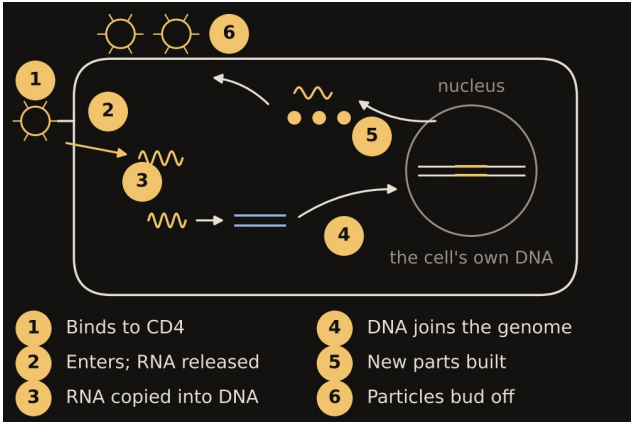


Figure F08-02. *Why this infection cannot be cleared. Step four is the difference: the virus writes a copy of itself into the cell's own genetic library, where nothing in the body can distinguish it from the cell's own instructions.*

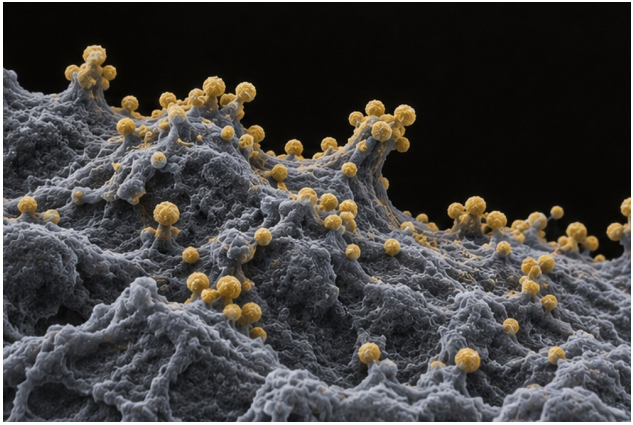


Figure F08-03. *The thing itself. Each particle leaving this cell carries a copy of a genome that is already written into the cell's own, which is why no immune system has ever cleared this infection and why no vaccine has yet been made against it.*

Paris and Bethesda

Finding the virus took two years and produced one of the least edifying arguments in modern science.

At the Pasteur Institute in Paris, a team around Luc Montagnier took a sample from a patient with swollen lymph nodes in early 1983. Françoise Barré-Sinoussi, working with Jean-Claude Chermann, detected the signature of a retrovirus in the culture: the activity of the enzyme that copies RNA into

DNA. They isolated the virus and called it LAV, for lymph-adenopathy-associated virus, and published it in *Science* on 20 May 1983. The paper was carefully hedged. They had a new retrovirus from a patient with the right symptoms; they did not yet claim it caused AIDS. Montagnier presented data pointing that way in September 1983 to a small group of experts, among them Robert Gallo.

Gallo ran a laboratory at the National Cancer Institute in Bethesda and had discovered the first two human retroviruses known, HTLV-I and HTLV-II. In the spring of 1984 he submitted a series of papers announcing a new retrovirus, HTLV-III, as the probable cause of AIDS, and on 23 April 1984 the American Secretary of Health and Human Services, Margaret Heckler, announced that the cause of AIDS had been found by Gallo's laboratory and that a blood test would follow. On the same day, Gallo filed a patent for that test.

It became clear quite quickly that HTLV-III and LAV were the same virus, and in January 1985 that they were not merely the same kind of virus but the same isolate: the sample Gallo's laboratory had grown had come from Paris. Whether that happened by contamination or otherwise was argued over for years, through investigations of scientific misconduct and a dispute between two governments over patent royalties. It was settled in 1987 by an agreement

brokered at the level of heads of government, naming Gallo and Montagnier co-discoverers and splitting the proceeds. The two men disliked each other for the rest of their lives.

In 2008 the Nobel Prize in Physiology or Medicine went to Montagnier and to Barré-Sinoussi, who had done the bench work and had been the least prominent figure in the quarrel. The other half went to Harald zur Hausen, whose work is the subject of the next chapter. Gallo was not included, which was its own kind of verdict.

> **HOW WE KNOW — That this virus caused this disease**
>
> Finding a virus in a sick person does not show that the virus caused the illness. A collapsing immune system invites many viruses in, so one found in an AIDS patient might be a passenger rather than the driver. Four lines of evidence settled it, and none would have been enough alone.
>
> Consistency: once a reliable antibody test existed, from 1985, the virus was found in essentially everyone with the syndrome and in very few people without it. Sequence: infected but healthy people could be followed, and the order was the same everywhere, antibodies first, then a long silence, then a slow fall in the CD4 cells the

virus attacks, and only then the opportunistic infections. Routes: the virus appeared in exactly the groups the epidemiology had already identified, and in the right proportions. Intervention: when the blood supply was screened, transfusion-acquired cases stopped, and when drugs were developed that attacked the virus specifically, the immune system recovered.

The last is the strongest, and it is the general principle: the most convincing demonstration that something causes a disease is that removing it stops the disease. That is also, as chapter 15 will discuss, exactly the kind of evidence unavailable when the question is where a virus came from rather than what it does.

The blood supply

For most of the 1970s, haemophilia care had been transformed by a product called clotting factor concentrate.

Haemophilia is an inherited shortage of one of the proteins that make blood clot. Without it a knock produces bleeding into a joint, and a lifetime of such bleeds destroys the joints and often the person. Concentrate changed that: a freeze-dried powder, made from human plasma, kept at home in a

fridge and injected at the first sign of a bleed. Children treated this way grew up with usable knees.

The method of making it contained the disaster. To extract enough clotting factor, plasma from thousands of donors was pooled into a single batch, and the concentrate made from that pool went to many patients. One infected donation contaminated the lot. For hepatitis this had been understood and accepted as a price worth paying. For a virus that destroyed the immune system it meant that the treatment which had given haemophiliacs ordinary lives would kill a large proportion of them.

The first cases of AIDS in haemophiliacs were reported in 1982. Heat treatment, which inactivates the virus in the concentrate, was available from 1984 and a test to screen donors from 1985. The gap between knowing and acting has been examined by inquiries in a dozen countries, which found much the same thing everywhere: the risk was recognised earlier than it was acted on, and patients were not told.



Figure F08-04. *The treatment that transformed haemophilia and then devastated it. Each batch was made from plasma pooled from thousands of donors, so a single infected donation reached everyone who received that batch.*

Ireland’s haemophiliacs

Ireland has about four hundred people with haemophilia, which makes it possible to state the damage exactly.

A tribunal of inquiry under Judge Alison Lindsay was established in 1999, heard evidence from May 2000 to November 2001, and reported in September 2002 in a document of more than five hundred pages. It found that 106 people with haemophilia had been infected with HIV and 255 with hepat-

itis C through contaminated blood products. Many had both. By the twentieth anniversary of the report, 114 of them had died.

Set that against the size of the community. Roughly a quarter of all the people in Ireland with haemophilia were infected with HIV, and well over half with hepatitis C, by the treatment that was supposed to let them live normally.

Most of the concentrate was imported from the United States, where plasma was collected from paid donors, and the Irish authorities had little control over what they were buying. But not all of it. The Blood Transfusion Service Board made its own Factor IX, and two batches of it, issued in 1985, infected seven Irish haemophiliacs with HIV, five of whom died. By June 1986 the head of the national haemophilia treatment centre, Professor Ian Temperley, believed that five or six of his patients had been infected by the Irish product. That is a domestic failure, in a product made at home, after heat treatment was available.

The tribunal's report was criticised on publication for the proportion it gave to the people it was about: the personal testimony of sixty-four infected people and their families was summarised in three pages. It found that no individual was responsible.

Anti-D and hepatitis C

The second Irish blood scandal is less well known outside the country and is in some ways worse, because it was entirely home-made.

Anti-D immunoglobulin is given to women with rhesus-negative blood after they have carried a rhesus-positive baby. Without it, the mother's immune system can be primed against the blood type, and a later pregnancy can end in a severely anaemic or stillborn child. It is one of the quiet triumphs of twentieth-century obstetrics, and the Blood Transfusion Service Board had made it in Dublin since 1970 from the plasma of donors carrying the relevant antibodies.

In 1976 a woman being treated at a Dublin hospital, known in the record as Patient X, was found to have unusually high levels of exactly the antibody the product needed, and her plasma was used in manufacture. She had also, it emerged, been infected with hepatitis through the blood given to her during her own treatment. Five batches of anti-D made with her plasma were distributed between January and April 1977.

From July 1977 the Board began receiving reports of women who had developed jaundice some weeks after an anti-D injection: at least seven, possibly as many as ten. Six were

confirmed. Production continued. The connection was not publicly acknowledged for seventeen years.

In February 1994 the Board announced that women who had received anti-D should be tested. About sixty thousand were screened, and around a thousand had evidence of hepatitis C. Counting later cases from a second contaminated period around 1991, roughly 1,200 women are known to have been infected; by 2024 about five hundred of them were still alive.

The tribunal that examined it, under the former Chief Justice Thomas Finlay, reported in 1997. Its terms of reference were drawn narrowly around anti-D, which meant it did not examine how the State had conducted itself in the courts against the women suing it, and that conduct is the part most people in Ireland remember.

Brigid McCole was a farmer's wife from Donegal, a mother of twelve, infected by an anti-D injection in November 1977. She sued. She asked to bring the case under a pseudonym, because of how women with hepatitis C were being treated in their own communities, and was refused. The State's lawyers warned her that if she pursued the case and lost she would face costs. She was by then dying of liver failure. The Board conceded liability and a settlement was reached on 1 October 1996. She died the next day, aged fifty-four, and she

was the first woman publicly acknowledged to have died from the contaminated product.

What links the two Irish scandals is not a virus; one was HIV and the other hepatitis C. It is a pattern. In both, a blood product that was a genuine medical advance was made in a way that carried a known risk of viral contamination; in both, evidence of harm arrived early and was not acted on; and in both, the institutions concerned protected themselves before they protected the people they had injured. Chapter 18 will ask how a modern system detects and admits a rare harm, and the answer owes a good deal to how badly this was done.

1996

For fifteen years the treatment of HIV was a holding action. The first drug, zidovudine, licensed in 1987, slowed the virus for a while and then stopped working, because a virus that copies itself carelessly and in enormous numbers will throw up a variant resistant to any single drug within months.

The answer, when it came, was arithmetic. If a patient takes three drugs at once, each attacking the virus at a different point, the virus would have to acquire three separate resistance mutations simultaneously to survive, and the odds against that are long enough that it does not happen. The

drugs that made this possible, in particular a class called protease inhibitors, arrived in 1995 and 1996, and the results were presented at the International AIDS Conference in Vancouver in July 1996.

The effect was immediate and visible in the mortality figures. In the United States, deaths among people with AIDS had climbed from 451 in 1981 to 50,628 in 1995. By 1998 the figure was 18,851, a fall of 63 per cent in three years, and it has stayed at or below that level ever since. Doctors who had spent a decade watching their patients die described wards emptying.

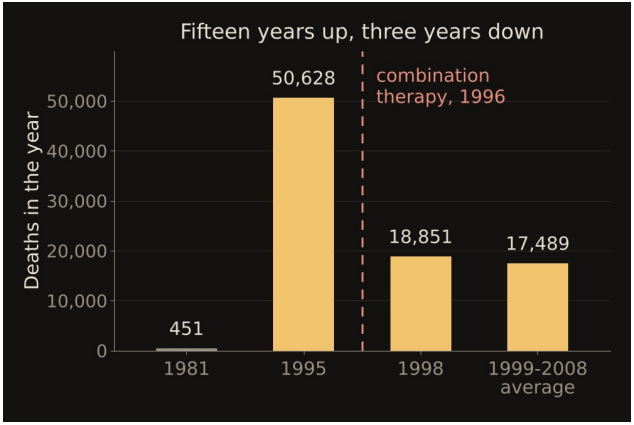


Figure F08-05. *The year the curve turned. Combination therapy arrived in 1996, and deaths among Americans with AIDS fell by nearly two-thirds within three years, having risen without interruption for fifteen. Drawn from United States surveillance figures.*

One drug used in those wards deserves a sentence here, because it returns later in this book under a different reputation. Thalidomide, withdrawn worldwide in 1961 after causing catastrophic birth defects, was found to heal the mouth ulcers that made it impossible for AIDS patients to eat, and to help with the wasting of late disease. It was approved in the United States in 1998, under the strictest distribution controls ever applied to a medicine. Chapter 11 tells that story

whole, because it is the clearest case in this book of why harm and benefit cannot be sorted into different drugs.

THE NUMBER — From two years to a normal life

Before 1996, the median survival after an AIDS diagnosis was one to two years. The diagnosis was, for practical purposes, a date.

In 2023 an analysis of eighteen cohorts across Europe and North America reported that a forty-year-old starting treatment after 2015, with a healthy immune cell count, could expect to live to about seventy-seven if a man and seventy-nine if a woman: within a few years of the general population. Taking a tablet daily, a person infected with HIV in 2026 in a country with a functioning health service can expect to die of something else, at roughly the age they would have died anyway.

Three things the number does not say. It applies to people who are diagnosed, treated and remain on treatment; the remaining gap is largest for those diagnosed late, and late diagnosis is commonest among the people least well served by health systems. It describes high-income countries; of the 41 million people living with HIV in 2025, around 9 million were not on treatment at all. And it is a statement about mortality, not about the ill-

> ness: a daily medicine for life, with its own side effects, in a condition that is still stigmatised, is not the same as not having been infected.

Why there is still no vaccine

Forty-five years after the first report, there is no vaccine against HIV. This is not for want of trying: more money has been spent on it than on any other vaccine in history. The reasons are all consequences of the box above.

Every successful vaccine in this book copies something the body already does. Smallpox vaccine works because people recover from cowpox; measles vaccine works because people recover from measles and are then immune for life. The vaccine shows the immune system an enemy it would eventually learn to beat, and lets it learn in advance. Nobody recovers from HIV. There is no natural immunity to imitate, because the virus hides inside the genome of the very cells that would organise the attack on it.

The surface of the virus is also unusually well defended. The only part exposed to antibodies is a single protein, and that protein varies enormously between strains and even within one infected person; the parts that do not vary are tucked into recesses that antibodies struggle to reach, and much of

the surface is coated in sugars the immune system reads as the body's own. A small number of people do eventually make antibodies that neutralise many strains, but it takes years of continuous infection, and no vaccine has yet been able to coax the immune system into making them to order.

The trials record the consequences. One in Thailand in 2009 produced a modest and much-debated effect; four since have produced none. The last of them, Mosaico, enrolled around 3,900 people and was stopped in January 2023 because there was no evidence the vaccine prevented infection. It was widely described as the end of a whole approach.

What has worked instead is a different kind of prevention. Drugs that treat HIV also prevent it: a person on treatment whose virus is suppressed does not pass it on sexually, and an uninfected person taking the same class of drugs is protected. The newest, lenacapavir, needs an injection twice a year rather than a daily tablet, and in a trial among young women in southern Africa in 2024 there were no infections at all among the 2,134 who received it. That is not a vaccine: it teaches the immune system nothing, and the protection stops when the injections stop. But it is the closest thing that exists, and chapter 21 returns to it.

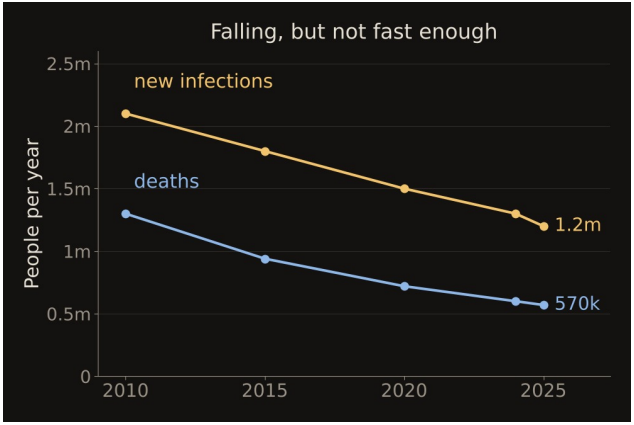


Figure F08-06. *What treatment did without a vaccine. New infections and deaths have both fallen by roughly half since 2010, almost entirely because of drugs rather than immunity, and the decline has slowed well short of the targets set for 2030. Drawn from UNAIDS estimates.*

Where this leaves us

HIV broke the method that Part I built.

Jenner's trick, Pasteur's method, Enders' cell culture and Francis's trial had produced, between them, a way of dealing with viruses that worked so reliably that by 1980 it was reasonable to expect it to work on anything. Chapter 7 ended with a disease deliberately driven to extinction. Eighteen months later a virus arrived that the method could not touch,

and forty-five years and many billions of pounds have not changed that.

What has changed the disease is chemistry rather than immunology: drugs that suppress the virus so completely that the person is well and cannot pass it on. That is a genuine triumph and a different kind of triumph from the one this book has been describing. It requires a diagnosis, a supply chain, a clinic and a lifetime of adherence, which is why 9 million of the 41 million people living with HIV in 2025 were not receiving it, and why the response is faltering rather than finishing.

Two things carry forward. The blood scandals recur in chapter 11, as a question about how institutions behave when their own product has harmed someone, and in chapter 18, as a question about how a modern system is supposed to behave instead. And the research itself built most of the tools Part III depends on: fast genome sequencing, the trial networks, the regulatory pathways for rapid approval, and a generation of virologists. When a new virus appeared in 2019, the people who went to work on it had learned their trade on this one.

The next chapter stays with viruses that take a long time to kill, and with a surgeon from Enniskillen who noticed that a

children’s cancer in Africa followed the pattern of the rainfall.

A Cancer You Can Catch

> *"the first demonstration of persistent tumour virus DNA in human malignancies" — Harald zur Hausen, on finding Epstein-Barr virus in Burkitt's lymphoma*

A chicken with a lump

On 1 October 1909 a farmer from Long Island brought a hen to the Rockefeller Institute in New York. The bird was a barred Plymouth Rock and it had a lump in its breast. The young pathologist who took it, Francis Peyton Rous, found that the lump was a sarcoma, a cancer of the connective tissue, and did what pathologists did: he ground some of it up and injected it into other hens of the same brood, to see whether the tumour would grow in them. It did.

That alone was not remarkable. Transplanted tumours sometimes take. What Rous did next was. He took the ground-up tumour, passed it through a filter of the kind chapter 3 described, fine enough to hold back every bacterium and every

cell, including every cancer cell. He injected the clear fluid that came through. New tumours grew.

By 1911 he had published it. The cancer could be transmitted by something in the fluid that was smaller than a cell: by whatever it was that passed through the filter. Rous was cautious about what that meant, which was wise, because almost nobody believed him. Cancer, as everyone knew, was not an infection. You could not catch it. A chicken tumour was a chicken tumour, and whatever it proved about chickens proved nothing about people. Rous turned to other work for the better part of thirty years.

He was awarded the Nobel Prize in 1966, fifty-five years after the paper, at the age of eighty-seven. He had lived long enough to be right.

The surgeon from Enniskillen

Denis Parsons Burkitt was born on 28 February 1911 at Lawnakilla outside Enniskillen, in County Fermanagh, a few months before Rous published. His father was the county surveyor and an ornithologist of some note, the first person to band birds in order to map how they held territory. That is worth noting, because the son's whole method was his father's: to establish where something happens before asking why.

Burkitt trained in Dublin, served as an army surgeon in the war, and in 1946 took a post with the colonial medical service in Uganda. By the mid-1950s he was at Mulago Hospital in Kampala, a general surgeon dealing with whatever came through the door.

In 1957 a boy was brought in with swellings at all four corners of his jaw. Burkitt had never seen anything like it. The biopsy came back as a small round cell sarcoma, which was a description rather than a diagnosis, and there was nothing he could offer. The face of that child stayed with him. Then he saw another, and another. He began collecting cases, and found that the jaw tumours were common, that they appeared in children, and that they were often accompanied by tumours in other unusual places. He published the series in 1958.

Then he did the thing nobody else had thought to do. He sent out a simple leaflet with a picture of the tumour to mission hospitals across central and southern Africa, asking one question: do you see this? The replies came back, and when he put them on a map a shape appeared. The tumour was common in a broad band across equatorial Africa and absent outside it. The edges of the band were sharp.

To find those edges exactly, Burkitt and two missionary doctors set out from Kampala in October 1961 in a 1954 Ford

station wagon. They were gone ten weeks and covered about ten thousand miles, stopping at hospital after hospital with the same leaflet and the same question. They came back with a line drawn on a map.



Figure F09-01. *Ten weeks and ten thousand miles. Burkitt's method was not laboratory work at all: he established where the tumour was and where it was not, and let the shape of that answer ask the next question.*

> **HOW WE KNOW — What a map can tell you**
>
> The belt Burkitt mapped did not follow borders, tribes, languages, diets or anything else human. It followed two things, and he worked out which by taking the map to a

meteorologist: the tumour appeared where the temperature never fell below about 60 degrees Fahrenheit and where there was at least twenty inches of rainfall a year. At the edges, altitude mattered more than latitude. Go high enough up a mountain inside the belt and the tumour vanished.

That is the signature of something carried by an insect. Temperature and rainfall are what decide where mosquitoes can live, and they decide nothing at all about inherited disease or diet. Burkitt could not say what the insect carried or how it worked. He could say that the pattern of a children's cancer matched the pattern of an infection, and that is a claim with a testable consequence: look for an infectious agent in the tumour.

The limitation is worth stating as plainly as the finding. A map shows association, never cause, and the association here turned out to be more tangled than it looked. The insect-borne thing in the belt was malaria, and malaria is not what causes the tumour; it is what weakens the immune control of the virus that does. Burkitt's map pointed at the right answer by way of the wrong mechanism. That is common in science and is not a failure. A map does not have to be right about why in order to send people to the right place to look.

The belt that resulted covered something like five million square miles, a band across equatorial Africa with a southern extension down the coast of Mozambique where the warm wet conditions reached further from the equator. Inside it, the tumour was the commonest childhood cancer there was, accounting for about half of all cancers in children in some districts. Outside it, in places a few hours' drive away, surgeons had never seen a case. The highland exceptions were the detail that convinced people: pockets within the belt, at altitude, where the climate was cooler and the tumour simply did not occur.

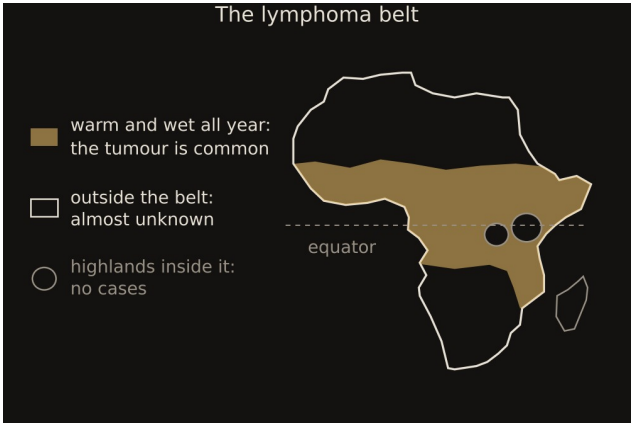


Figure F09-02. *The shape of an answer. The band follows temperature and rainfall rather than any human boundary, which is what told*

Burkitt he was looking at something carried by an insect; the gaps inside it are highlands, cool enough to break the pattern.

He gave a lecture on the tumour at the Royal College of Surgeons in London in May 1961. Twelve people came. One of them was a virologist named Anthony Epstein.

A virus in the tumour

Epstein had been working on Rous's chicken virus and was looking for a human equivalent. He asked Burkitt to send him tumour material from Kampala, which arrived by air in a thermos flask.

Growing it was the hard part. The cells would not form the kind of sheet that virologists knew how to work with, and for nearly three years nothing came of the samples. The breakthrough was a change of method: the cells grew as free-floating lymphoblasts in suspension, not stuck to glass. Once there was a culture, Yvonne Barr and Bert Achong put it under an electron microscope, and in a small proportion of the cells they found virus particles with the unmistakable shape of a herpesvirus. The paper ran in the *Lancet* in March 1964, under the three names: Epstein, Achong and Barr.

It was the first virus ever found in a human cancer. We call it Epstein-Barr virus, and it turned out to be one of the com-

monest infections there is. Most people catch it in childhood and notice nothing; catch it as an adolescent instead and it causes glandular fever. More than nine in ten adults alive today carry it, permanently, in a small population of their own immune cells.

Which raises the problem that occupies the rest of this chapter. If almost everyone has the virus and almost nobody gets the cancer, in what sense does the virus cause the cancer?

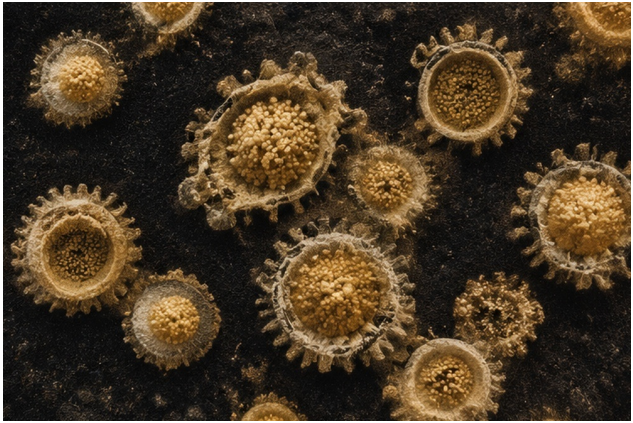


Figure F09-03. *The first virus found in a human cancer. It is also one of the commonest infections on earth, which is the difficulty: a cause that nearly everybody carries cannot be the whole explanation for a disease that hardly anybody gets.*

KEY CONCEPT — How a virus causes a cancer

A cancer is a cell that has stopped obeying the instructions to stop dividing. Those instructions are enforced by a set of proteins inside every cell that act as brakes: they check for damage, they halt division when something is wrong, and if the damage cannot be repaired they instruct the cell to destroy itself. Cancer needs several of those brakes disabled at once, which is why it is mostly a disease of long life. Each one fails rarely; waiting for several to fail in the same cell takes decades.

A tumour virus is a shortcut. It does not create cancer cells. It disables brakes, and it does so in every cell it infects, for as long as the infection lasts.

Why would a virus do such a thing? Not out of malice, and not to its own advantage, since killing the host is no use to it. A virus needs its host cell alive and dividing in order to be copied along with it. The proteins that keep an infected cell alive and dividing are, unavoidably, proteins that interfere with the machinery for shutting a damaged cell down. The cancer is a side effect of the virus making itself comfortable.

Three things follow, and they explain the shape of every story in this chapter. The infection has to be persistent,

because brakes disabled for a fortnight do no lasting harm. The delay is long, decades usually, because the virus supplies only some of the steps and the rest still have to happen by chance. And most infected people never get the cancer, because most never accumulate the rest. The virus loads the dice. It does not throw them.

In Burkitt's belt, malaria is what tips the odds further. Chronic malaria keeps the immune system permanently occupied and the number of infected cells permanently high, and in a few children one of those cells acquires the additional change that finishes the job.

Zur Hausen's long road

Harald zur Hausen was born in Gelsenkirchen in 1936 and began his career in Philadelphia working on Epstein-Barr virus, where he showed that the viral DNA was present in Burkitt lymphoma cells that were producing no virus at all. The genetic material was in the tumour permanently, sitting there, doing something. That finding shaped everything he did afterwards.

In the early 1970s he turned to cervical cancer. It had been suspected of being infectious for a century, on the plain epidemiological grounds that it tracked sexual activity and

barely occurred in women who had none. The favoured candidate was herpes simplex type 2, and zur Hausen set out to find its DNA in cervical tumours. By 1976 he had enough results to know it was not there consistently, which meant it was not the cause.

So he changed candidate, on an argument that sounded thin at the time. Genital warts are caused by papillomaviruses; warts occasionally turn malignant; therefore some papillomavirus might be involved. The objection was obvious. Warts are common and harmless and had been known for centuries, and papillomaviruses were a dull corner of virology nobody worked in. He spent the better part of a decade in it.

The difficulty was that there is no single papillomavirus. There are hundreds of types, each adapted to its own patch of skin or membrane, and the type causing a common hand wart shares too little with the type in a cervical tumour to be found by looking for it. Zur Hausen's laboratory had to work its way through them. Two of his students, Mathias Dürst and Michael Boshart, cloned the sequences that mattered: HPV16 in 1983, HPV18 in 1984. Between them those two types account for about seven in ten cervical cancers.

In the tumours, as in Burkitt lymphoma, the viral DNA was found integrated into the cell's own chromosomes, in the

same place in every cell of the tumour. That is the signature of a tumour grown from one originally infected cell. Zur Hausen shared the 2008 Nobel Prize for it, with Barré-Sinoussi and Montagnier, whose virus chapter 8 described.

By then something had followed from the work that almost nothing in cancer research ever produces: not a treatment, but a prevention.

A vaccine against a cancer

The vaccine is made by a trick worth understanding, because it recurs in Part III. The outer coat of a papillomavirus is built from copies of a single protein. Make that protein alone, with no viral genes anywhere near it, and the copies assemble themselves into an empty shell the right shape and size. The immune system cannot tell the shell from the real thing and raises antibodies against it. The shell can do nothing at all: there is no genome inside it, so there is nothing to infect anyone with.

The first such vaccine was licensed in 2006, covering HPV16 and HPV18 and two wart-causing types. Later versions cover nine types. The schedule was aimed at girls of twelve and thirteen, before most would encounter the virus, which is the reason for a choice that caused trouble almost

everywhere: a vaccine against a sexually transmitted infection, offered to children.

The trials could show that the vaccine prevented infection and prevented the pre-cancerous changes that screening is designed to catch. They could not show that it prevented cancer, because cervical cancer takes decades to appear and no trial runs that long. That gap was filled in by populations rather than trials, and it took until the 2020s.

Scotland started vaccinating girls aged twelve to thirteen in 2008, and keeps linked national records of vaccination, screening and cancer. In January 2024 a team at Public Health Scotland published what those records showed for every woman born between 1988 and 1996: 447,845 of them, of whom 239 had developed invasive cervical cancer.

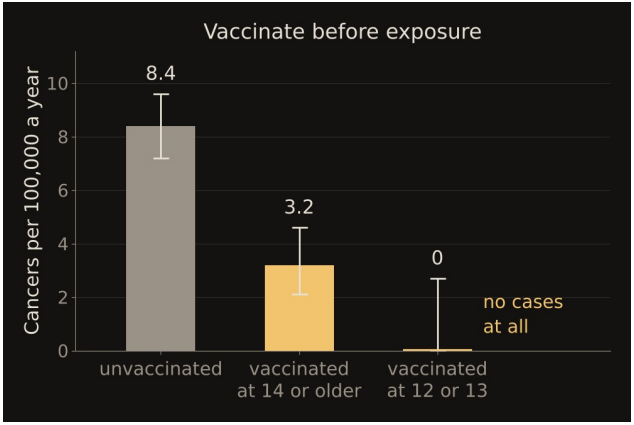


Figure F09-04. *What vaccinating before exposure does. The bar for girls vaccinated at twelve or thirteen is not short; it is absent, because no cancers occurred in that group at all. The upper whisker is what the data can rule out, not a value observed. Drawn from Scottish national linkage records.*

THE NUMBER — Zero, and what zero means

Among the Scottish women who were unvaccinated, invasive cervical cancer occurred at 8.4 cases per 100,000 women per year. Among those vaccinated at fourteen or older, 3.2. Among those vaccinated at twelve or thirteen, no cases occurred: not a low rate, but none.

Zero is the most easily misread number in medicine, so take it carefully. It does not mean the vaccine makes cervical cancer impossible. It means that in this group, over this period, with roughly 136,000 women-years of follow-up, no cancer was observed. The published uncertainty range runs from zero up to 2.7 cases per 100,000 per year, and that upper figure is the honest statement: the true rate could be anything from nothing at all to about a third of the unvaccinated rate, and more years of data will narrow it.

Two further cautions, both from the authors. The women vaccinated at fourteen or older were, many of them, already exposed, which is why their rate is lower but not zero: a vaccine prevents an infection, it does not clear one. And the vaccine studied was the early two-type version. The one in use now covers nine.

What the number does establish is the thing the trials could not. A vaccine given to a twelve-year-old prevents a cancer she would otherwise have been at risk of in her forties. Nothing else in oncology works like that.

Laura Brennan

Ireland began offering the vaccine to girls in first year of secondary school in 2010. Uptake went above the eighty per cent the programme needed and kept climbing, reaching 86.9 per cent of eligible girls completing the course in 2014/15.

Then it collapsed. Campaigning groups formed in 2015, and through social media and public meetings they circulated accounts of girls left chronically ill by the vaccine. The symptoms described were real and the distress of the families was real; what was missing was any evidence connecting them to the vaccine, and the European Medicines Agency reviewed the question in 2015 and found none. Evidence of absence is a poor match for a parent's testimony, and the parents were the ones on the television.

Completion fell to 72.3 per cent in 2015/16. The following year it was 51 per cent. In the space of two academic years, Ireland had gone from one of the best-vaccinated countries in Europe to barely half of a cohort protected, with no change whatsoever in the evidence.

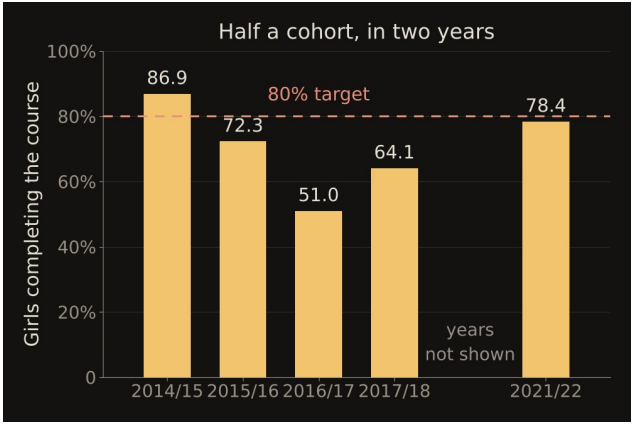


Figure F09-05. *A collapse with no evidence behind it, and a recovery with no new evidence either. Nothing about the vaccine changed across these years. What changed was who was talking about it. Drawn from Irish national immunisation statistics.*

What turned it round was not a study. The Health Service Executive, the professional bodies and the cancer charities formed an alliance in 2017 and ran a campaign, and the campaign needed a voice.

Laura Brennan was born in Ennis, County Clare, on 30 September 1992. She was diagnosed with cervical cancer in December 2016, at twenty-four. In September 2017 she was told it was terminal, and she contacted the HSE to say she wanted to speak publicly in favour of the vaccine. She was

twenty-five, she had a disease the vaccine prevents, and she had been born a few years too early to have been offered it.

She appeared first in a campaign video in March 2018, then on the *Late Late Show*, then at the World Health Organization. She was very good at it: funny, direct, impatient with sentimentality, and entirely unwilling to let her own illness be used as anything but an argument. She died on 20 March 2019, aged twenty-six. A documentary she had been filming, *This Is Me*, was broadcast that May.

Uptake had already begun to recover before she started and kept rising afterwards, reaching 64.1 per cent for her first campaign year and 78.4 per cent by 2021/22, with the programme extended to boys in 2019 and the dose schedule simplified in 2023. The Irish recovery is now cited internationally as the case study, because almost nowhere else that lost an HPV programme has got one back.



Figure F09-06. *Where a cancer is prevented. The entire benefit of this room arrives thirty years later, in women who will never know which of them it was for.*

Where this leaves us

Rous's chickens, Burkitt's map, Epstein's thermos flask and zur Hausen's decade in an unfashionable corner took most of a century to produce a vaccine that stops a cancer before it starts. Something like one in ten cancers worldwide is caused by an infection, and the list keeps growing: cervical and several head and neck cancers from HPV, liver cancer from hepatitis B and C, stomach cancer from a bacterium, and a handful of others.

Two of those now have vaccines. Chapter 6 described the hepatitis B vaccine, licensed in 1981 and the first of its kind; this chapter has described the second. A child vaccinated against both is protected from a substantial share of the cancers that would otherwise have killed her, decades before anyone could tell which ones they would have been.

That is as complete a victory as this book has to offer, and it is also where the trouble starts. A vaccine whose benefit arrives in thirty years and whose cost arrives this afternoon, given to a healthy twelve-year-old who is not ill and will never know she was saved, is the hardest thing in medicine to ask people to accept. Ireland lost half a cohort to that difficulty in two years, over claims that had been investigated and dismissed, and it took a dying woman to get it back.

The next chapter is about a scare of the same kind that went very much further, and about what it cost in a city where children had not died of measles for thirty years.

Twelve Children

> *"the statements in the paper were utterly false" — Richard Horton, editor of the Lancet, on retracting it, 2 February 2010*

A press conference at the Royal Free

On a dull morning in late February 1998, television crews and reporters assembled at the Royal Free Hospital in north London. The hospital's medical school had called a press conference to accompany a paper appearing in the *Lancet* that week, and had taken the unusual step of issuing a video news release with it.

The paper was published on 28 February 1998. It had thirteen authors, and the one who spoke was a gastroenterologist named Andrew Wakefield. What he said at the press conference went considerably beyond what the paper said. He told the journalists that he could not support the continued use of the combined measles, mumps and rubella vaccine, and that

until the question was resolved parents should ask for the three vaccines separately, a year apart.

That was the story. Not the paper, which few journalists read and fewer understood, but a doctor at a London teaching hospital standing up in front of cameras and advising parents against a vaccine their children were about to be given.

What nobody in the room knew was that Wakefield had been retained two years earlier by a solicitor assembling a class action against the manufacturers of MMR, and was being paid out of the legal aid fund at £150 an hour. Nor that several of the twelve children described in the paper were clients of that solicitor. Nor that within days of the press conference Wakefield and business partners would be meeting the medical school's managers to discuss a company to sell products arising from his claims, among them a replacement vaccine. None of this was disclosed, to the journalists, to the *Lancet*, or to the paper's readers.



Figure F10-01. *Where it started. The paper reported on twelve children and reached no conclusion about the vaccine; what reached the public came from the room, not from the journal.*

What the paper claimed

Read now, the paper is a slighter thing than its reputation. It describes twelve children referred to the Royal Free with intestinal symptoms and developmental problems, mostly autism. In nine, the parents or the doctor associated the onset of the behavioural change with MMR vaccination. The authors reported the bowel findings, noted the association the parents had made, and wrote that they had not proved a connection between the vaccine and the syndrome they described.

That last sentence is the one that matters. The paper's own authors said in print that they had not demonstrated a link. What they had was twelve children, a set of symptoms and a group of parents who believed the vaccine had caused them.

There was a proposed mechanism, which is worth setting out because it is rarely stated and because its collapse is part of the story. The suggestion was that measles virus from the vaccine persisted in the lining of the gut, producing a previously undescribed inflammatory bowel condition; that the inflamed gut became abnormally permeable; and that substances which should have stayed in the intestine crossed into the bloodstream and from there affected the developing brain. Each step was a hypothesis. None was demonstrated in the paper, and the chain was long enough that every link needed to hold for the conclusion to follow.

Over the following decade each link was tested and failed. Laboratories could not find vaccine-strain measles virus in the intestinal tissue of affected children when the samples were examined blind, and the finding that had appeared to show it turned out to be a known artefact of the technique. The distinctive bowel condition was not found by other gastroenterologists looking for it. The permeability step was never shown at all.

This is a kind of study called a case series, and the box later in this chapter explains exactly what it can and cannot do. In brief: a case series is a description. It is how doctors raise a question. It is not how anyone answers one, because a group of patients selected because they have both of two things cannot tell you whether either caused the other.

There is also the matter of the timing, which no amount of careful study design can fix. The first dose of MMR is given at around twelve to fifteen months. The behavioural signs of autism are typically noticed by parents in the second year of life. Those two things will coincide in a great many children for no reason other than that they happen at the same age. Any parent whose child regressed shortly after a vaccination would have a vivid and entirely sincere memory of the sequence. The sequence is real. The causation is the open question, and twelve children cannot close it either way.

The fall in uptake

Uptake fell, and it fell for years.

In England, coverage of MMR by a child's second birthday had been just over 90 per cent in the mid-1990s. It declined every year from 1995-96 onwards and reached its lowest point in 2003-04 at 79.9 per cent. It did not return to its

former level for the better part of a decade: 84.9 per cent in 2008-09, 88.2 in 2009-10, 89.1 in 2010-11.

Ireland was worse, and was worse sooner. Irish MMR uptake at twenty-four months was 76 per cent in 1999 and 79 per cent in 2000, against the 95 per cent that stopping measles requires. In the fourth quarter of 2001 it reached 69 per cent. Regional figures were worse still: in mid-2000 the health boards were reporting 76 per cent in the east, 68 in the midlands, 73 in the mid-west, and in parts of north Dublin below 70.

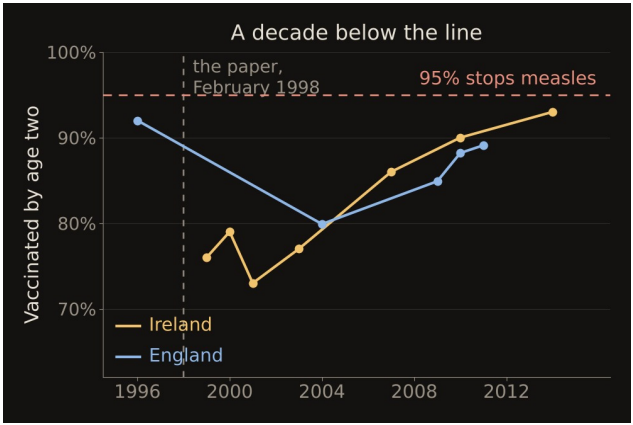


Figure F10-02. *Two countries, one trough. Both series are uptake at twenty-four months, plotted at the years for which published figures exist. Neither country was near the 95 per cent that stops measles spread-*

ing, and Ireland was below it well before 1998. Drawn from national immunisation statistics.

It is worth being careful here, because the chart invites a simpler story than the evidence supports. Irish uptake was already short of the target before the press conference; there were other reasons, including a long-running difficulty with the childhood schedule and with how the vaccine was delivered. The scare did not create the gap. It widened it, held it open for years, and made it fashionable, and the consequence was that by 2000 Ireland had a large population of unvaccinated children in a small number of places.

That is all measles needs.

Dublin, 2000

The outbreak began in December 1999 and ran until July 2000. Most of it was in north Dublin.

Measles is the most contagious disease in this book. One infected person in a susceptible population infects something like twelve to eighteen others, which is why the threshold for stopping it is so high. In a district where fewer than seven children in ten were protected, it moved through the schools and the flats and the waiting rooms at something close to its theoretical maximum speed.

By the tenth of June 2000 the Department of Health had been notified of 930 cases, against 148 in the whole of the previous year. The final national figure for the year was 1,603.



Figure F10-03. *What measles looks like when it is not mild. In this outbreak one admitted child in eight needed intensive care, and more than half of those needed a ventilator.*

THE NUMBER — What happened at one hospital

The Children's University Hospital in Temple Street served the worst-affected part of Dublin, and its staff entered every case into a database as it arrived. The fig-

ures below are from the paper they published afterwards, and they describe one hospital in one outbreak.

355 children came to the hospital with measles. 111 of them were admitted, most for dehydration, nearly half with pneumonia or inflammation of the lungs, a third with infection of the windpipe. 13 of those admitted, one in eight, needed intensive care. 7 of the 13 had to be put on a ventilator. 3 children died.

Set the last number against the thing it is usually compared with. Measles in a well-nourished child in a wealthy country kills roughly one in a thousand to one in several thousand. That is the figure people have in mind when they call it a mild childhood illness, and as a description of the average case it is not wrong. It is also the reason the disease is dangerous in exactly the way it is: an illness that is unpleasant for almost everyone and lethal for a few needs a great many cases before it kills anyone, and once it is spreading freely it very quickly has them.

The arithmetic cuts the other way too. A one in a thousand risk is invisible to any single parent, and three deaths in a city is a small number against any national statistic. It is not a small number if you are one of the families, and it was not an unavoidable number. All

three children were too young or too unprotected to have been covered, in a country where the vaccine had been free and routine for years.

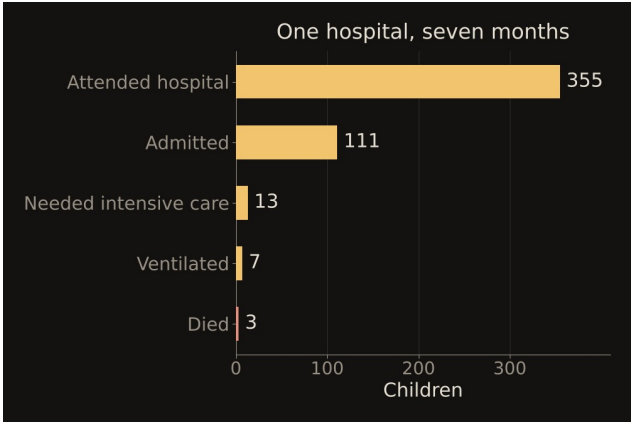


Figure F10-04. *One hospital, seven months. Each step down is a smaller number and a worse outcome, and the shape is what a common illness with a rare severe form looks like when enough people have it at once. Drawn from the hospital's published series.*

Public health responded by lowering the age of the first dose to six months in north Dublin, which is what is done when a population's immunity has fallen far enough that the usual schedule leaves babies exposed for too long. It is a measure of how far things had gone that a European capital in the

year 2000 was vaccinating infants on a schedule designed for epidemic conditions.

What the Danes counted

Meanwhile, elsewhere, the question was being answered properly.

> **KEY CONCEPT — Why twelve is not a small version of six hundred thousand**
>
> A case series and a cohort study are not the same kind of thing done at different scales. They answer different questions, and only one of them can answer this one.
>
> A case series starts with people who already have the condition and describes them. The twelve children at the Royal Free were there because they had both bowel symptoms and developmental problems; that is why they had been referred. Asking what proportion of them had been vaccinated tells you nothing, because there is nothing to compare it with. In a country where nine children in ten receive MMR, nine in ten of any group of children will have received MMR: the ill, the well, the ones with broken arms. A case series has no denominator and no control, and it cannot be fixed by collecting more cases.

Four hundred children selected the same way would be four hundred times as misleading.

A cohort study starts at the other end. You take everybody, vaccinated and not, before anyone has developed anything, and you follow them forward and count what happens. Now there is a comparison, and the comparison is the whole point: of the children who were vaccinated, what fraction were later diagnosed with autism, and of the children who were not, what fraction?

Denmark can do this in a way almost nowhere else can. Every resident has a personal identification number, and vaccinations, hospital diagnoses, births and family relationships are all recorded against it. Nobody is recruited, nobody volunteers, and nobody drops out, which removes the commonest way a study of this kind goes wrong: the people who agree to take part being different from the people who do not.

The result is a study that can do the thing a case series fundamentally cannot. It can find an absence.

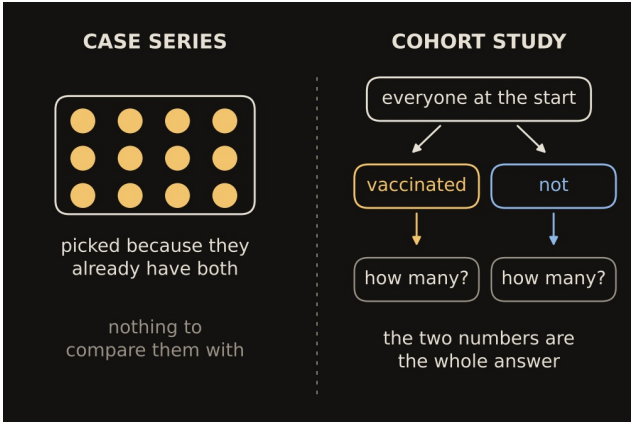


Figure F10-05. *The difference is the comparison. The design on the left begins with the answer already in the room, which is why it can raise a question and never settle one.*

HOW WE KNOW — The Danish cohorts

There have been two, built the same way on different generations of children.

The first, published in the *New England Journal of Medicine* in 2002, followed 537,303 Danish children and found 738 cases of autism among them. Comparing the vaccinated with the unvaccinated gave a rate ratio of 0.92 for autistic disorder and 0.83 for the wider group of autism spectrum disorders, both with ranges comfortably

spanning 1. A ratio of 1 would mean the two groups were identical. A ratio above 1 would mean more autism in the vaccinated. These were, if anything, slightly below.

The second, published in *Annals of Internal Medicine* on 5 March 2019, followed 657,461 children born between 1999 and 2010, of whom 6,517 were later diagnosed with autism. The fully adjusted hazard ratio was 0.93, with a range from 0.85 to 1.02.

The second study was designed to answer the objections raised against the first. It looked specifically at the children a vaccine would most plausibly affect if the hypothesis were true: those with a sibling already diagnosed with autism, those with other known risk factors. It looked at whether diagnoses clustered in the months after vaccination, as they should if the vaccine were triggering something. They did not.

Read the numbers honestly and they do not prove that MMR has never contributed to a single case of autism in any child anywhere; no study of anything can prove that. The upper end of the 2019 range, 1.02, leaves room for a risk increase of up to about two per cent, which in a cohort of that size would be a handful of children over a decade and is as likely to be statistical noise as anything

> real. What the studies establish is that the large effect the hypothesis required is not there. If MMR caused autism at anything like the rate the scare implied, 657,461 children would have shown it.

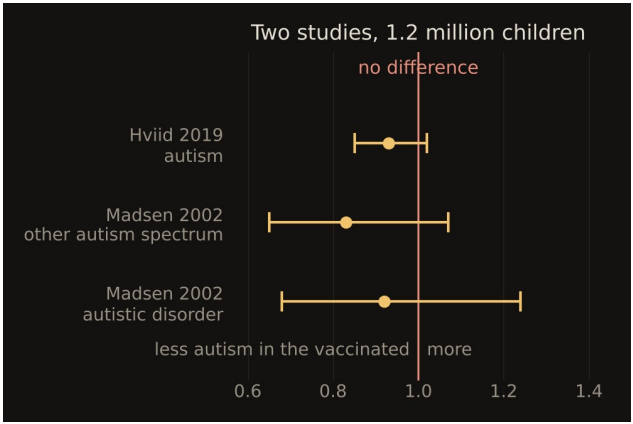


Figure F10-06. *What "no association" looks like when it is drawn. Each bar is a range the true value is likely to lie within; all of them cross or sit below the line where the vaccinated and unvaccinated are identical. Drawn from the published Danish results.*

The investigation

The paper's unravelling took twelve years and was not the work of the medical establishment.

In February 2004, six years almost to the day after the press conference, the *Sunday Times* published an investigation by Brian Deer revealing that the research had been funded through the Legal Aid Board, that several of the children in the paper were litigants in the action Wakefield was being paid to support, and that none of this had been declared. The *Lancet*'s editor called the paper fatally flawed. The following month, ten of the thirteen authors formally retracted the paper's interpretation. Wakefield did not.

Deer kept going. The General Medical Council opened what became the longest fitness-to-practise hearing in its history, obtaining the children's actual medical records and comparing them with what the paper had described.

That comparison is the part of the story that is usually skipped, and it is the part that matters most. The paper described a sequence: normal child, vaccine, regression within days. The records did not. For several children the developmental concerns had been documented before the vaccination the paper placed them after. For others the interval between the two was months rather than days. Diagnoses recorded in the hospital's own notes differed from the diagnoses printed in the journal. The problem was not that the twelve children had been over-interpreted. It was that the

twelve children in the paper were not, in several respects, the twelve children in the notes.

On 2 February 2010 the *Lancet* retracted the paper in full. On 24 May 2010 the GMC found Wakefield guilty of serious professional misconduct, describing his conduct as dishonest, unethical and callous, and struck him off the medical register.

One thing should be said for completeness, because leaving it out would be the kind of tidying-up this book is supposed to avoid. John Walker-Smith, the senior paediatric gastroenterologist who had been struck off alongside Wakefield, appealed. On 7 March 2012 the High Court quashed both the finding and the sanction, holding that the GMC panel's reasoning had been inadequate and superficial and in places simply wrong. That judgment is about the fairness of a disciplinary hearing and nothing else, and the judge said so in terms: there is now, he wrote, no respectable body of opinion supporting the hypothesis that MMR and autism are causally linked. Wakefield had by then abandoned his own appeal.

The GMC panel, for its part, had been careful to state that its case was not concerned with whether there is or might be any link between MMR and autism. It was concerned with how twelve children had been investigated and how the res-

ulting paper had been written. The two questions were settled separately, by different people, and both are settled.

A closed question, and what it cost

This book takes the view that contested questions deserve the strongest version of each side, set out fairly, with the reader left to judge. That is how chapters 15, 16 and 18 are written, and it is how this chapter would be written if the question were still open.

It is not. There is no serious scientific dispute about whether MMR causes autism. The hypothesis was investigated across several countries and tens of millions of children, by people with no stake in the answer, and it failed. Pretending otherwise in the name of balance would not be fair-minded; it would be inaccurate, and it would leave a reader worse informed than when they started.

Being clear about that is not the same as being contemptuous of the parents who believed it. They were presented with a doctor at a teaching hospital, speaking on camera, saying he could not support a vaccine their child was due to receive. The undisclosed payments, the litigation and the business plan took six years to surface, and the definitive evidence took longer still. A parent deciding in 1999 had a sincere expert warning on one side and reassurances from the same

health service that had got things wrong before on the other, and no way to tell which was which. The failure was not theirs.

The cost was real and it was not paid by the people who caused it. Britain lost measles elimination. Ireland had an outbreak of more than sixteen hundred cases and buried three children. Across Europe, outbreaks followed everywhere uptake had fallen, and chapter 22 returns to this because it is still happening in 2026, in countries that had eliminated the disease. A generation of children grew up unprotected because of a paper whose own authors wrote that it proved nothing.

Where this leaves us

What makes this story hard is that it is not a story about bad evidence being believed. The evidence was always thin, and said so. The paper was a description of twelve children that declined to draw the conclusion the world drew from it.

What was believed was a person, in a setting that conferred authority, saying something that matched what some parents had already observed with their own eyes, at an age when they were being asked to do something frightening to a healthy child for a benefit they could not see. Everything about the situation favoured the claim, and nothing about the

design of the health system was prepared for the claim to be made by one of its own.

That is the problem chapter 22 is about, and it is the reason this book exists. The science of vaccines was essentially finished by 1980, in the sense that chapter 7 described. What was not finished, and is not finished now, is the question of who people believe.

The next chapter is about real harms: the ones that are genuinely caused by medicines, how they are found, and what happens to the people they happen to. It begins with a drug that did more damage than any other in the history of medicine, and ends with that same drug being used to keep people alive.

The Pig, the Bird and the Needle

> *“absolute non-toxicity” — Chemie Grünenthal, promoting Contergan*

1957 and 1968, the forgotten pandemics

Two influenza pandemics happened within the lifetimes of people still alive, and almost nobody remembers them.

The first began in southern China in early 1957 and was reported from Singapore that February. It was a new combination of influenza genes carrying a surface protein, H2, that nobody’s immune system had met. It went round the world in six months. The United States Centers for Disease Control put the global death toll at about 1.1 million, of whom 116,000 were American; other published estimates run as high as four million, and the spread of those numbers tells you how little was counted at the time.

The second began in Hong Kong in July 1968, carried a new H3 protein, and killed about a million people worldwide by

the CDC's estimate, 100,000 of them in the United States. Both pandemics had vaccines within months, because the machinery built after 1918 and described in chapter 4 could now identify a new strain and grow a vaccine against it in eggs. Neither pandemic closed a school for long. Neither produced a public inquiry. People went to work.

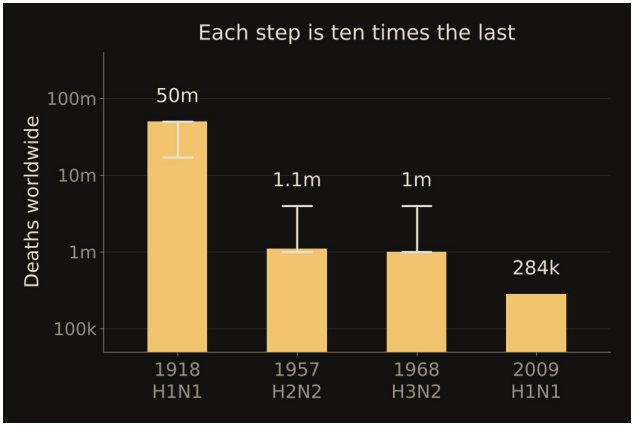


Figure F11-01. *Four pandemics, on a scale where each step is ten times the last. The ranges matter more than the bars: nobody counted 1918, and the 1957 and 1968 figures are reconstructions. Drawn from published estimates.*

That is the useful thing about them. A pandemic is not automatically a catastrophe; what decides which it will be is the particular virus, who it kills, and what the world has ready.

Chapter 20 returns to this with a virus that has not yet jumped.

This chapter is mainly about what happened in between, and what it taught medicine about its own harms. The lesson did not come from a vaccine. It came from a sleeping pill.

The pill that wrote the rulebook

Chemie Grünenthal was a West German firm in Stolberg, near Aachen, founded in 1946. In October 1957 it put a new sedative on the market under the name Contergan. It was promoted as remarkably safe: a sleeping tablet with, in the company's own phrase, absolute non-toxicity. You could not overdose on it. That claim was the selling point, and it was based on animal studies in which enormous doses had failed to kill anything.

It worked well on nausea, so it was prescribed for morning sickness. Fourteen pharmaceutical companies took licences and marketed it in more than forty countries under at least thirty-seven trade names: Distaval in Britain and Australia, Softenon in Austria and Switzerland, and so on. The active ingredient was thalidomide.



Figure F11-02. *An over-the-counter sedative, sold in more than forty countries under more than thirty names. In several of them, including Germany, it did not require a prescription.*

The first child was born with the characteristic malformations in 1956, before the drug was even on sale, to a Grünenthal employee who had been given samples. From 1958 German doctors began reporting an unusual number of babies born with shortened or absent limbs, a deformity so rare that most paediatricians had never seen one. Reports of nerve damage in adults taking the drug arrived from 1959. By April 1961 there were warnings in circulation about both.

What the drug did to a developing embryo is as specific as it is cruel. Taken between roughly the thirty-fourth and fiftieth day of pregnancy, counting from the first day of the last period, it interfered with the formation of limbs, and the pattern of damage depended on the exact day. A dose in one week produced arms that stopped at the shoulder; a dose a few days later affected the legs; others affected eyes, ears, the heart or the internal organs. Outside that window, nothing. A single tablet could be enough.

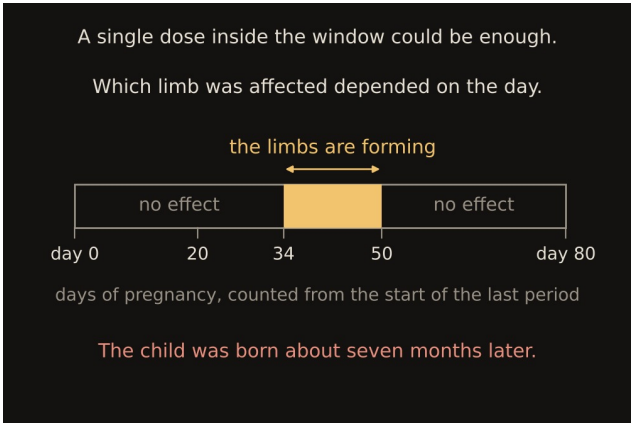


Figure F11-03. *Why the connection was so hard to see. A drug taken for a week in the second month and a child born five months later, with the damage depending on which days the tablets were taken, is not a pattern any individual doctor can perceive from their own patients.*

The figure most often given is about ten thousand children born worldwide with thalidomide malformations, around five thousand of them in Germany, of whom roughly half died at or soon after birth. Those are the survivors of the exposure. Work at Glasgow attempting the full count put the number of pregnancies lost before birth at something like 123,000. Nobody will ever know the true figure.

What went wrong, and what changed

Four separate failures had to happen together.

The first was the testing. Thalidomide was tested on animals as the standards of the time required, and those standards did not require testing on pregnant animals, because it was not then generally believed that a drug could cross from mother to embryo and cause malformation. The species used also happened to be ones that do not show the effect. The drug passed the tests that existed.

The second was the claim. “Absolute non-toxicity” was true in the narrow sense that the drug did not poison the person who swallowed it, and was used to mean something far broader. A sedative you cannot overdose on is genuinely valuable, and the valuable thing was allowed to stand in for a safety record nobody had.

The third was the reporting. There was no system. A doctor in Hamburg who saw two babies with absent limbs had no way to know that a doctor in Liverpool had seen three, and no obligation to tell anyone if he had.

The fourth was the response. Grünenthal had internal evidence by October 1961 that its drug was responsible, and continued to assure doctors it was safe.

> **HOW WE KNOW — Two doctors, six weeks apart**
>
> The connection was made twice, independently, in November 1961, by men who did not know each other and were working ten thousand miles apart.
>
> Widukind Lenz was a paediatrician in Hamburg. He had collected cases of the limb malformation and had been interviewing the mothers about everything they had taken. By his own account he first suspected thalidomide on 11 November 1961, and by 16 November was sure enough to telephone Grünenthal. The company argued. Ten days of discussion followed, by his account driven largely by press reports, and the drug came off the German market on 27 November.
>
> William McBride was an obstetrician in Sydney. He had seen three babies with the same unusual malformations

in his own practice, had noticed that all three mothers had taken Distaval, and reported it to the Australian distributor and then in a letter to the *Lancet* in December.

Two things are worth drawing out. Neither man ran a trial; both did what Burkitt did in chapter 9, noticing a pattern among their own patients and asking what those patients had in common. That is still how most rare drug harms are first spotted. And, less comfortably, the question was answerable only because the effect was so specific: absent limbs in a newborn are rare enough that a handful of cases in one practice is a signal, whereas a drug causing a small increase in something common would not have been found this way, and in 1961 would not have been found at all.

The registry studies later in this chapter are the same method made rigorous. Instead of one doctor's caseload, a whole country's records; instead of a suspicion, a rate against a baseline. The logic has not changed. The denominator has.

In the United States, the drug was never approved. An application to market it as Kevadon was submitted in September 1960 and landed on the desk of a medical officer who had been at the Food and Drug Administration for about a

month. Frances Oldham Kelsey read the file, judged the evidence more testimonial than scientific, and sent it back asking for proof. The company pressed, at increasing levels of seniority. She refused for over a year, citing the incomplete safety data and the reports of nerve damage. Seventeen American children were born damaged, from samples the company had distributed for “investigational” use. It would have been thousands.

What followed was the rulebook. On 10 October 1962 President Kennedy signed the Kefauver-Harris amendments, which for the first time required a manufacturer to prove that a drug worked, not merely that it was not obviously poisonous; required that proof to come from adequate and well-controlled investigations; required informed consent from trial subjects; brought prescription drug advertising under FDA control; and required manufacturers to report adverse reactions after approval. Kelsey stood beside the President as he signed.

> **KEY CONCEPT — Pharmacovigilance, and why a signal is not a cause**
>
> No trial can find a rare harm. A trial of thirty thousand people can reliably detect something that happens to one person in a thousand; it cannot detect something that

happens to one in a hundred thousand, because at that rate you would expect fewer than one case in the whole trial. Yet a vaccine given to fifty million people will produce five hundred cases of a one-in-a-hundred-thousand harm.

So the harm is necessarily found after approval, in use, and the system for finding it is called pharmacovigilance. It has two halves, and confusing them is the commonest mistake made about vaccine safety.

The first half is reporting. Anyone — a doctor, a pharmacist, a patient — can report that something bad happened after a medicine, and these reports are deliberately uncritical, because the point is to catch anything. A raw report establishes only that two things happened in that order to one person.

The second half is testing. When reports cluster, the question becomes whether this happens more often in people who received the medicine than in comparable people who did not: a rate against a baseline, which needs the kind of population records chapter 10 described. Most signals die here. In any large vaccinated population people will have heart attacks and be diagnosed with cancer at exactly the rates they would have

> anyway, and some of those will fall in the week after a vaccination by pure arithmetic.
>
> A system that never raised a false alarm would be one that was missing real harms. The alarms in the rest of this chapter are the ones that survived testing.

A poison returns as a medicine

In November 1964, in Jerusalem, a dermatologist named Jacob Sheskin was treating a man of forty-four who had a complication of leprosy called erythema nodosum leprosum: painful nodules across the skin, fever, weight loss, and such distress that the patient could not sleep. Sheskin had some thalidomide in the hospital pharmacy, left over from before the withdrawal. He gave it as a sedative, reasoning that there was no contraindication for a man.

The patient slept. He also, within three days, got better: the nodules resolved, the fever went. Sheskin stopped the drug and the condition returned; he restarted it and the condition cleared again. He treated six more patients and reported it in 1965, and a trial run by the World Health Organization in the 1970s established it as the best treatment available for the condition.

The two effects turn out to be related. Thalidomide interferes with the growth of new blood vessels and with a particular inflammatory signal, which is both how it damages a developing limb and how it helps here: the same property, in two contexts, with opposite moral weight.

From there it spread. In the AIDS wards of the early 1990s, described in chapter 8, it healed the mouth ulcers that made eating impossible. In 2000, results presented to an American oncology meeting showed it working against multiple myeloma, a cancer of the bone marrow then uniformly fatal within months. In 1998 the FDA licensed it formally, for the leprosy complication, under the strictest distribution controls ever written for a medicine: prescriber, pharmacy and patient all registered, mandatory contraception, mandatory pregnancy testing, limited quantities. In 2006 it was approved for myeloma, the first new drug for that disease in more than a decade.

So the drug that caused the worst disaster in the history of medicine is on the shelf today, keeping people alive, under rules that exist because of what it did. It is the clearest illustration this book can offer of something the rest of the chapter depends on: there is no such thing as a safe drug or an unsafe drug. There is a dose, a person, a condition and a set of controls, and the same molecule can be the best thing

available or the worst, depending entirely on which of those you are talking about.

Fort Dix, 1976

On 4 February 1976 a nineteen-year-old army recruit named David Lewis reported sick at Fort Dix, New Jersey, and insisted on joining his platoon's night march anyway. He collapsed and died hours later. Laboratories found an influenza virus in him that looked like a swine strain, which at the time was thought to be a close relative of the 1918 virus.

Eleven other soldiers tested positive and recovered, and hundreds more at the base had antibodies, so it had spread between people. The virus had, as far as anyone knew, done the thing everyone had feared since 1918.

The United States responded on an extraordinary scale. In March President Ford announced a programme to vaccinate the entire American population. Manufacturers refused to proceed without protection from liability, and Congress passed legislation in August providing it. By mid-December about 45 million Americans had been vaccinated, roughly a quarter of the country, in under three months.



Figure F11-04. *Forty-five million people in under three months, with jet injectors rather than needles. The speed was the achievement and also the problem: a rare harm among that many people is a large number of harmed people.*

Two things then happened. The pandemic did not. The virus never spread beyond Fort Dix, and nobody has ever satisfactorily explained why.

And in late November, reports began arriving of Guillain-Barré syndrome in people who had been vaccinated. It is an autoimmune condition in which the body attacks its own peripheral nerves, causing weakness and sometimes paralysis; most people recover, some do not, and a few die. Investigation found a real excess, which the CDC now puts at

roughly one additional case per hundred thousand people vaccinated. The programme was suspended on 16 December 1976 and never resumed.

The arithmetic is worth sitting with, because it is the arithmetic of this entire chapter. One in a hundred thousand is a very small risk. Forty-five million doses is a very large number. The two together are several hundred people with a serious neurological illness they would not otherwise have had, in a campaign against a pandemic that never came. Had the pandemic come, the same campaign would be remembered as a triumph and the same several hundred people would have been the price of it.

A pandemic that fizzled

In the spring of 2009 a new influenza virus appeared in Mexico and the United States and spread worldwide within weeks. It was an H1N1, the same family as 1918 and as Fort Dix, and it had been assembled from pig, bird and human influenza genes that had met in the same animal.

That assembly is worth a paragraph, because chapter 20 depends on it. Influenza's genome comes in eight separate pieces rather than one, so if two different influenza viruses infect the same cell, the pieces can be dealt out between them like two packs of cards shuffled together, and the cell

can assemble a virus carrying some segments from each. It is called reassortment, and it is why pigs matter: their airways carry receptors that both bird and human influenza viruses can latch onto, so a pig can be infected with both at once. The result can be a virus with a bird's surface protein, which no human immune system has seen, and a human virus's ability to spread between people. That combination is what a pandemic is, and it is exactly what has not yet happened with the virus in chapter 20.

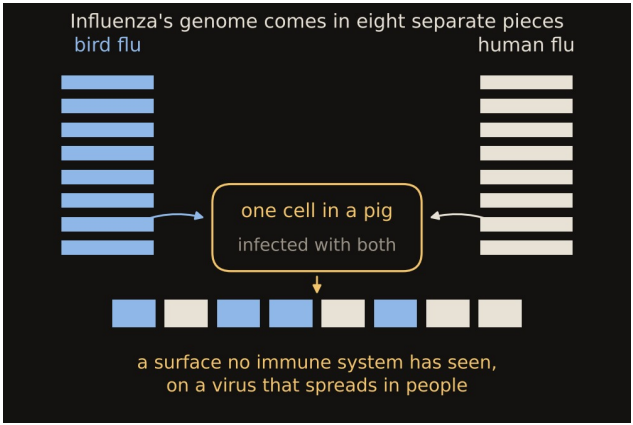


Figure F11-05. *Eight segments, two parents. A flu virus does not have to mutate its way to novelty; it can acquire it in a single step, in a single animal.*

The 2009 virus turned out to be mild in most people, though it fell unusually hard on the young and on pregnant women. The eventual global death estimate was around 284,000, which is within the range of a bad ordinary flu season. By the time the vaccine arrived in late 2009, the worst was over in much of Europe.

Europe used several pandemic vaccines. One of them, Pandemrix, contained an adjuvant: a substance added to provoke a stronger immune response from a smaller amount of virus protein, which matters when you are trying to make hundreds of millions of doses quickly. About 30.5 million people across Europe received it.

The sleeping children

In August 2010 the Swedish and Finnish medicines agencies both reported the same thing: an unexpected number of children and teenagers being diagnosed with narcolepsy, and a striking proportion of them had received Pandemrix.

Narcolepsy is not sleepiness. It is the loss of the small population of brain cells producing a signal called orexin, which holds the boundary between sleeping and waking. Without them the boundary dissolves: people fall asleep without warning, dream while awake, and in the form called cataplexy lose muscle tone when they laugh or are startled, col-

lapsing while fully conscious. It is incurable, it is for life, and it most often begins in adolescence.

The signal was tested, in the way the box above described, and it survived. A Finnish study published in 2012 found a thirteen-fold increase among vaccinated four to nineteen year olds, almost all with onset within six months, and no increase at all among the unvaccinated in the same period, which ruled out the pandemic itself as the cause. England found a fourteen-fold increase in the under-nineteens and put the attributable risk at 1.9 additional cases per hundred thousand doses; in adults the increase was about nine-fold and the risk 0.6 per hundred thousand. Norway, France, Denmark and Ireland found the same thing.

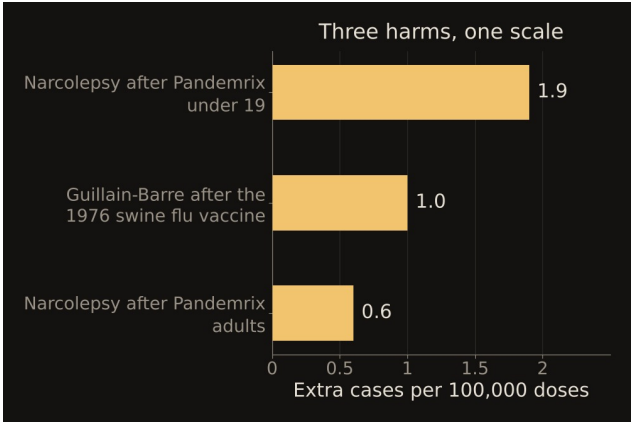


Figure F11-06. *Three real harms, on one scale. Each is rare enough that no trial could have found it and common enough that a national campaign produced hundreds of cases. Drawn from the published attributable risks.*

THE NUMBER — One in how many, and against what

Put the Finnish figure the other way up. Roughly one vaccinated child in sixteen thousand developed narcolepsy. Said like that it sounds negligible, and for any individual child it very nearly is.

Now scale it. Pandemrix went into about 30.5 million Europeans. Even at the lower adult rate, that is several

hundred people with a permanent disabling condition. The disease the campaign was defending against killed an estimated 284,000 people worldwide, and in the European countries that used Pandemrix heavily it killed comparatively few, partly because those countries were wealthy and partly because the virus turned out to be mild.

So was it worth it? The question cannot be answered forwards, which is the whole difficulty. In the autumn of 2009 nobody knew the virus would be mild; the planning assumption, built on 1918, was that it might not be. The reasonable choice given what was known in 2009 produced, given what was known by 2012, several hundred children who will never drive a car.

What can be said is narrower. The harm was real, it was found, it was measured, and the vaccine was withdrawn from use in the age groups affected. That is the system working. Whether the people it worked on find that a comfort is a separate question, and the next section is what happened when they asked.

Ireland’s Pandemrix cases

Ireland began vaccinating with Pandemrix in October 2009, and between seventy-five and a hundred Irish children and young adults were subsequently diagnosed with narcolepsy.

Aoife Bennett was sixteen, at school in County Kildare, when she was vaccinated in December 2009. She was later diagnosed with narcolepsy and cataplexy. She sued the Minister for Health, the Health Service Executive, the medicines regulator and the manufacturer, and her case was taken as the test for all the others.

It ran for twenty-one days of hearings over five weeks, and the State fought it. The court heard, among other things, that safety data on the vaccine had not appeared in Irish medicines reporting. In November 2019 the State settled, on confidential terms, without admitting liability. It was left paying her legal costs of €6.3 million, which is a figure worth setting beside the ten years she had spent waiting.



Figure F11-07. *The afternoon. Narcolepsy is not tiredness; it is the loss of the brain cells that keep sleeping and waking apart, and it does not resolve.*

A second case followed, and around eighty remaining claims were then settled under a framework benchmarked at about half the value of each claim. Families objected publicly to being given twenty-eight days to accept, and to a decade of litigation against a State that had asked them to take the vaccine in the first place.

This is the point of putting this chapter where it is. Chapter 8 described a harm that was real and denied, and chapter 10 one that was not real and could not be disproved to the satis-

faction of the people who mattered. Here the harm was real, the State had caused it by doing something defensible, and it still spent ten years in court.

The Irish thalidomide survivors have had a longer version of the same experience. More than twenty-seven of them are alive and receive lifelong support from the Irish State and from the German foundation set up in the manufacturer's country. Their campaign for a fuller acknowledgement has been running for most of their lives.

Where this leaves us

Three harms in this chapter were real: limbs, nerves, sleep. Each was found by the same method, which is to notice that a rare thing is happening more often than it should and then do the work of establishing against what baseline. Each was rare enough that no trial could have found it in advance. Each was found faster than the one before, because of machinery built in response to the one before.

That is the honest case, and it is not a comfortable one. It does not say that medicines never harm anyone. It says the opposite: that they sometimes do, rarely, that the system which finds out is the same system that approves them, and that it has got steadily better at this over sixty years from a baseline of nothing at all.

What the system is still bad at is the part that comes after. Finding a harm is a technical problem and has largely been solved. Acknowledging it, compensating the people it happened to, and doing so without making them spend a decade proving it, is not a technical problem, and the record in this chapter is poor. Chapter 18 asks what that costs in trust, and the answer is most of what this book has been about.

The next chapter leaves all of this behind and goes into the forest, to look at where new viruses come from, and at the one occasion when a vaccine was proved to work in the middle of an outbreak it was stopping.

Out of the Forest

> *“Ebola ça Suffit” — Ebola, that’s enough. — the name the Guinea team gave their vaccine trial, 2015*

A mission hospital on the Ebola river

In 1976 there was a Belgian mission hospital at Yambuku, in the north of what was then Zaire, near a tributary called the Ebola. It had a small staff of nuns, a few hundred beds, and five syringes.

The syringes were reused. They were rinsed in warm water between patients and boiled at the end of the day, which was ordinary practice in a hospital that could not afford otherwise and which worked well enough for most of what came through the door. Every morning the outpatient clinic saw a queue of people, many of them pregnant women receiving routine injections, and the same five syringes worked their way down the line.

In August a schoolteacher arrived with a fever and was given an injection of an antimalarial. Within weeks the hospital was the centre of an outbreak of a disease nobody had seen, which killed most of the people it touched within two weeks: fever, vomiting, bleeding, collapse. Of the nuns, most died. The hospital closed, and the closing of the hospital is what stopped it. In all there were 318 cases and 280 deaths, a case fatality of 88 per cent.

The virus was isolated by teams in Antwerp, Porton Down and Atlanta, and under the electron microscope it had a shape nobody had seen in a human pathogen: not a sphere or a rod but a long filament, bent back on itself like a shepherd's crook. It was named after the river, on the reasoning, which turned out to be sensible, that naming it after the village would mark the village for ever.



Figure F12-01. *A shape no human pathogen had shown before. The filaments vary in length but not in width, which is what told the first observers they were looking at one virus rather than debris.*

Two things about Yambuku set up the rest of this chapter. The outbreak was extraordinarily lethal and it went nowhere: it burned through one hospital and a ring of villages and stopped, because a disease that incapacitates people within days and kills most of them does not travel far. And it came out of the forest. Nobody knew from what.

The Metropole Hotel

On 21 February 2003 a sixty-four-year-old medical professor from Guangdong province named Liu Jianlun travelled to Hong Kong for a family wedding and checked into room 911 on the ninth floor of the Metropole Hotel in Kowloon.

He had been treating patients at a hospital in Guangzhou with an atypical pneumonia that had been appearing in the province since the previous November. He felt unwell. He went sightseeing anyway, and then to hospital, where he died two weeks later.

During the day or so he spent at the Metropole he infected other guests: people who had done nothing more than share a corridor and a lift with him. Those guests then flew home. Within weeks there were outbreaks in Singapore, Vietnam, Taiwan and Toronto, each traceable to that floor of that hotel, and something like half the world's eventual cases can be traced back to the one stay.

China had notified the World Health Organization on 10 February of 305 cases in Guangdong, 105 of them in health-care workers, and five deaths, with the assurance that the situation was under control. The WHO issued a global alert on 12 March. The disease was named severe acute respiratory syndrome, SARS, and the agent turned out to be a

coronavirus, a family previously known mostly for causing colds.

SARS ended with 8,096 probable cases and 774 deaths across some thirty countries and territories. One person in ten who caught it died, and in Hong Kong, where there were 1,755 cases, it was nearly one in six.

The hotel is worth dwelling on, because it introduced a word the whole of Part III depends on. Most people with SARS infected nobody at all. A small number infected a great many, and the outbreak was built almost entirely out of those few. One housing complex in Hong Kong, Amoy Gardens, accounted for 329 cases on its own, apparently through a fault in the drainage. A disease that spreads this way is not well described by an average: saying that each case infected about three others is arithmetically true and practically useless, because what actually happened was that most cases infected nobody and a handful infected dozens.

That has a hopeful consequence. If transmission is concentrated in a few events, preventing those few events stops the outbreak, and the events tend to be identifiable: crowded indoor places, hospitals, confined spaces with poor ventilation.

It was stopped, and how it was stopped matters for chapter 16. SARS makes people obviously, unmistakably ill before

they become substantially infectious, which means that isolating the visibly sick isolates most of the transmission. Find cases, isolate them, trace and quarantine their contacts, and the chain breaks. That worked, and by July 2003 it was over. It is the reason the world's pandemic plans in 2019 were built around a disease that announces itself, and the reason those plans were wrong.

Why bats

Both of these viruses, and most of the others in this chapter, came out of bats.

> **KEY CONCEPT — Spillover**
>
> Almost every new human virus is an old animal virus that has changed address. The event is called spillover, and understanding why it is both common and usually harmless is most of what this chapter is for.
>
> A virus is adapted to its host: to the particular shape of the proteins on that host's cells, to that host's body temperature, to that host's immune defences. A virus of bats, meeting a human, is in a place it does not fit. Almost always nothing happens at all. Occasionally it manages to get into a cell and make a few copies, and the person's immune system clears it and they never know. Occasion-

ally it makes someone ill and goes no further, because it cannot get from that person to the next one. Only very rarely can it do all three things: infect a person, make them ill, and transmit onward.

So spillover is not rare. It happens constantly, wherever people and animals are in close contact: hunting, butchering, farming, living alongside. What is rare is the full set of abilities required to turn one spillover into an epidemic, and the rarity is why the arithmetic is unforgiving. If one spillover in ten thousand can spread, then reducing the number of spillovers matters enormously, and increasing it — by moving deeper into forests, by farming more animals more densely, by moving live animals further — matters just as much in the other direction.

There is also a halfway stage. A virus often reaches people through an intermediate animal rather than directly: one that is infected by the reservoir species and has more contact with us. For SARS it was probably civets and raccoon dogs in live animal markets; for the Middle East respiratory syndrome coronavirus it is camels; for Nipah virus it has been pigs and, in Bangladesh, date palm sap that bats have fed on overnight.

Bats recur because of what they are. There are over a thousand species of them, a fifth of all mammals, and they roost in enormous dense colonies in which a virus can circulate indefinitely. They fly, which spreads whatever they carry over distances no rodent manages. And they appear to tolerate infections that would kill other mammals, carrying viruses without obvious illness, probably as a side effect of an immune system adapted to the extraordinary metabolic demands of flight.

The result is a reservoir: a large, mobile, long-lived population in which viruses persist without the deaths that would otherwise extinguish them. Rabies, Nipah, Hendra, Marburg, the SARS coronavirus and almost certainly Ebola all come from bats.



Figure F12-02. *A reservoir, in the technical sense. The colony is large enough, mobile enough and tolerant enough that a virus can circulate in it indefinitely without ever killing it out.*

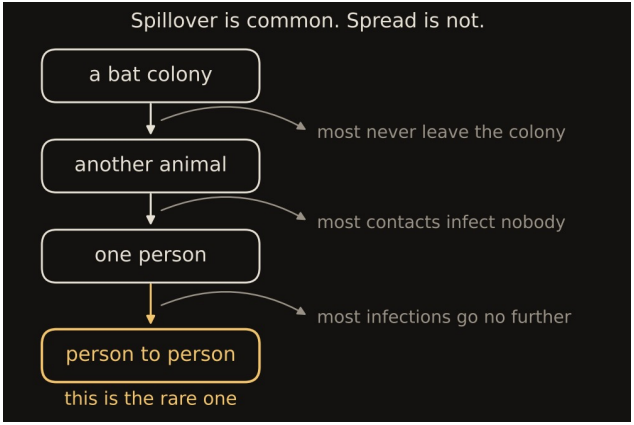


Figure F12-03. *Most spillovers end before the last arrow. The chain that matters is not the one that gets a virus into a person; it is the one that gets it from that person into the next.*

West Africa, 2014

Every Ebola outbreak between 1976 and 2013 had behaved like Yambuku. There were twenty-four of them in that period, and between them they produced about 2,400 cases. The largest single outbreak, in Uganda in 2000, reached 425 people. The pattern was: a village, a hospital, a lethal few weeks, containment.

In December 2013 a toddler died in a village in the forested southeast of Guinea. Three months passed before anyone

identified what it was, and by the time the WHO reported an outbreak on 23 March 2014 the virus had reached the capital and crossed two borders.

What was different was not the virus. It was a Zaire ebolavirus like the others. What was different was the setting.

The forest region where it began is where Guinea, Liberia and Sierra Leone meet, and people cross those borders daily to trade and to visit family. The health systems of two of the three had been dismantled by civil wars within living memory, and there were very few doctors. Funeral practice in the region involves washing and touching the body, which is the single most dangerous thing a person can do with someone who has died of Ebola, because the viral load at death is at its highest. And the thing that had ended every previous outbreak — the fact that a village is a long way from the next village — did not apply, because the roads led to cities of a million people.

The response was slow and then enormous, and the slowness is what the subsequent inquiries dwelt on. Three months passed between the first death and the identification. Months more passed before the scale was acknowledged internationally. By the time treatment units were being built at speed, the outbreak was in three capitals.

It was also, in places, resisted. Families hid the sick, because people who went into the treatment units mostly did not come out, which to a family looks like a place that kills people rather than a place that fails to save most of them. Burial teams were attacked. None of this was irrational given what people could see; it was a reasonable inference from a terrible set of facts, and it is the same problem chapter 22 is about in a different setting.

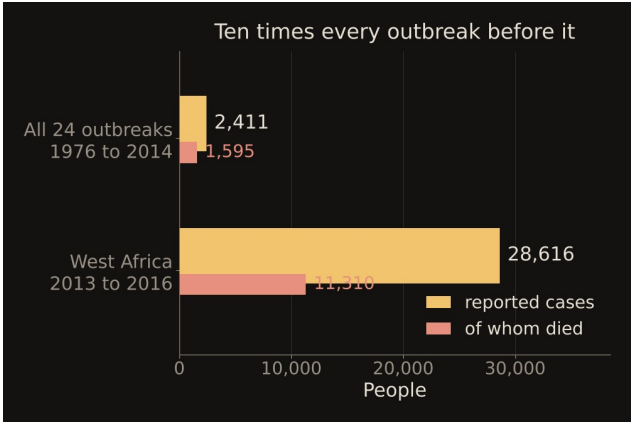


Figure F12-04. *Not a bigger version of the usual thing. In two years the West African outbreak produced more than ten times as many cases as every previous Ebola outbreak in history put together.*

By the time it ended in 2016, there had been 28,616 reported cases and 11,310 reported deaths in Guinea, Liberia and Sierra Leone, with exported cases in four more countries. The WHO's own view is that the true figures were substantially higher. Among confirmed cases with a recorded outcome, nearly two-thirds died.

The reported case fatality across all cases was about 40 per cent, well below Yambuku's 88, and that difference is not the virus becoming milder. It is the difference between counting the cases that reach a hospital in a village of two hundred and counting the cases in a city, where many more mild and uncertain cases are recorded and where some people, eventually, received supportive treatment that kept them alive.



Figure F12-05. *What the response looked like. The suits are the reason the people wearing them survived and the reason they could not comfort anyone.*

A trial inside an outbreak

There had been a candidate Ebola vaccine sitting in development for a decade. It was built on the chassis described in chapter 17: a harmless animal virus, vesicular stomatitis virus, modified to carry the gene for the protein on Ebola's surface. It had worked in monkeys in 2005. Nobody had taken it further, because Ebola outbreaks were small, remote and over before a trial could start, and there was no market.

In 2014 there was a reason. There was also a problem, which was that the conventional way of testing a vaccine — half the volunteers get it, half get a placebo, everyone waits — is close to unconscionable in the middle of an outbreak killing most of the people it touches.

The solution came from chapter 7. Smallpox was eradicated not by vaccinating everyone but by ring vaccination: find a case, vaccinate everyone around them, and let the ring of immunity contain it. That strategy can be turned into a trial design.

> **HOW WE KNOW — Ebola ça Suffit, Guinea, 2015**
>
> The design works like this. Whenever a case of Ebola was confirmed, the response team identified a ring around that person: their contacts, and the contacts of those contacts, typically eighty or so people. That ring is the unit. The rings, not the individuals, were then randomly assigned to one of two groups: vaccinated immediately, or vaccinated twenty-one days later.
>
> Nobody is denied the vaccine. Everyone gets it; the randomisation is only about when. If the vaccine works, cases will appear in the delayed rings during the three weeks before their vaccination and not in the immediate

rings, and the comparison is made among people who all had the same exposure to the same confirmed case.

The trial ran in Basse-Guinée, around Conakry and eight surrounding prefectures, and later in Sierra Leone. In the interim analysis, ninety rings had been randomised: forty-eight to immediate vaccination, covering 4,123 people, and forty-two to delayed, covering 3,528.

In the immediate rings there were no cases of Ebola at all with onset ten or more days after randomisation. In the delayed rings there were sixteen, in seven different rings. The estimated efficacy was 100 per cent, with a confidence interval running from about 75 per cent up to 100.

On that result, in July 2015, the randomisation was stopped, because continuing to delay anyone's vaccination could no longer be justified. Everyone in every subsequent ring was vaccinated immediately. The final published analysis, with more rings, gave the same answer and a slightly narrower interval.

Two cautions, both from the trialists. A hundred per cent is an estimate, not a guarantee; the interval is what matters, and it allows for real efficacy as low as about three-quarters. And the design tests protection in people who

> have just been exposed, in rings, which is a narrower question than how a vaccine performs in a general population.
>
> The achievement is nonetheless what the title of the box says. A vaccine was proved to work in the middle of the outbreak it was stopping, using a design that did not require anyone to go without it.

The vaccine was licensed in 2019, and it has been used in every Ebola outbreak since, deployed within days of confirmation. The disease that killed eleven thousand people in two years is now one that an outbreak response can usually ring-fence in weeks.

The ones that did not spread

It is worth ending on the viruses that did not become pandemics, because they are the common case and the one nobody hears about.

The Middle East respiratory syndrome coronavirus appeared in 2012 and is still appearing. It comes from camels, it kills roughly a third of the people it infects, and it has never learned to spread efficiently between people: almost all hu-

man-to-human transmission has been in hospitals, where the usual precautions had not been taken.

Nipah virus emerged in Malaysia in 1998, from bats, through pigs. It kills between forty and seventy-five per cent of those infected, and in Bangladesh it recurs almost every year, usually through raw date palm sap collected overnight in pots that bats have drunk from. It transmits between people only poorly, and almost entirely in households and hospitals.

Zika had been known since the 1940s as a trivial illness until it reached the Americas in 2015 and it emerged that infection in pregnancy could cause catastrophic damage to a developing brain. It spread widely and then subsided, as mosquito-borne infections do when enough of a population has been infected.

These are usually described as near misses, which is the wrong frame. MERS is not a pandemic that nearly happened; it is a virus that has had more than a decade of opportunities and has not acquired the one ability it would need. Nipah has had the same opportunity every winter for twenty-five years. The reasonable conclusion is not that we have been lucky thirteen years running, but that the last step in the chain — reliable transmission between people — is genuinely difficult, and that most viruses never manage it.

What that leaves is a surveillance problem rather than a prediction problem. Nobody can say which spillover will be the one. What can be done is to notice them quickly: to sequence unusual clusters of illness, to have the sequence in the open within days, and to have the platform ready to make a vaccine against whatever the sequence turns out to be. All three of those are about machinery that exists before the event, not insight during it.

THE NUMBER — Lethality is not danger

Put four outbreaks beside each other and the relationship inverts.

Yambuku in 1976 killed 88 per cent of the people it infected and reached 318 of them. A DRC outbreak in 2014 killed 74 per cent and reached 66. The West African outbreak killed about 40 per cent of its reported cases and reached more than 28,000. SARS killed about one in ten and reached about 8,000.

The pattern is not a coincidence and it is not quite a law. A virus that kills most of its hosts quickly tends not to travel, because the sick are in bed or dead rather than on a bus. It is the less dramatic ones that go furthest. Seasonal influenza kills a tiny fraction of those it infects and kills far more people every year than Ebola has in its en-

tire recorded history, because it reaches hundreds of millions.

The lesson is that the question “how deadly is it?” is the wrong first question about a new virus. The right first question is “how well does it spread, and can it spread before people know they have it?” A virus that scores badly on lethality and well on silent transmission is the dangerous combination, and chapter 14 is about one.

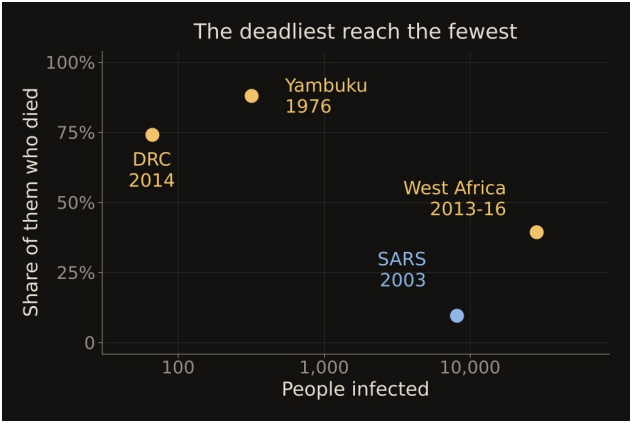


Figure F12-06. *The deadliest outbreaks are the smallest. Each point is one outbreak; the horizontal axis runs in multiples of ten. Drawn from published outbreak figures.*

Where this leaves us

Part II has been about viruses that were harder than the ones in Part I. HIV defeated the vaccine method entirely. Tumour viruses took a century to understand and produced a vaccine against a cancer. The MMR scare showed that the hardest part of a vaccine programme is not making the vaccine. Thalidomide and Pandemrix showed what it costs when a medicine harms somebody and an institution would rather not say so.

This chapter adds the last piece, which is where the next virus comes from. It comes from an animal, through a chance contact, in a place where people and animals meet; it probably comes from a bat; it almost certainly will not spread, and occasionally one does. The Ebola trial showed that the response to such an event can now be fast enough to matter, provided the vaccine already exists.

Which raises the question the last chapter of this part is about. In December 2019 the world had a vaccine for Ebola, none for SARS or MERS, and no way at all to make one quickly for something new. It also had, in two companies with no approved products and a handful of laboratories nobody was paying much attention to, almost everything it would need.

The next chapter is about how that happened, and about a scientist who spent twenty years being demoted.

The Molecule Nobody Wanted

> *“Suppression of RNA recognition by Toll-like receptors” — the title of the paper, 2005*

A demotion in Philadelphia

In 1995 a biochemist at the University of Pennsylvania was told that her work had not come to anything and that she would be moved off the faculty track. She lost the title, the laboratory space that went with it, and the prospect of ever being given tenure there. She was forty, she had come from Hungary some years earlier with her husband, her small daughter and the family’s savings carried out in a teddy bear, and the work in question was messenger RNA.

Katalin Karikó stayed. She found a series of people willing to keep her on in junior positions, worked out of other people’s laboratories, and went on applying for grants that were turned down. The objection was not that her idea was wrong. The objection was that it was obviously impractical, and the people making it had good reasons.

The idea was this. If you could get a synthetic piece of messenger RNA into a person's cells, the cells would read it and make whatever protein it encoded. You would not need to purify a protein, grow a virus, or culture anything; you would just need the sequence. For a drug, that would mean instructing someone's own cells to make the medicine. For a vaccine, it would mean instructing them to make the piece of the virus their immune system needed to learn.

There were two reasons nobody thought it would work, and both were real.



Figure F13-01. *Thirty years of this. The work that made the fastest vaccine in history possible was done, for most of its life, by people with no funding and no title, in buildings where nobody was watching.*

What messenger RNA does

KEY CONCEPT — From gene to protein

Everything a living cell does, it does with proteins. They are the enzymes, the structure, the signals, the pumps and the receptors. A cell is a machine for making proteins and then being run by them.

The instructions are kept in DNA, in the nucleus, and they are kept there permanently and never leave. DNA is a reference library: durable, double-stranded, repaired when damaged, and copied only when the cell divides.

When a cell needs a particular protein, it does not take the library book to the factory. It makes a working copy of the one page it needs, in a related but different material called RNA, and sends that copy out of the nucleus. That working copy is messenger RNA. Out in the body of the cell, structures called ribosomes read it three letters at a time, and each triplet specifies one amino acid; the amino acids are joined in order, and the result is a protein, which folds into its shape and goes to work.

Then the messenger is destroyed. That is the crucial point for this chapter. Messenger RNA is deliberately short-lived, usually minutes to hours, because it is how a

cell controls how much of each protein it has at any moment. It is single-stranded, chemically fragile, and the cell is full of enzymes whose job is to chew it up.

So an mRNA vaccine works like this. You synthesise a messenger carrying the instructions for one protein from a virus. You get it into some cells. Those cells read it and make that protein, display it, and the immune system learns to recognise it. Then the messenger is destroyed, as all messengers are, and the cells go back to what they were doing.

Note what is not happening. The messenger never goes into the nucleus, where the DNA is, and nothing in the cell copies RNA back into DNA — except in retroviruses, which carry their own enzyme to do it and which chapter 8 was about. A synthetic messenger carries no such enzyme. It is a page of instructions that is read and then shredded.

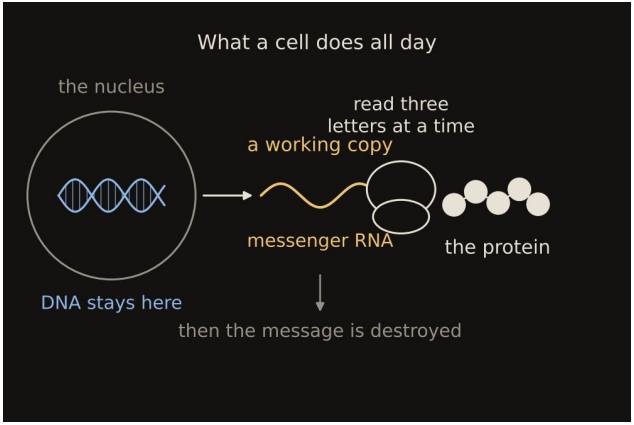


Figure F13-02. *The traffic in every cell, every second. A vaccine of this kind inserts one extra message into a process that is already running, and the message is read and destroyed like any other.*

It is worth pausing on the distinction that separates this from the things it is often confused with, because the confusion is the source of most of what chapter 18 has to deal with.

Gene therapy changes the library. It puts new DNA into a cell's permanent genetic store, either to replace a faulty gene or to add a missing one, and the change is intended to be permanent and to be passed on to that cell's descendants. That is a serious intervention and it is regulated as one.

A messenger RNA vaccine does not do that and cannot. It delivers a working copy, not a library entry. The working copy never reaches the nucleus where the library is kept, and there is no mechanism in a human cell for turning RNA back into DNA. What happens is the same thing that happens with every other message a cell receives: it is read, the protein is made, and the message is destroyed. The distinction is not a reassurance offered after the fact; it is the entire basis of the design, and it is why the technology was attractive in the first place. A permanent change would have required a far higher bar of evidence than anything available in 2020.

The two obstacles followed directly from that description.

The first was fragility. A synthetic messenger injected into a person would be destroyed almost immediately, and even if it survived, it could not cross a cell membrane on its own. RNA is large and negatively charged, and the outside of a cell is not designed to let such things in.

The second was worse, and it was the one Karikó spent a decade on. When you put synthetic messenger RNA into cells, the cells attack it. Not by degrading it, but by raising an alarm: inflammation, interferon, the full response to a viral invasion. The dose that would have been needed to make a useful amount of protein also made the animal ill.

This is not a flaw in the technique. It is the immune system doing precisely what it evolved to do, because loose single-stranded RNA in the wrong part of a cell is the signature of a virus. Chapter 3 described what a virus is: genes in a coat. Loose RNA is what a virus leaves behind. The body is built to notice it.

The 2005 paper

The question Karikó and Drew Weissman asked was why the body does not raise the same alarm about its own messenger RNA, which is present in every cell in enormous quantity.

The answer turned out to be chemical. The RNA a cell makes for itself is not plain: a proportion of its letters are modified after it is made, decorated with small chemical groups. Synthetic RNA made in a test tube has none of these modifications. It is chemically naked, and the sensors that detect viral RNA are, in effect, looking for nakedness.

So they made synthetic messenger RNA with the modifications included, and tested it.

> **HOW WE KNOW — Karikó and Weissman, 2005**
>
> The paper appeared in the journal *Immunity* in August 2005, under a title that gave no hint of what was coming:

“Suppression of RNA recognition by Toll-like receptors: the impact of nucleoside modification and the evolutionary origin of RNA.”

Toll-like receptors are the sensors that detect the molecular signatures of infection. The paper showed that RNA activates three of them, and then that substituting any of several modified building blocks for the ordinary ones abolished that activation. The one that mattered most was pseudouridine, a naturally occurring variant of the letter U that cells use in their own RNA.

The experiment is a comparison, and its strength is how direct it is. Take the same message, make it two ways — once with ordinary uridine, once with pseudouridine — and give it to immune cells. The ordinary version provokes the alarm. The modified version does not. Nothing else about the two is different.

There were two further findings, which only became obviously important later: the modified RNA was more stable, and the cells made substantially more protein from it. One change had solved the immune problem and improved the yield at the same time.

The paper was not ignored, exactly. It was read by the few hundred people who worked on RNA, and it made

very little impression elsewhere. Karikó and Weissman's patent applications were handled with no great urgency. Grant applications continued to be turned down. The University of Pennsylvania did not reinstate her, and in 2013 she left for a small German company with no approved products, where she became a senior vice-president.

In 2023 she and Weissman were awarded the Nobel Prize in Physiology or Medicine, for discoveries concerning nucleoside base modifications that enabled the development of effective mRNA vaccines against Covid-19.

Wrapping it in fat

That solved the second obstacle. The first — getting the thing into a cell at all — was solved by a different set of people working on a different problem.

The answer is a lipid nanoparticle: a droplet of fatty material, a hundred nanometres or so across, with the RNA inside it. Cell membranes are made of lipids, so a lipid particle arriving at a cell is something the cell is prepared to take in, by the ordinary process of swallowing small things from out-

side. Once inside, the particle breaks down and the RNA is released.

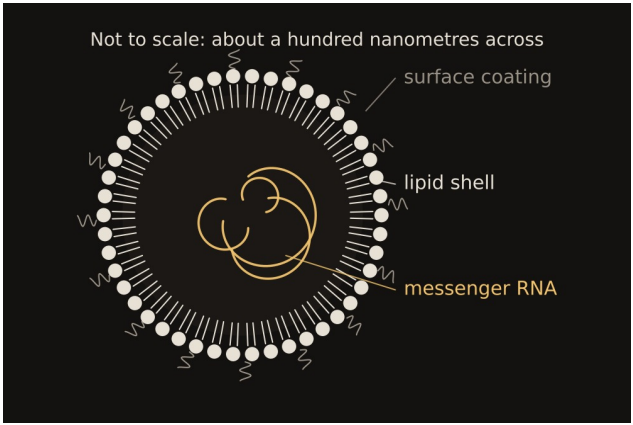


Figure F13-03. *The delivery problem, solved by chemistry rather than biology. Each of the four kinds of lipid in the particle does a different job, and getting the proportions right took longer than designing anything the particle carries.*

This sounds simpler than it was. The particle has to be stable in a vial and unstable inside a cell; it has to be neutral in the bloodstream and acquire charge in the compartment where it needs to break open; it has to be made consistently, at scale, in particles of the right size. Several of the lipids used are not natural substances but designed ones, and the formula-

tions were the product of twenty years of work on delivering genetic material, most of which had failed at something else.

It is worth noticing that the two halves of the technology came from entirely separate traditions. The chemistry of the messenger came from immunology; the chemistry of the envelope came from drug delivery. Neither group was working on a pandemic vaccine. Neither was working on vaccines at all.

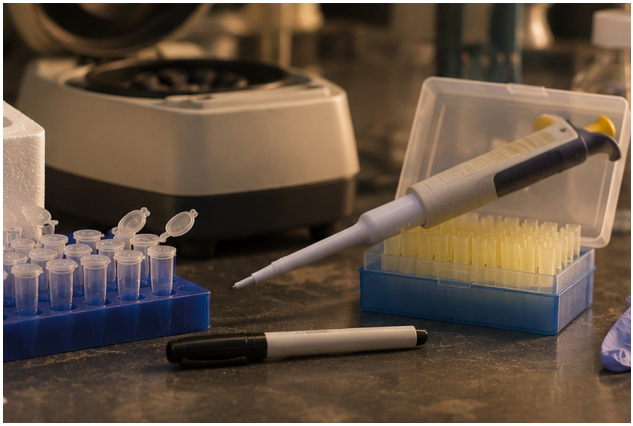


Figure F13-04. *Twenty years of this, in laboratories whose funders had been promised something else.*

Two companies with no products

BioNTech was founded in Mainz in 2008 by Uğur Şahin, Özlem Türeci and Christoph Huber, out of work at the university there. Moderna was founded in Cambridge, Massachusetts, in September 2010. Both were built on messenger RNA, and neither was founded to make vaccines against infectious disease.

They were founded to treat cancer.

The reasoning was the one chapter 21 returns to. Every tumour carries mutations that are unique to it, and those mutations produce proteins that no healthy cell in the body makes. If you could sequence a patient's tumour, identify those proteins, and then instruct the patient's own cells to display them, you would be training that patient's immune system to recognise that patient's cancer. Messenger RNA is the only technology that makes such a thing imaginable, because it is the only one where making a new product means typing a new sequence.

Infectious disease was, for both companies, a sideline: something the platform could obviously also do, worth a few programmes, useful for demonstrating that the technology worked in people. BioNTech had a collaboration on influenza. Moderna had worked on a handful of viruses including

a coronavirus, the one that causes Middle East respiratory syndrome, which chapter 12 described.

The economics are worth stating, because they explain why this was a private undertaking at all. A personalised cancer vaccine is a product with a market of one. It cannot be stockpiled, tested on a population, or manufactured in advance, and no part of the pharmaceutical business was built for it. What such a company must therefore build is not a product but a factory capable of producing a different product every time: sequencing, design, synthesis, formulation and release, repeatedly, for individual patients. Both companies spent the 2010s building exactly that, for reasons that had nothing to do with influenza or coronaviruses.

It is also why Karikó's move mattered more than a hiring decision. BioNTech was the first organisation whose entire plan depended on the problem she had solved, and it funded Weissman's laboratory as well. The 2005 finding had sat for eight years as an interesting result in a small field; it now had a company behind it.

By the end of 2019 neither company had a single approved product. Between them they had spent well over a decade and several billion dollars, and what they had was a method.

Holding a protein still

One more thing had to be in place, and it came from a third direction entirely.

Respiratory syncytial virus is the commonest cause of hospital admission in small children, and for half a century it had no vaccine.

The reason for the delay was a disaster. In the 1960s a killed RSV vaccine was tested in young children, and the children who received it did not merely fail to be protected; when they subsequently met the virus, they became more severely ill than the unvaccinated, and some died. Nothing in Part I had prepared anyone for that. A killed vaccine was the safe, conservative option: it was how polio had been done. The episode left the field with a specific fear — that showing the immune system the wrong thing could be worse than showing it nothing — and that fear, which was well founded, slowed RSV work for forty years.

The problem was eventually understood to be about shape. The protein that RSV uses to fuse with a cell is a spring-loaded machine: it starts in one configuration, and when it does its job it snaps irreversibly into another. The antibodies that actually neutralise the virus recognise the first shape, the one that exists before the spring is released. But the first

shape is unstable, and a protein made in a laboratory flips into the second one almost immediately. Every vaccine attempt had been showing the immune system the wrong shape.

In 2013 a team at the National Institutes of Health led by Jason McLellan and Barney Graham published the solution. They solved the structure of the protein in its pre-fusion shape, identified a site on it that only exists in that shape, then designed and tested more than 150 engineered variants to find ones that stayed locked in the shape they wanted.

That is a technique, not a result. The technique is: work out the structure, find the shape the useful antibodies recognise, engineer mutations that hold the protein in that shape, and make a vaccine out of that. It works for any virus with a fusion protein, which includes every coronavirus.

The same group then applied it to coronaviruses, finding that substituting two prolines at a particular point locked the spike proteins of MERS, the original SARS and a common cold coronavirus into the shape they wanted. So when the sequence of a new coronavirus appeared in January 2020, the people who had done that work knew exactly which two amino acids to change, and where. Chapter 17 gives the timing. It is shorter than you would guess.

December 2019

Here is where things stood at the end of 2019.

A method existed for making a messenger RNA that cells would read without raising an alarm, published fourteen years earlier by somebody who had been demoted for pursuing it. A method existed for getting that messenger into cells, developed by a separate field for a different purpose. A method existed for designing a viral protein that the immune system would learn properly from, published six years earlier by a team working on a children's respiratory virus. Two companies existed with the manufacturing to put the three together, and neither had an approved product, and both had been built for cancer.

Nobody had assembled them. There had never been a reason urgent enough.

> **THE NUMBER — From years to days**
>
> A conventional vaccine starts with the organism. You have to grow it, which means having it, which means getting a sample from a patient and persuading it to multiply in eggs or in culture. Then you weaken it or kill it or purify a protein from it, and each of those is a laboratory programme of its own. Chapter 6 described Maurice

Hilleman swabbing his daughter's throat in 1963; the mumps vaccine was licensed four years later, and that was regarded as extraordinarily fast.

An mRNA vaccine starts with a sequence of letters. There is nothing to grow and nothing to purify from. The design step is choosing which protein to encode and which modifications to make, and it is done on a computer.

The numbers are not comparable, because the design step was never the slow part of a conventional vaccine and the slow parts — the trials, the manufacturing, the regulators — do not go away. But the design step went from years to days, and that turns out to matter more than it sounds, because it moves the start of the clock. Everything else can only begin once there is something to test.

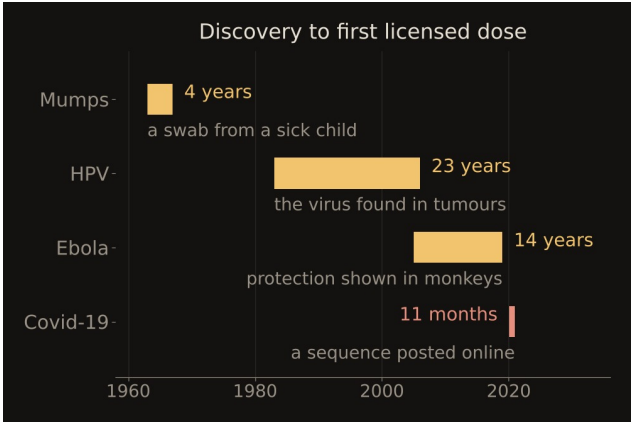


Figure F13-05. *Four vaccines, from the discovery that made each possible to the first licensed dose. The starting points are different kinds of event, which is part of the point: the last one starts from a sequence posted on the internet.*

Where this leaves us

Part II has been about difficulty. HIV broke the method that Part I built. Tumour viruses took a century. A scare built on twelve children cost a generation of measles protection. Three medicines harmed people, were found out, and the finding out took a decade of court time to acknowledge. New viruses keep arriving out of the forest.

This chapter is the exception, and it is worth being clear about what kind of exception it is. It is not a story about foresight. Nobody built this technology for a pandemic. It was built by an immunologist pursuing a question about why cells tolerate their own RNA, by chemists trying to deliver genetic material for reasons unrelated to infection, by structural biologists working on a childhood respiratory virus, and by two companies that wanted to treat cancer. Every one of them was doing something else, and several of them were doing it against the advice of people who controlled their funding.

That is what basic research looks like from inside. It looks like waste, right up until the moment it does not.

In the last week of December 2019 a cluster of pneumonia cases was reported in a city in central China. Within three weeks the genome of the virus causing them would be posted on the internet by a Chinese laboratory, and within forty-eight hours of that, a sequence for a vaccine would have been designed in Massachusetts. None of the people in this chapter had planned for that. All of it worked because they had not been able to stop.

Part III begins there.

PART THREE

Covid-19

December in Wuhan

> *“This is the first pandemic caused by a coronavirus.” — Tedros Adhanom Ghebreyesus, Geneva, 11 March 2020*

Written in November 2026.

A seafood market and a doctor’s message

At about half past five on the afternoon of 30 December 2019, an ophthalmologist at Wuhan Central Hospital received a message from a colleague. Thirteen minutes later he forwarded it to a WeChat group of his old medical school classmates, under the name he used there, with a line to the effect that seven cases of SARS had been confirmed in patients from the Huanan seafood market.

He was wrong about the SARS, and right about everything that mattered. Li Wenliang was thirty-four, an eye doctor rather than an infectious disease specialist, and what he had was a laboratory report and an instinct. He asked the group

not to circulate it beyond themselves, except to tell their families to be careful.

It circulated. Within days he was summoned to the Public Security Bureau and made to sign a statement agreeing that he had made false statements that disturbed public order. He was one of eight people in Wuhan reprimanded that week for spreading rumours. On 28 January, by which time the rumour had become the national emergency, China's highest court publicly criticised the police for reprimanding them.

By then Li had been infected at work. He died at the hospital where he worked in the early hours of 7 February 2020, and the grief that followed on Chinese social media was of a kind the authorities had not anticipated and did not know how to manage. Five more doctors from the same hospital were dead by that June.



Figure F14-01. *Where the first recognised cluster was found. A market of this kind is where chapter 12's spillover chain reaches people: live animals, close quarters, and a great many human beings passing through every day.*

It would be comfortable to make Li a martyr to censorship and leave it there, and the comfort would be misleading. The same day he sent his message, the Wuhan Centre for Disease Control was issuing emergency warnings about the same cluster to the city's hospitals. Officials knew. What the reprimand establishes is not that nobody in authority had noticed, but that the system's first instinct was to control the information rather than the infection, and that instinct cost the world weeks it did not have.

Chapter 15 returns to the question of where the virus came from. This chapter is about what happened next, which is answerable.

A genome on the internet

A laboratory at Fudan University in Shanghai, led by Zhang Yong-Zhen, received a sample from a Wuhan patient on 3 January 2020. By two o'clock in the morning on 5 January they had a complete genome of the virus. Zhang reported it to the authorities and submitted it that day to GenBank, the American public sequence database.

Submitted is not the same as published. The sequence sat in a queue.

What happened on 11 January is one of the hinge moments of the whole story, and it happened because of a friendship. Zhang had collaborated for years with Edward Holmes, an evolutionary virologist at the University of Sydney. Sitting on a plane in Shanghai waiting to take off for Beijing, Zhang gave Holmes permission to release the sequence. Holmes telephoned Andrew Rambaut in Edinburgh, who runs an open scientific forum called virological.org, and at five minutes past one in the morning, Greenwich time, the complete genome of the virus went up on a website where anyone in the world could read it. The post is headed with the

previous day's date, and there is a continuing dispute about whether other sequences had appeared on a different database some hours earlier.

The day after the genome was released, Zhang's laboratory was visited by officials and closed, under what was described as a rectification order.

Hold the two halves of that together, because the rest of the book depends on both. A virus unknown three weeks earlier had been sequenced and published, free, to the entire scientific world, within six days of a sample reaching a laboratory. And the laboratory that did it was shut for doing so.

What the sequence meant, in practical terms, is what chapter 13 described. It is a string of letters. Within forty-eight hours of its appearing, a vaccine had been designed in Massachusetts by people who never saw the virus and never needed to. Chapter 17 is about what happened to that design.

The ship in Yokohama

The *Diamond Princess* left Yokohama on 20 January 2020 with about 3,700 passengers and crew aboard. On 25 January a passenger who had been feeling unwell left the ship in Hong Kong, was tested, and was found to have the new coronavirus. On 3 February the ship was quarantined in the

port of Yokohama and everybody aboard was confined to their cabins for fourteen days.

It was a horrible thing to happen to the people on that ship, and it produced, by accident, the single most useful piece of early evidence anyone had.

HOW WE KNOW — A closed population

Every early estimate of how dangerous this virus was came from counting sick people, which gives you a biased answer. People who are seriously ill go to hospital and get tested; people who feel slightly off do not. So the cases you count are the worst ones, the proportion who die looks terrifyingly high, and nobody can tell whether the gap is large or small.

The *Diamond Princess* removed the bias. It was a closed population of known size with a known period of exposure, and almost everybody aboard was eventually tested whether they felt ill or not.

Of the 3,711 people on board, 712 tested positive — about one in five. Of those 712, 331, which is 46.5 per cent, had no symptoms at the time they were tested.

That figure was widely reported and widely misunderstood, and the correction is the point of this box. “No

symptoms when tested" is not the same as "never had symptoms". Many of those people were tested early and became ill afterwards. A Japanese group later worked through the timing and estimated that the proportion who genuinely never developed symptoms at all was 17.9 per cent, with a range of about 15 to 20 per cent.

Both numbers are true and they answer different questions. If you want to know how many infected people are walking around with no idea, at any given moment, the high figure is closer. If you want to know how many will get through the whole infection without ever feeling ill, it is the low one.

The other finding was the severity. Among the 381 people who did develop symptoms, 37 needed intensive care and nine died. These were mostly elderly passengers, and the ship was not a cross-section of humanity. But it established, in February 2020, that this was a virus that a substantial minority carried silently and that killed around one or two per cent of those it infected on a wealthy cruise ship with a hospital alongside.



Figure F14-02. *An accidental experiment. A closed population of known size, almost all of it eventually tested, is the thing epidemiology almost never gets and spent the rest of 2020 wishing for.*

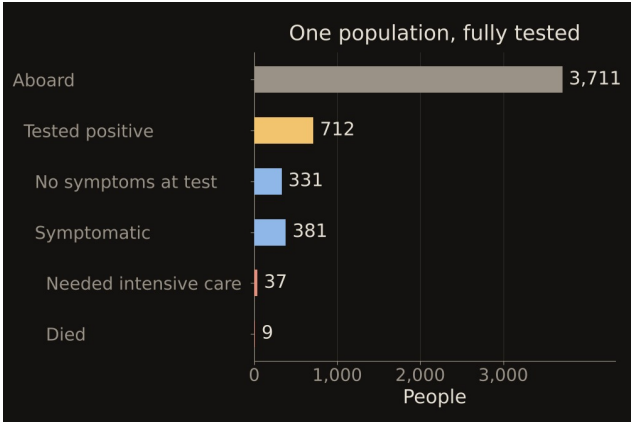


Figure F14-03. *What one fully tested population showed. The second step down is the one that mattered: a disease where nearly half of those infected are not visibly ill cannot be controlled by isolating the visibly ill. Drawn from published surveillance figures.*

That last sentence is why the world's pandemic plans failed. Chapter 12 described how SARS was stopped in 2003: find the sick, isolate them, trace their contacts. That works when illness and infectiousness arrive together. The *Diamond Princess* said, in February, that they did not.

Bergamo

Italy was the first European country to understand what was coming, because it was the first to have it.

The province of Bergamo, thirty miles north-east of Milan, took the worst of it. By the middle of March the crematorium was running twenty-four hours a day. Coffins filled two hospital morgues, then a cemetery morgue, then were laid out inside a cemetery church, where a priest blessed them one after another, by his own account about every ten minutes. On the nights of the eighteenth and nineteenth of March, army convoys — fifteen trucks and fifty soldiers — drove through the town taking coffins to be cremated in other cities, because Bergamo could not deal with its own dead. The pictures of that convoy went round the world.

The local newspaper, *L'Eco di Bergamo*, normally carried two or three pages of death notices. In the middle of March it was carrying ten or eleven, sometimes more than 150 names in a day, and its editor said about nine in ten of them were the virus. He compared the pages to war bulletins.

Behind the coffins was the thing that actually broke. A modern hospital's capacity to keep a person alive through severe pneumonia depends on a ventilator, a bed in an intensive care unit, and a trained person to watch both. Lombardy has one of the better health systems in Europe and it had, like every health system, roughly as many of those as it normally needed. When demand arrives at several times the normal rate and does not recede, the shortage is not of drugs or of

knowledge. It is of rooms and of people, and neither can be manufactured in a fortnight.

What follows from that is the decision nobody wants to describe. When there are more patients who would benefit from a ventilator than there are ventilators, somebody chooses. Italian intensive care doctors published guidance in March 2020 on how to make that choice, and the fact of its publication was as shocking in Europe as the convoy.

Nothing in the previous forty years of this book had looked like that in a wealthy European country. The scenes in Bergamo did more than any model to convince other governments that what was coming was real, which matters for chapter 16, because the decisions taken across Europe in the following fortnight were taken by people who had just watched it. It also fixed, in a great many minds, a picture of what the disease did — and that picture was of the first wave, in the oldest population in Europe, before anyone knew how to treat it.

Dublin, February

The first case on the island of Ireland was confirmed in Northern Ireland on 27 February 2020, in a woman who had returned from Italy. Two days later, on 29 February, the Republic confirmed its first: a male student in the east of the

country who had arrived from northern Italy. The school he attended closed.

What followed in Ireland over the next twelve days is a compressed version of what happened everywhere, and it is worth setting out by date because the speed is the thing people forget.

> **THE NUMBER — Doubling**
>
> On 29 February, Ireland had one confirmed case. By 9 March it had twenty-four. By 11 March it had thirty-four.
>
> Twenty-four to thirty-four in two days is a growth rate that, continued, doubles the total every three or four days. That sounds survivable until you do the arithmetic, so here it is. Starting from thirty-four and doubling every three days: a week later, about two hundred. A fortnight later, about a thousand. Three weeks, five thousand. A month, twenty thousand. Five weeks, a hundred thousand.
>
> Nothing about the growth rate changes across that month. The only thing that changes is that the numbers become large, and they become large suddenly, because that is what repeated doubling does. The difference

between acting on a Monday and acting the following Monday is not one week of cases. It is more cases in that week than in the entire epidemic to that point.

Three cautions, because this number has been misused in both directions. Confirmed cases are not infections; they are infections that were found, and in March 2020 testing was expanding at the same time as the epidemic, so some of the rise is better looking rather than more disease. Exponential growth never continues, because it runs out of susceptible people or runs into the measures taken against it. And a doubling time calculated from twenty-four to thirty-four cases is calculated from very small numbers, where chance matters.

What survives all three cautions is the shape. A disease that grows by a constant percentage looks like nothing for a long time and then like everything at once, and governments in March 2020 were being asked to act at the point where it still looked like nothing.

On 9 March the Irish government cancelled every St Patrick's Day parade in the country, eight days before the seventeenth, along with a three billion euro support package. It was an extraordinary thing to do. The parade is the one day Ireland presents itself to the world, and Ireland cancelled

it with twenty-four confirmed cases and no deaths, before most of Europe had closed anything.

On 11 March the Taoiseach, Leo Varadkar, flew to Washington for the traditional St Patrick's Day meeting with the American president. On 12 March he spoke to the country from there, on camera, in a short address that began by saying he needed to talk to them about coronavirus, and announced that every school, college and childcare facility in Ireland would close that day. Three days later the pubs were asked to shut.

What this virus does

It is worth saying plainly what the thing actually is, because almost all of the argument in the next five chapters depends on it.

> **KEY CONCEPT — How a coronavirus gets in**
>
> Coronaviruses are named for their appearance: a roughly spherical particle studded with projections that, in the first electron micrographs, looked like a crown or a corona. Those projections are a protein, and it is called the spike.

The spike has one job: to open a door. Every cell in your body is covered in proteins that have their own purposes — receiving signals, transporting molecules, regulating blood pressure. One of these, called ACE2, is present on cells in the lining of the nose and airways, in the lungs, and in blood vessels and several other organs. It has nothing to do with viruses. It regulates blood pressure.

The tip of the spike happens to fit ACE2. The virus arrives, the spike binds to that protein as though it were the molecule it was designed to receive, and the cell, following its own ordinary procedures, draws the whole particle inside. Once in, the virus releases its genetic material and the cell begins making more virus.

So the “door” in this metaphor is not a weakness. It is a normal, essential part of a cell that a virus has evolved a key for. That is also why the infection is not confined to the lungs: wherever that protein is present, the virus can get in, which is part of why Covid-19 turned out to be a disease of blood vessels and clotting as well as of breathing.

Two consequences for everything that follows. First, the spike is the part of the virus the immune system most needs to recognise, because antibodies that stick to its tip stop it fitting the door. That is why every vaccine in

chapter 17 is a way of showing the body the spike and nothing else. Second, the spike is also the part that changes most as the virus evolves, which is chapter 19.

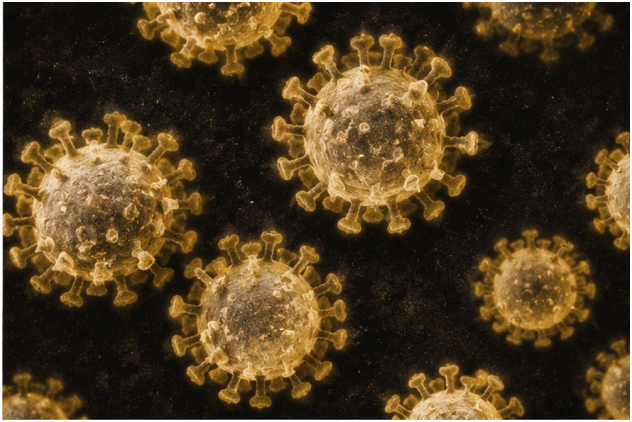


Figure F14-04. *The crown that gives the family its name. Each projection is a copy of the spike protein, and each is a key cut for a door that every one of us has.*

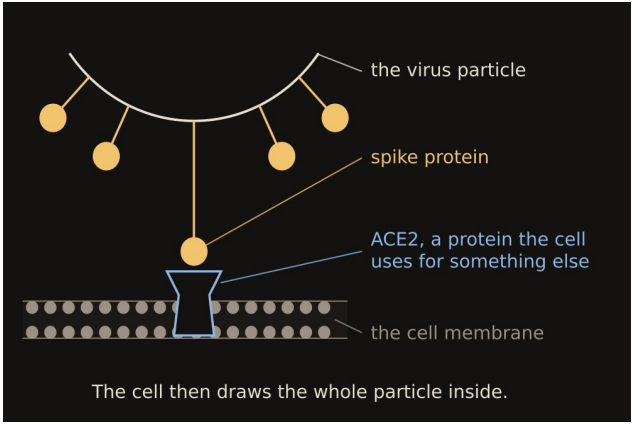


Figure F14-05. *The whole of the problem, and the whole of the solution. Block this one interaction and nothing else happens; every vaccine in the next chapter is an attempt to teach the body to block it.*

What made this virus dangerous was not that it was unusually lethal. Chapter 12 set out why that is the wrong first question, and by the standards of that chapter Covid-19 was mild: something under one per cent of those infected died, against Ebola's forty and SARS's ten.

What made it dangerous was the combination. It spread through the air, efficiently, indoors. It spread from people who felt well, which the *Diamond Princess* had shown by February. And a population of eight billion people had no

immunity to it whatsoever, because no human immune system had ever seen it.

The first of those took longer to accept than it should have, and the delay had consequences. The established model of respiratory transmission divided it into droplets — heavy, falling to the ground within a metre or two, which is why the advice was distance and hand washing — and airborne particles, small enough to remain suspended and drift, which was thought to be the exception rather than the rule. Covid-19 turned out to behave much more like the second. The evidence accumulated across 2020 from outbreak after outbreak in which people had infected each other across a room: a choir practice, a restaurant with a particular pattern of air conditioning, a bus. Guidance eventually caught up, and the practical difference is large. If a virus falls to the ground, you clean surfaces and keep two metres apart. If it hangs in the air, you open windows and change the ventilation, and two metres indoors is not much better than one.

This is the first of several places in Part III where the official position lagged the evidence and was then corrected. Recording that is not an attack on the people who held the earlier position; it was the orthodoxy, it had decades behind it, and the new evidence arrived in pieces. But the lag was real,

and pretending otherwise would make the rest of this part dishonest.

A low fatality rate multiplied by almost everybody is a very large number. That is the whole arithmetic of 2020.

11 March 2020

The World Health Organization had declared a public health emergency of international concern on 30 January. On 11 March, in Geneva, its Director-General said that the organisation had made the assessment that Covid-19 could be characterised as a pandemic, and added that this was the first pandemic caused by a coronavirus.

He also said what the word was being used to convey. In the preceding fortnight, cases outside China had increased thirteenfold and the number of countries reporting them had tripled. He said the WHO was deeply concerned by the levels of spread and severity, and equally by the levels of inaction.

The declaration changed nothing legally. What it did was give every government in the world a date to point at.



Figure F14-06. *The seventeenth of March 2020. The parade had been cancelled eight days earlier, with twenty-four cases in the country and nobody yet dead.*

Where this leaves us

In eleven weeks the world went from a doctor's message in a group chat to a declared pandemic, and almost everything that was going to be argued about for the next six years was already in place by the end of it.

The virus had come from somewhere, and the question of where had already become political rather than scientific, which is chapter 15.

Governments had begun closing things, without exit plans, on the basis of models and of what they had seen in Bergamo, which is chapter 16.

A vaccine had been designed on 13 January, from a sequence, before most people had heard of the disease, and the design would take the rest of the year to become a dose, which is chapter 17.

The vaccine would turn out to have rare real harms and a great many claimed ones, and telling them apart would become the central argument of the decade, which is chapter 18.

And the virus itself was already changing, in ways nobody was yet looking for, which is chapter 19.

What nobody had on 11 March 2020, and what the following chapters are written with, is hindsight. The rules this book follows for contested questions are set out at the front of this part, and the first of them is the hardest to keep: a decision made in March 2020 has to be judged by what was known in March 2020. Almost everything written about this period fails that test in one direction or the other.

Where It Came From

> *"As things stand, all hypotheses must remain on the table." — Tedros Adhanom Ghebreyesus, 27 June 2025*

Written in November 2026.

Two stories, one city

Wuhan is a city of eleven million people in central China. In the last weeks of 2019 a new coronavirus began spreading in it, and seven years later nobody can say with confidence how it got there.

There are two serious accounts, and this chapter is going to give each of them the strongest case its ablest defenders make, which is not the case its opponents attribute to it.

The first is that the virus came out of an animal, as the viruses in chapter 12 did: circulating in bats somewhere in southern China or south-east Asia, passing into another species, reaching people through the wildlife trade, and igniting at a market in Wuhan where live mammals were sold. On

this account Wuhan matters because it is a transport hub with a market, not because of anything in its laboratories.

The second is that the virus came out of a laboratory. Wuhan is home to the Wuhan Institute of Virology, which held one of the world's largest collections of bat coronaviruses and did research on them, some of it involving modifying viruses to study what they might become. On this account the coincidence of the world's foremost bat coronavirus laboratory and the world's first bat coronavirus pandemic in the same city of eleven million is the thing requiring explanation, and an accident during research is the explanation.

Both accounts are coherent. Both have serious people defending them. They cannot both be right, and after six years the evidence that would distinguish them has not been produced.

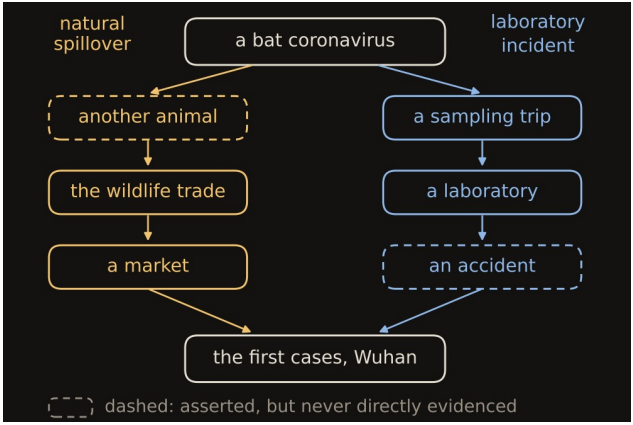


Figure F15-01. *Where the two accounts differ and where they agree. Both begin with a bat virus and end in Wuhan in late 2019; the dispute is entirely about the middle, and in both chains the middle is the part with least direct evidence.*

KEY CONCEPT — Three disputes wearing one coat

Almost everything written about this question is confusing because three separate arguments are being conducted in the same sentences. Separating them is the single most useful thing a reader can do.

The first is a scientific dispute about evidence. Where did the earliest cases cluster, what does the genetic diversity of the first sequences imply, what was in the

market, what was in the laboratory? These are factual questions with factual answers, and they are answerable in principle by anyone, regardless of politics.

The second is a policy dispute about risky research. Should laboratories collect dangerous viruses and modify them to see what they could become? That question does not depend on where this virus came from. If the answer is that such research is too dangerous, it was too dangerous in 2018, and a natural origin for Covid-19 would not make it safe. If the answer is that it is justified, a laboratory origin would not by itself settle the argument either, though it would weigh heavily on it.

The third is a dispute about trust. Did Chinese authorities conceal things? Did Western scientists and officials dismiss the laboratory hypothesis too early, and if so, why? Did a reasonable concern get treated as a conspiracy theory because of who was voicing it?

Mixing them makes every one of them worse. A person convinced that officials behaved badly in 2020 — and some did — can slide into treating that as evidence about where a virus came from, which it is not. A person convinced that the scientific evidence favours an animal origin — and it does — can slide into treating anyone who raises the laboratory question as acting in bad faith,

> which many are not. This chapter keeps them apart, and it is the only way to make the question tractable.

What was known in January 2020

The first recognised cases clustered around the Huanan Seafood Wholesale Market, which sold live mammals as well as fish. The market was closed and disinfected on 1 January 2020, which destroyed evidence that might have answered the question and which was, at the time, the obvious public health action to take.

Chapter 14 described the sequence being published on 11 January. What the sequence showed was a coronavirus most closely related to viruses found in horseshoe bats, with no known close relative in any database.

By the end of January 2020, scientists had noticed something about the spike protein. It carries a feature called a furin cleavage site, a short insertion at a particular junction which makes the spike easier to activate and the virus better at infecting human cells. It was not present in the closest known relatives of this virus, and its presence in a sarbecovirus was novel.

This single feature did more to shape the argument than any other. To one set of eyes it was a red flag: exactly the sort of thing a laboratory might add. To another it was unremarkable: furin cleavage sites occur across the coronavirus family, arise naturally, and if you were engineering one you would probably not build the particular version that appeared here.

In March 2020 a group of virologists published a paper in *Nature Medicine* arguing that the genome's features were inconsistent with deliberate engineering. That paper was influential at the time and has been heavily criticised since, both for its reasoning and because correspondence released later showed several of its authors expressing more uncertainty in private than the paper conveyed. Whatever is made of that, it is a dispute about how scientists behaved, which is the third kind. It is not by itself evidence about the virus.

The case for the market

The strongest version of the natural-origin case is not "most new viruses come from animals", true though that is. It is specific, quantitative, and rests on three independent lines of evidence that point at the same few hundred square metres.

The first is geography. A 2022 analysis mapped the residential addresses of 155 cases from December 2019. The cases

clustered around the market — and, crucially, so did the cases with no known connection to it. If the market were merely a place where an outbreak that started elsewhere happened to be noticed, the unconnected cases should have been spread across the city. They were not.

The second is the two lineages. The earliest sequences fall into two groups, called A and B, separated by two mutations. Analysis published in 2022 concluded that these represent at least two separate introductions of the virus into people, both around late November 2019. That is awkward for a single laboratory accident, which would be expected to seed one lineage. It is exactly what repeated exposure at a wildlife interface looks like. Both lineages have been detected in material from the market.

The third is what was in the building. Chinese researchers took environmental swabs from the market in early 2020 and eventually released the raw data. International groups analysing it found that the positive samples were concentrated in one corner — the corner where live wildlife was sold — and that samples from a particular stall which tested positive for the virus also contained the genetic material of raccoon dogs, civets and other species known to be susceptible. Huanan was the most active of only four markets in Wuhan selling live mammals.



Figure F15-02. *The corner of the market where the positive samples clustered. What was never done, and can no longer be done, is the obvious thing: test the animals themselves. They were gone before anyone looked.*

> **HOW WE KNOW — What a family tree can and cannot tell you**
>
> When a virus copies itself it makes occasional mistakes, and those mistakes are inherited. Compare enough genomes and you can reconstruct how they are related, and roughly when their common ancestor existed, because the mistakes accumulate at a roughly steady rate.

That is how the two-lineage argument works. If every early sequence descended from one introduction, they would form a single tree with one root. Instead the early sequences form two groups whose common ancestor sits further back than either, which is most simply explained by two introductions from the same animal source within a short period.

Here is what the method cannot do. It cannot tell you where the introduction happened, only when, and roughly how many times. It cannot distinguish an animal that infected two people from a laboratory worker who infected two people, if the laboratory samples were themselves genetically diverse. And the dating carries real uncertainty: an estimate of “late November 2019” is a range, not a date.

The honest summary is that the genetic evidence constrains the answer without determining it. It makes a single point-source introduction less likely, which counts against the simplest version of the laboratory account and in favour of the market. It does not exclude more complicated versions of either.

The counter-argument made by the laboratory hypothesis’s defenders is straightforward and cannot be dismissed. A mar-

ket is a place where a great many people from across a city gather in an enclosed space. An infected person — a laboratory worker, a contact of one — visiting that market could seed exactly the cluster that was observed, and the market would then amplify it, producing a pattern that looks like origin but is only superspreading. The spatial analysis has also been challenged on statistical grounds, with the original authors responding in turn.

What would have settled it is the animals. They were removed and the market disinfected before anyone sampled them.

The case for the laboratory

The strongest version of this case is also not the one usually attributed to it. It does not require that the virus was engineered as a weapon, and its serious proponents do not claim that.

It rests on four points.

The first is location. The Wuhan Institute of Virology held one of the largest collections of bat coronaviruses in the world and had spent years sending expeditions to caves in southern China to find more. The nearest known natural relatives of this virus are in bats over a thousand kilometres

from Wuhan. A pandemic virus of bat origin emerging first in the one city with a major bat coronavirus laboratory is, at minimum, a coincidence that requires accounting for.



Figure F15-03. *Where both accounts begin. The nearest known relatives of this virus live in bats more than a thousand kilometres from Wuhan, which is a fact both hypotheses have to accommodate.*

The second is the research itself. Laboratories including the one in Wuhan were doing work that involved constructing modified coronaviruses to assess their potential. This is not secret and was published at the time.

The third is the furin cleavage site, and here the case became considerably stronger in 2021 when a grant proposal became

public. In 2018 a consortium including EcoHealth Alliance, the Wuhan Institute and the University of North Carolina submitted a bid to the American defence research agency DARPA, called DEFUSE, which among other things proposed inserting furin cleavage sites into novel bat sarbecoviruses. DARPA declined to fund it. But the proposal establishes that researchers in this collaboration were, in 2018, specifically contemplating putting into a bat sarbecovirus the very feature that appeared in SARS-CoV-2 and in none of its known relatives. In congressional testimony the North Carolina researcher confirmed his laboratory would have performed that insertion.

That is a genuinely striking coincidence and it should be stated as such rather than minimised.

The fourth is the absence on the other side. Six years on, and despite extensive searching, no infected intermediate animal has been found. For SARS in 2003 the civet link was established within months; for MERS, camels. Here there is nothing.



Figure F15-04. *The records that would test this hypothesis — staff health records, biosafety inspection reports, laboratory audits — have been requested since November 2021 and have not been provided.*

The counter-arguments are equally real. No evidence has ever been produced that the Wuhan Institute held a virus close enough to SARS-CoV-2 to be its progenitor, and a laboratory cannot leak what it does not have. The DEFUSE proposal was rejected, and no evidence has emerged that the work was done. The particular furin cleavage site in SARS-CoV-2 is, several virologists argue, not the one a designer would have chosen. And the coincidence of location is weaker than it sounds: laboratories studying bat coronaviruses are concentrated in exactly the regions where bat coronaviruses

are, which is where such an outbreak would be most likely anyway.

What the WHO group concluded

In November 2021 the World Health Organization convened the Scientific Advisory Group for the Origins of Novel Pathogens, twenty-seven scientists from twenty-seven countries, chaired by the South African virologist Marietjie Venter. It worked for more than three years and reported on 27 June 2025.

Its conclusion was that the weight of available evidence suggests zoonotic spillover, either directly from bats or through an intermediate host.

Its second conclusion was that a laboratory-related incident could not be investigated or excluded, because the information needed to do so had not been provided. The group described the evidence advanced for the laboratory scenario as remaining speculation. It also recorded that it had been unable to see relevant American and German intelligence reporting.

The group did not agree with itself. In the week the report was published one member resigned and three others asked for their names to be removed, which is why the chapter can

say that twenty-three of the twenty-seven stood behind it. The Director-General, publishing it, said that all hypotheses must remain on the table.

In February 2026 those twenty-three published a commentary in *Nature* setting out where things stood. Most of the peer-reviewed scientific evidence, they wrote, supports an animal origin, though it does not establish when or where the spillover happened. They found no convincing genetic evidence of deliberate manipulation and no concrete proof of a laboratory breach. And they stated, in their own words, that the Chinese government had not provided the information they had requested since November 2021 in order to investigate a potential biosafety breach: health records of laboratory staff, biosafety inspection reports, and audit documents from the relevant facilities.

That is the position of the body with the broadest access and the clearest mandate, and it should be read carefully, because it says two things at once. The scientific evidence that exists leans towards an animal origin. The evidence that would test the alternative has been withheld, and the withholding is a fact about the Chinese government rather than about the virus.

What intelligence agencies could and could not say

Several Western intelligence services have reached the opposite conclusion from SAGO, and their reasoning is not public.

In January 2025 the CIA stated that it assessed, with low confidence, that a research-related origin was more likely than a natural one. The assessment was commissioned under the previous administration and was explicitly not based on new intelligence; it was a re-reading of existing material, and the agency said both scenarios remained plausible. The FBI had previously reached a similar conclusion with moderate confidence, and the Department of Energy with low confidence. Other American agencies favour a natural origin, several also with low confidence. In March 2025 German newspapers reported that Germany's foreign intelligence service had assessed the probability of a laboratory origin at eighty to ninety-five per cent, in an analysis completed around 2020 and not made public at the time. A congressional select subcommittee concluded in a long report that a laboratory origin was most likely.

THE NUMBER — What “low confidence” means

These phrases are terms of art and they are routinely reported as though they were not.

In the American intelligence system, a confidence level describes the quality of the evidence and reasoning behind a judgement, not the probability of the judgement being true. High confidence means multiple high-quality sources agreeing. Moderate confidence means credible sourcing that is not corroborated enough to be more. Low confidence means the information is scanty, questionable, fragmented or so plausibly alternative-explained that the agency has real doubts.

So “the CIA assesses with low confidence that a research-related origin is more likely” is not a statement that a laboratory origin is likely. It is a statement that, on the evidence the agency has, it leans that way, and that it does not think much of the evidence. An agency saying this is explicitly inviting you to hold the conclusion loosely.

Two further things these figures are not. They are not independent votes. Agencies read overlapping material, and when one changes its view the others’ positions are not thereby confirmed. And they are not scientific as-

sessments: intelligence agencies assess human reporting, intercepted communications and patterns of behaviour, which answer questions like “did someone conceal something” far better than “where did a virus come from”.

The German figure of eighty to ninety-five per cent is of a different kind again — a probability estimate from a service that does not use the American confidence vocabulary, in an assessment whose underlying material has not been published. A number without its workings is not more informative than a word, only more quotable.

What would settle it

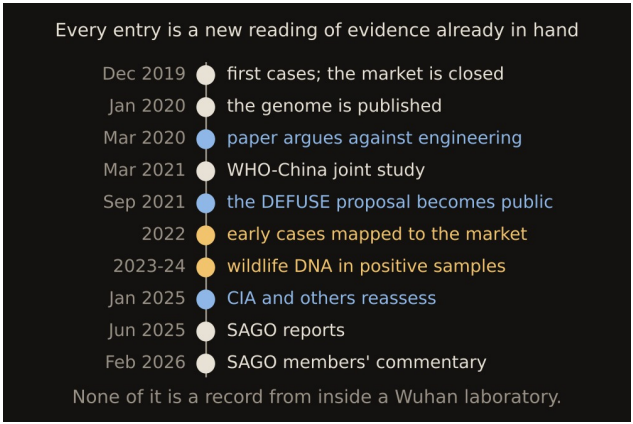


Figure F15-05. *Six years of findings, and the one thing none of them is. Almost every entry is an analysis of material already in hand; the items that would decide the question are records nobody outside China has seen.*

It is worth being concrete about this, because a question that cannot be settled and a question that has not been settled are different things.

A laboratory origin would be effectively established by any of: a record of a virus closely matching SARS-CoV-2 in a Wuhan collection before December 2019; evidence of the DEFUSE work having been carried out; or documented ill-

ness in laboratory staff in autumn 2019 with retained samples confirming it. All three would be in records that exist or existed.

A natural origin would be effectively established by finding the progenitor virus in an animal population, or by finding archived samples from the wildlife supply chain feeding the market that contain it. Both are searches that can still be done, and the Chinese authorities did sample farms supplying the market; the results have not been released in full.

What makes the question unusual is that it is not short of evidence in principle. It is short of access. The market was cleared and disinfected within days. The laboratory records have been withheld for five years. The farm sampling results are unpublished. Several of these would likely have resolved the matter, and the fact that none of them is available is itself a fact — though it is a fact about an institution's behaviour, and institutions conceal things for many reasons, including embarrassment about matters unrelated to the question being asked.

Where this leaves us

Here is the summary this chapter is prepared to defend.

The peer-reviewed scientific evidence leans towards a natural origin, and it is not a small lean: the clustering of early cases, the two independent introductions, and the wildlife DNA in the positive market samples are three distinct lines of evidence pointing at one place. The body that reviewed all of it over three years concluded that the weight of evidence suggests zoonotic spillover.

The laboratory hypothesis has not been refuted, remains entirely plausible, and is favoured by several intelligence services, all of them at low or moderate confidence and none of them showing their working. No intermediate animal has ever been found. The coincidence of the DEFUSE proposal and the furin cleavage site is unexplained.

The evidence that would distinguish the two has been requested for five years and not supplied.

Anyone telling you this question is settled, in either direction, is telling you something the available evidence does not support. Anyone telling you it is unknowable is going too far the other way: it is knowable, the information exists, and it is being withheld.

What the question does not need is for the answer to matter for what comes next. Chapter 12 established that spillovers happen constantly and that most pandemics in history began

that way; chapter 20 is about one that may be starting now. Laboratory accidents are also real, have happened repeatedly with other pathogens, and the policy argument about dangerous research stands on its own evidence. Both risks need managing, and neither case for doing so depends on how this particular virus got into this particular city.

The next chapter turns from where the virus came from to what was done about it, and it is a harder chapter to write, because the evidence there is abundant rather than withheld, and people disagree about it anyway.

The Year Indoors

> *“Too little, too late.” — Baroness Hallett, summarising the UK Covid-19 Inquiry’s findings, 20 November 2025*

Written in November 2026.

The week of 16 March 2020

By the middle of March 2020 the governments of western Europe were all looking at the same problem and none of them had a plan for it.

The plans they did have were pandemic influenza plans. Those assumed a virus that would arrive, move through the population over a few months, be partly slowed by school closures and hand hygiene, and be met by a vaccine within six months because influenza vaccines can be remade from an existing process. The plans did not contain the word lockdown. Nobody had modelled one, costed one, or written down how one would end.

What they were now facing was the thing chapter 14 set out: a virus spreading from people who felt well, against which nobody had any immunity, doubling every three or four days, in a continent whose intensive care capacity was already visible in Bergamo.

The sequence across Europe took about a fortnight. Italy closed first, then Spain, then France. Ireland closed its schools on 12 March and ordered people to stay at home on 27 March. The United Kingdom announced its lockdown on 23 March, days after most of its neighbours.

It is worth registering how little precedent there was for any of this. No western European government had confined its entire population to their homes in peacetime. There was no legal machinery for it in most countries, no modelling of what it would do to an economy, no plan for schools, no agreed criterion for when it would end, and no example to copy except China's, which was widely assumed to be both unreliable in its reporting and impossible in a democracy. Every government in Europe improvised the largest peacetime restriction of movement in its history inside a fortnight, and then had to keep improvising for two years.



Figure F16-01. *What a lockdown looked like from the outside. The thing being bought with this was time: time for hospitals not to be overwhelmed, and, though nobody could be sure of it in March 2020, time for the vaccine in chapter 17 to arrive.*

What the models said

On 16 March 2020 a team at Imperial College London published a report that changed what several governments did that week. Its central estimate was that in an unmitigated epidemic — one in which nothing at all was done — approximately 510,000 people would die in Great Britain and 2.2 million in the United States, and it noted that this did not ac-

count for the additional deaths that would follow from health systems being overwhelmed.

That number did the work. Whatever is now thought of the model, no government reading it could have done nothing, and the UK Inquiry's eventual finding was that the governments of the UK, in imposing a mandatory lockdown, acted in the genuine and reasonable belief that it was required.

KEY CONCEPT — What a model is for

A model of an epidemic is not a prediction. It is a sentence of the form: if these assumptions hold and these things are done, this is roughly what follows.

The assumptions are the whole of it. How infectious is the virus? How long between catching it and passing it on? What fraction of infected people die, by age? How many contacts does an average person have, and how much does closing schools change that? In March 2020 every one of those was uncertain, several were unknown, and the model's output inherits all of that uncertainty and then some.

So the correct reading of "510,000 deaths" was never "510,000 people will die". It was "on our current best guesses, doing nothing leads to a number in this region,

which is intolerable". That is a useful sentence. It is the sentence that justifies acting.

Two failure modes follow, and both happened. The first is treating a scenario as a forecast, which governments and newspapers did, and which makes a model look wrong when the thing it warned of does not happen — even though the thing not happening is usually because people acted on the warning. The second is the reverse: dismissing all modelling because one scenario did not come to pass, which discards the only tool available for thinking about what has not happened yet.

A model can be useful and wrong at the same time, and most of them are. The question to ask of one is not "was the number right?" but "were the assumptions reasonable given what was known, and did the conclusion survive changing them?"

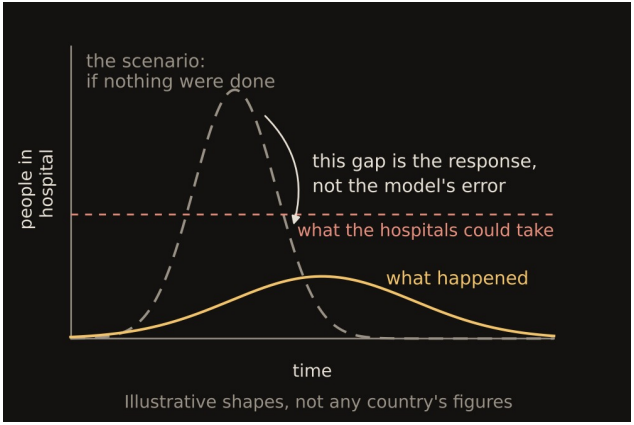


Figure F16-02. *Why a model that warns of something which then does not happen has not necessarily failed. The warning and the response are not independent, and that makes models of this kind almost impossible to score after the fact.*

There is a real criticism to make of the modelling, separate from this. Where a model's assumptions can be tested against a country that did not do the thing, it should be. The same group's work projected around 35,000 first-wave deaths for Sweden under the scenario closest to what Sweden actually did. Sweden had fewer than 6,000. That is a large miss on the one available test, and it deserves to be stated as plainly as the 510,000 figure.

Suppression or mitigation

The choice in March 2020 was usually described as between two strategies.

Mitigation meant slowing the epidemic without stopping it: protect the vulnerable, flatten the peak so that hospitals cope, and accept that most of the population would be infected on the way to immunity. Suppression meant driving transmission down hard enough that case numbers fell, and holding them there.

The argument for mitigation was that suppression had no end. You cannot hold a population indoors indefinitely, and the moment you stop, the epidemic resumes in a population still susceptible. On this reasoning suppression merely postpones the same deaths at enormous cost.

The argument for suppression was that postponing was the point, for two reasons. Hospitals that are not overwhelmed save people that overwhelmed hospitals cannot, so the same number of infections produces fewer deaths when spread out. And the assumption that nothing would change was wrong: treatment improved rapidly through 2020, and a vaccine arrived in December.

That second point is the one that settled the argument in retrospect, and it is important to notice that in March 2020

nobody could rely on it. No coronavirus vaccine had ever been licensed. Deciding to suppress was, in part, a bet that something would turn up.

The phrase that carried the policy was “flatten the curve”, and it is worth examining because it was both the clearest piece of public health communication of the pandemic and quietly misleading. The picture showed two epidemic curves of equal area, one tall and narrow and one low and wide, with a horizontal line marking hospital capacity. The message was that the same number of infections spread over longer would stay under the line.

Equal area was the problem. The picture said, in effect, that the same people would be infected either way, and that the benefit of restrictions was only in the timing. That was the mitigation argument drawn as a diagram, and it was what governments showed the public while actually pursuing suppression, whose entire premise is that the areas are not equal. A great deal of later confusion about what lockdowns were supposed to achieve traces back to a population having been shown one strategy and governed by another.

Sweden

Sweden did something closer to mitigation than its neighbours and became the most argued-about country in Europe.

What it actually did is usually overstated in both directions. Sweden did not carry on as normal: gatherings were limited, universities and upper secondary schools moved online, visits to care homes were banned, and a great deal of voluntary distancing happened. What it did not do was close primary schools, order people to stay at home, or make most of its measures legally binding.

Then the argument about what happened began, and it has not stopped, because both sides are quoting real numbers.

> **HOW WE KNOW — Why Sweden is read two ways**
>
> A study published in 2024 compared the four largest Nordic countries using the same method throughout: weekly all-cause deaths from national registers, with 2010 to 2019 as the baseline, adjusted for population size, age and sex. Using one consistent method on four similar countries is about as good as this kind of comparison gets.
>
> For 2020, Sweden had 75 excess deaths per 100,000 people. Denmark, Finland and Norway had almost none. In the first wave specifically, Sweden's per-capita Covid death rate was something like ten times Finland's and Norway's and five times Denmark's.

For 2022, the order reverses. Sweden had 25 excess deaths per 100,000, the lowest of the four; Finland had 130, Norway 88, Denmark 52.

Across the whole three years the four are hard to separate: Denmark 86, Sweden 117, Norway 121, Finland 190 per 100,000 — and the uncertainty ranges on those are so wide that they overlap almost entirely.

So: did Sweden do badly or well? In 2020, clearly worse than its neighbours, and that cost several thousand lives that Denmark, Norway and Finland did not lose. By 2022, clearly better. Over three years, indistinguishable from Denmark and Norway and better than Finland, with intervals too wide to call.

Every one of those statements is true and each has been quoted on its own as the verdict. The choice of window decides the answer, which means anyone who gives you a single number for Sweden has chosen a window, and you should ask which one and why.

Three cautions the authors themselves raise. Excess mortality depends entirely on the baseline you expect; a mild flu season before 2020 lowers the baseline and inflates the apparent excess. Deaths deferred from one year appear as a deficit later, which flatters whichever

country had its deaths early. And none of this isolates the effect of policy, because the four countries differ in population density, household structure, care home organisation and how the virus arrived.

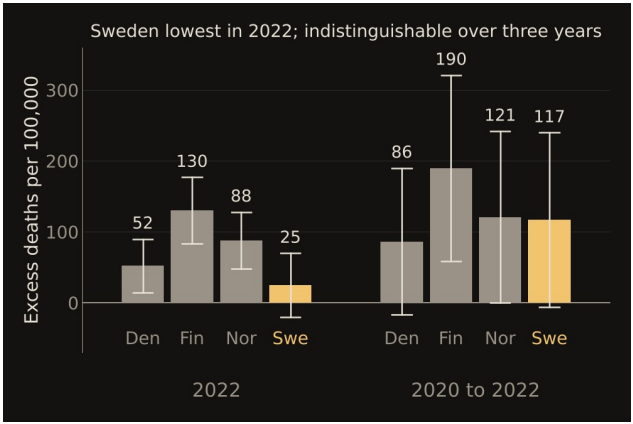


Figure F16-03. *The same four countries, three windows, three different answers. The intervals on the three-year totals are drawn because without them the ranking looks meaningful, and it is not. Drawn from published Nordic register data.*

What can be said with reasonable confidence is narrower than either camp wants. Sweden's 2020 was worse than its neighbours' and its care homes were hit hard. Sweden's approach did not produce the catastrophe some models projec-

ted. And across three years the Nordic countries ended up in much the same place by very different routes, which is an argument for humility rather than for either side's conclusion.

The nursing homes

Wherever you look, the same thing happened, and it is the part of the pandemic that has been least examined relative to its size.

In Ireland, about sixty per cent of everyone who died in the first wave died in a nursing home. Nursing home residents are about five per cent of the population.

That concentration was not inevitable and it was not purely a consequence of age. An Irish analysis presented to the national evaluation found that mortality among over-seventies living in nursing homes was twenty-one times higher than among over-seventies who were cocooning at home — and that for those cocooning at home, mortality was no higher than in the general population. The same work found that fewer than half of Irish long-term residential care facilities had any infection prevention and control training before the pandemic.



Figure F16-04. *The policy almost every country adopted. Visiting bans protected residents from infections brought in by families, and did not protect them from infections brought in by staff, which is how the virus mostly arrived.*

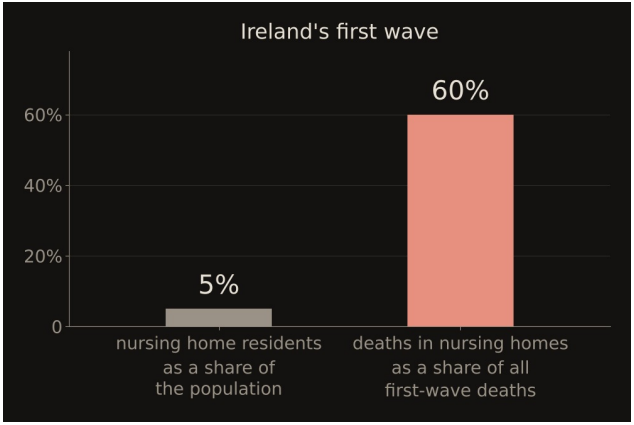


Figure F16-05. *Five per cent of the people, sixty per cent of the deaths. This is the single starkest number in the Irish pandemic, and it concerns a population that was locked down harder than anyone else and died anyway.*

There is a bitter logic here worth stating plainly. The general lockdown was justified largely by the need to protect the most vulnerable. The most vulnerable were overwhelmingly in care homes. Care homes were where the protection failed worst. Whatever else the policy achieved, it did not achieve that.

Ireland’s long lockdown

Ireland’s first stay-at-home order came on 27 March 2020, and Ireland’s restrictions turned out to be among the longest-lasting in Europe, particularly for hospitality and retail. Irish pubs were closed longer than any others on the continent.

Ireland also had, through 2020 and much of 2021, one of the lower death rates in western Europe. Both of those facts are usually deployed as arguments, and both are true at once.

The Irish decision-making structure was distinctive. A National Public Health Emergency Team, chaired by the Chief Medical Officer, met in private and gave advice to the Minister for Health, which went to government for decision. That arrangement concentrated expertise and moved quickly. It also meant that the reasoning behind advice was not public at the time, and that the division between advice and decision became blurred in public understanding, which is a problem the Irish evaluation has heard about repeatedly.

That evaluation is the Covid-19 Expert Evaluation, established by the government in late 2024 and chaired by Professor Anne Scott. It is independent and non-statutory, it has held public consultations and public hearings through 2026, and it is due to report by the end of this year. It has not reported at the time of writing, which is a limitation of this

chapter and will be a limitation of this book if it is read after the report appears.

What has emerged publicly from its hearings is not comfortable. At a forum in 2026 the chair summarised it by saying that six years on, Ireland is probably no better prepared than it was in 2020. The gaps identified were the ones the nursing home figures imply: infection control in long-term residential care, public health structures, and coordination.

Counting the costs

The costs of what was done fall into categories that are measured with very different degrees of confidence, and honesty requires saying which is which.

Some are well measured. Cancer diagnoses fell sharply in 2020 across every comparable country, because screening programmes paused and people did not present. Nordic registry work covering 2020 and 2021 found declines in early-stage breast cancer diagnoses in Denmark, Norway and Sweden with partial catch-up the following year, and declines in prostate cancer rates across four countries. A diagnosis delayed is not a diagnosis avoided; it is the same cancer found later.

Some are real but harder to size. Children lost months of schooling, unevenly, with the losses falling hardest on those who had least at home — the Irish evaluation heard this put as the psychological experience of restrictions being entirely different for someone with secure employment and adequate housing than for someone in poverty and overcrowding. Mental health deteriorated, though disentangling the effect of the restrictions from the effect of the pandemic itself has proved very difficult.

Some are genuinely contested. The economic effects, the long-run effects on children, and the effects on social trust are all argued about by serious people using the same data.

The education loss deserves a paragraph of its own, because it is the cost most likely to be large and is the one least visible in any mortality figure. A child who missed a term of school did not experience an event that ended when the school reopened. Learning is cumulative; a gap in the foundations shows up years later in what can be built on them. And the loss was not distributed like the virus. A child with a quiet room, a laptop, broadband and a parent working from home lost comparatively little. A child sharing a room, with one phone between a family and a parent who had to leave the house to work, lost a great deal. The pandemic's mortality fell hardest on the old; this cost fell hardest on the young

and poor, and it will be paid over decades rather than months.

THE NUMBER — Excess deaths, and what it hides

Counting Covid deaths directly is unreliable: it depends on testing, on how death certificates are completed, and on whether a country counts deaths with the virus or from it. Excess mortality avoids most of that. You take the deaths that actually occurred and subtract the number you would have expected from previous years. What is left is the net effect of everything the pandemic did.

Everything is the important word. Excess mortality includes deaths from Covid, deaths from treatment people did not get, deaths from care that was delayed, and deaths from any other consequence of the response. It also subtracts the deaths that did not happen: fewer road deaths, fewer deaths from seasonal influenza, which was almost absent in 2020 and 2021.

So excess mortality is the right measure for “what did the whole thing cost in lives”, and it is the wrong measure for “how many people did the virus kill” or “how many did the lockdown kill”, because it cannot separate them. Anyone using a single excess mortality figure to prove that restrictions killed people, or to prove that they

> saved people, is using a number that by construction cannot answer that question.
>
> What it can do is bound the argument. A country whose total excess mortality over three years is close to zero did not suffer a hidden catastrophe from either the virus or the response.

What the inquiries found

The United Kingdom's inquiry is the most thorough examination any country has conducted. It has been running since 2022, is expected to cost the British taxpayer up to £200 million, and will produce ten reports. Its second, on government decision-making, was published on 20 November 2025 and runs to about 800 pages.

Its chair, Baroness Hallett, summarised the response as too little, too late. The report made twenty-two key findings and nineteen recommendations, described a toxic and chaotic culture at the centre of government, and concluded that had a mandatory lockdown been imposed a week earlier than 23 March 2020, deaths in England in the first wave up to 1 July would have been lower by about forty-eight per cent, or roughly 23,000 people.

Two things about that finding are consistently left out when it is quoted.

The first is what the report says alongside it. Hallett did not conclude that lockdowns should be the default. She concluded the opposite: that if earlier and better action had been taken, a lockdown might have been avoided altogether, and that an earlier lockdown could have been a shorter one. The finding is about the cost of indecision, not about the merit of confinement.

The second concerns the number itself. In February 2026 the Office for Statistics Regulation wrote to the inquiry to say that it had conveyed too great a level of certainty in presenting the modelling. The executive summary had said that modelling "has established" the 48 per cent figure. The inquiry accepted the criticism: it agreed that "established" was too definitive, that it did not reflect the uncertainty inherent in modelling, and that the executive summary had not sufficiently expressed that uncertainty — while noting that the main body of the report had used the more cautious "could have", and maintaining that the substantive finding stands.

That exchange is worth more than the headline. A public inquiry stated a modelled result too firmly in its summary, a statistics regulator said so, and the inquiry agreed. That is a system working, and it is also a demonstration of how a

number acquires a certainty in transmission that its authors never claimed for it.



Figure F16-06. *The cost that is hardest to measure and most certain to be real. Schooling lost in 2020 and 2021 was not evenly lost, and the children who lost most had least at home to replace it with.*

Where this leaves us

Here is what this chapter is prepared to defend.

That the restrictions reduced transmission is not seriously disputed by anyone working with the data; the argument is about magnitude, duration and cost, not about direction. That governments acted on the best information they had in

March 2020 is the finding of the most thorough inquiry conducted. That several of them, including the British, were slower than their own advice and that the delay cost lives is also the finding of that inquiry.

That the measures were held in place longer than the evidence supported in some countries, Ireland among them for hospitality, is probable and is what the Irish evaluation appears to be moving towards, though it has not reported.

That the costs were real, large, unequally distributed, and insufficiently weighed at the time is, as far as this book can tell, simply true, and is the thing most likely to be judged harshly.

And that the care homes were the central failure is the conclusion the numbers force, in Ireland and almost everywhere else. The policy's stated purpose was to protect the most vulnerable. The most vulnerable died anyway, in the places built to look after them, at twenty-one times the rate of people the same age at home.

What this chapter will not do is tell you whether it was worth it, because that is not a question evidence can answer on its own. It requires deciding how to weigh a death against a year of a child's schooling, a bankruptcy, a cancelled funeral, and a cancer found too late, and reasonable people weigh

those differently. What evidence can do is make sure the weighing is done with the right numbers.

The next chapter is about the thing that ended it, and it is the one unambiguously good story in this part of the book.

Three Hundred Days

> *“It’s the best early birthday present I could wish for.” — Margaret Keenan, Coventry, 8 December 2020*

Written in November 2026.

The eleventh of January

On 11 January 2020 the genetic sequence of the new coronavirus was posted to a public database. Chapter 14 described that moment from the point of view of the people trying to work out what the virus was. This chapter describes it from the point of view of the people trying to build something against it, and from that angle the posting of the sequence was not information. It was a starting pistol.

Within about twenty-four hours, a team at the Vaccine Research Center in Maryland had taken the part of the sequence that codes for the spike protein and made two small changes to it. Two days after that, on 13 January, the same team and a

company in Massachusetts called Moderna had settled on the final design of a vaccine.

Nobody had been to a laboratory bench. Nothing had been grown, purified or tested. The design existed only as a string of letters in a file. That is worth sitting with, because almost everything else in this book has involved somebody going out and looking at something — Jenner at a milkmaid's hand, Burkitt in a Ford station wagon, Montagnier's team under a microscope. Here, the thing that was going to be put into several billion arms was specified in a text file in two days, by people who had not yet seen the virus.



Figure F17-01. *Empty vials on a filling line. The design took two days. Everything after it took ten months.*

The reason this was possible is the subject of the next section, and it is not that anybody worked very fast in January 2020. It is that the work had already been done.

What was already on the shelf

A virus gets into a cell by fusing with it, and a fusion protein is a loaded spring. Before it does its job it is held in one shape, tense and awkward; afterwards it has snapped into another, which is stable and quite different. The problem for vaccine makers is that the immune system needs to recog-

nise the first shape, the one it will actually meet on an incoming virus, and the first shape is the unstable one. Purify the protein and it springs. You end up immunising people against the sprung version, which is the version that no longer matters.

In the 2000s a group working on respiratory syncytial virus solved this for RSV by finding mutations that pinned the protein in its unsprung shape. An earlier vaccine against RSV had failed badly in the 1960s, which is why the problem had anyone's attention at all. The same group then went looking to see whether the trick generalised, and found that it did. Substituting two prolines at a particular point in the stalk of a coronavirus spike locked it in the shape they wanted. It worked on the MERS coronavirus, on the original SARS coronavirus, and on one of the coronaviruses that causes common colds.

> **KEY CONCEPT — Two prolines**
>
> Proline is an amino acid with an unusual property: its side chain loops back and joins the protein's own backbone, which makes it rigid. Put a proline where a protein needs to flex and it cannot flex there any more.
>
> The spike protein has a hinge. After the spike grabs a cell, that hinge folds and the whole protein collapses into

> its post-fusion shape. Put two prolines in the hinge and the collapse cannot happen. The spike stays in the shape the immune system needs to learn.
>
> The modification is called 2P. It is in the Moderna vaccine, in the Pfizer-BioNTech vaccine, and in several others. It is one of the quiet reasons the vaccines worked as well as they did, and it was worked out years before anyone had heard of SARS-CoV-2, on a different virus, for a different reason.

This is what people mean by preparedness, and it is a less exciting thing than the word suggests. It does not look like a stockpile. It looks like a handful of researchers spending a decade on the structural biology of a family of viruses that, at the time, caused colds and two frightening but contained outbreaks. Nobody could have told you in 2016 what that work was for. In January 2020 it turned out to be for this.

Moderna, meanwhile, had spent years failing to get mRNA to do anything commercially useful. Chapter 13 told that story. What matters here is the residue of it: by January 2020 the company had already taken nine other infectious disease vaccines into clinical trials, including ones against two pandemic influenza strains and three other respiratory viruses. It had a manufacturing process that it had run many times. The

coronavirus vaccine was its tenth. It was not a leap into the unknown; it was the tenth time down a path they had already worn.

Sixty-three days

The sequence was chosen on 13 January. The first clinical batch of vaccine — actual vials, filled and finished — was completed on 7 February, twenty-five days later. It was shipped to the National Institutes of Health on 24 February. On 16 March the first participant in a phase 1 trial was injected.

Sixty-three days from a decision on a text file to a needle in an arm.

For comparison, the fastest vaccine development in history before this was mumps, which took four years. The usual figure quoted for a vaccine from scratch is ten to fifteen. The Covid vaccines were authorised for emergency use in December 2020, about three hundred and thirty days after the sequence was posted.

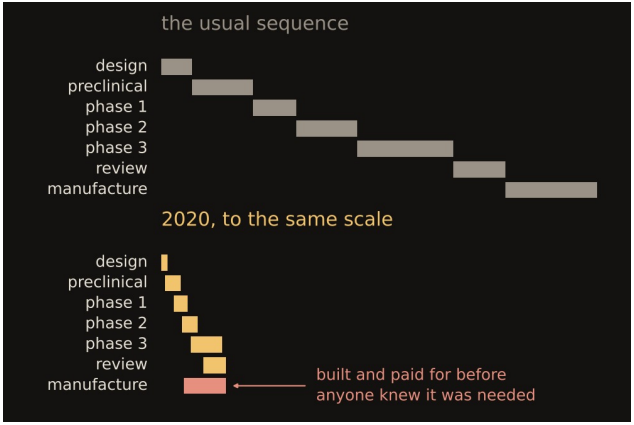


Figure F17-02. *How the time was found. The stages were not shortened; they were overlapped, and someone paid for the ones that might have been wasted.*

The obvious question is what was skipped, and the answer, as far as the record shows, is nothing. The phase 1 trial ran. The phase 2 trial ran. The phase 3 trial ran, and was larger than most phase 3 vaccine trials in history. The regulators reviewed the data.

What changed was the shape of the process rather than its content. Normally the stages run one after another, because each one is expensive and nobody spends the money on the next until the last has worked. Normally manufacturing be-

gins after approval, because building a factory for a product that might fail is a good way to lose a great deal of money. Normally the regulator receives a completed dossier and begins reading.

In 2020, phases overlapped. Manufacturing began while the trials were still running, at scale, for vaccines that might have turned out to be useless — and some did. Regulators read data as it arrived rather than at the end, a procedure called rolling review. And the trials themselves recruited faster than any trials ever had, because tens of thousands of people volunteered in a matter of weeks for something that would normally take a year to enrol.

Every one of those changes costs money or has a cost in risk. None of them costs safety data. That distinction is the single most important thing to be clear about in this chapter, and it is the thing that was least well explained at the time.

Paying for failure in advance

The money is worth being specific about, because the speed was bought.

Governments and foundations signed contracts that paid manufacturers whether or not the vaccine worked. They funded factory capacity for candidates still in phase 1. They

placed advance orders for hundreds of millions of doses of products that did not exist and might never exist. From a narrow accounting view this was a scandalous way to spend public money, and in the ordinary run of things it would be. Several of the candidates that were paid for did fail, and the money spent on them was lost.

That was the point. The reason vaccines normally take a decade is not that the science takes a decade. It is that the risk is carried by whoever is developing them, and a rational developer manages risk by doing one thing at a time and stopping when it stops working. Move the risk somewhere else — onto a state, which can absorb it — and the sequence collapses into something much shorter.

It is a useful way to think about the whole year. The question "how did they do it so fast?" has a technical answer, which is the 2P substitution and twenty years of mRNA work, and a financial answer, which is that somebody agreed to pay for all the failures. Both are true. The second is the one that could be repeated tomorrow for some other disease, if anyone decided to.

How you prove a vaccine works

The trial that produced the first result was run by Pfizer and BioNTech. Between 27 July and 14 November 2020 it

screened 44,820 people and randomised 43,548 of them at 152 sites around the world. Half got two 30-microgram doses of the vaccine three weeks apart. Half got salt water.

Then everybody went back to their lives, and the trial waited.

HOW WE KNOW — The trial turned on 170 people

Forty-three thousand people were enrolled. The result rested on 170 of them.

A vaccine trial does not measure anything about the 43,000 who stay well. It counts the people who get ill and asks which arm of the trial they were in. If the vaccine does nothing, the cases should split roughly evenly. The further from even, the better the vaccine.

At the point of analysis there were 170 cases of Covid-19 with onset at least a week after the second dose. Eight were in the vaccinated group. One hundred and sixty-two were in the placebo group.

That split corresponds to 95 per cent efficacy, with a confidence interval from 90.3 to 97.6 per cent. There were also ten cases of severe Covid after the first dose: nine in the placebo group, one in the vaccinated group.

> This is why trials have to be large and why they have to wait. The size is not there to accumulate 43,000 measurements. It is there so that enough infections occur, in a reasonable time, to make the split mean something. In a place with little transmission the same trial would have taken years. One of the grim conveniences of the autumn of 2020 was that there was a great deal of transmission.

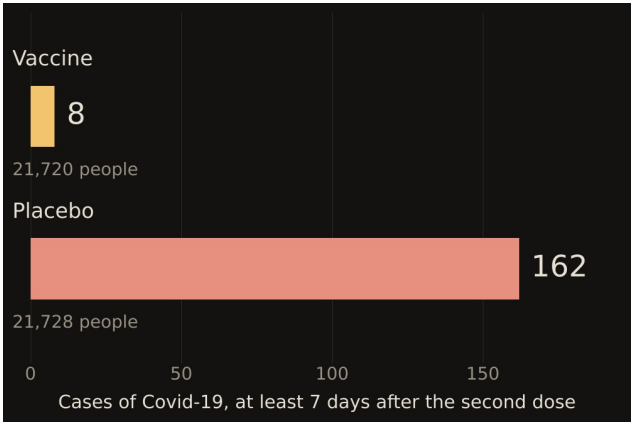


Figure F17-03. *The whole result. Eight cases against a hundred and sixty-two.*

Ninety-five per cent was far better than anyone had expected. Regulators had said in advance that they would accept

50 per cent, which is roughly where a good influenza vaccine sits in a good year. Several of the scientists involved said afterwards that they had been hoping for 60 or 70. The figure that came back was high enough that the first reaction in several places was to check the arithmetic.

It is worth noting what that 95 per cent was measuring, because it was widely misread later. It was efficacy against symptomatic, confirmed Covid-19, at least seven days after the second dose, in a population with no evidence of prior infection, measured over a median of two months, against the virus circulating in the second half of 2020. It was not a claim about transmission, which the trial did not measure. It was not a claim about how long protection would last, which the trial could not yet know. And it was not a claim about variants that did not exist yet. Chapter 19 returns to what happened to that number.

The ninth of November

Pfizer and BioNTech announced the result on 9 November 2020. Moderna announced a comparable result a week later. The Oxford-AstraZeneca vaccine, a different design using a modified chimpanzee adenovirus rather than mRNA, reported shortly after that.

Three hundred and three days had passed since the sequence was posted.



Figure F17-04. *Minus seventy. The Pfizer vaccine's cold chain was the hardest logistical problem of the rollout and the main reason the Oxford vaccine mattered.*

Then came the part nobody had rehearsed. The Pfizer vaccine had to be kept at around minus seventy degrees, which is colder than any ordinary medical freezer, and it came in trays of nearly a thousand doses that, once thawed, had to be used within days. A product with those properties is straightforward to deliver to a large teaching hospital in a rich country and close to impossible to deliver to a rural clinic without

reliable electricity. The Oxford vaccine, which kept in an ordinary fridge, was the one designed for everywhere else, and it was sold at cost. Its later difficulties, which chapter 18 deals with, have to be read against that.

Coventry, and then James's Street

The United Kingdom authorised the Pfizer vaccine first, and on the morning of 8 December 2020 a ninety-year-old woman named Margaret Keenan was injected at University Hospital in Coventry, a week before her ninety-first birthday. She was the first person in the world to receive a Covid-19 vaccine outside a clinical trial. She had been born in Enniskillen, in County Fermanagh, and had lived in Coventry for more than sixty years. The nurse who gave her the injection, May Parsons, had come to Britain from the Philippines.

Three people stood in that photograph, and between them they were a reasonable summary of how health services in these islands actually work.

Three weeks later, on 29 December, Ireland began. The first person vaccinated in the Republic was Annie Lynch, seventy-nine, a grandmother of ten, at St James's Hospital in Dublin.



Figure F17-05. *A vaccination centre in a hall. Within months the exercise had moved out of hospitals and into whatever large rooms were available.*

What followed in Ireland was, by most measures, one of the better national rollouts in Europe. Uptake was very high — persistently among the highest in the European Union — which is a striking thing for a country that, as chapter 10 described, had watched its own measles coverage fall below the threshold a few years earlier. The same population behaved quite differently about the same kind of intervention within a decade. Chapter 18 tries to explain why, because the explanation is not that the Irish became more trusting of vaccines in general.

The sequencing was argued about, as it was everywhere. Ireland vaccinated nursing home residents and staff first, which chapter 16 suggests was the right call and was in any case the call almost every country made. After that the arguments started: age bands against occupational groups, teachers against the clinically vulnerable, the long dispute about whether to delay second doses in order to give more people a first. The United Kingdom made that last decision on 30 December 2020, against the dosing schedule the trials had used, and was criticised for it at the time by people with good reasons. It appears in retrospect to have been correct. It was also, strictly, a decision to give millions of people a regimen that had not been trialled, taken on the basis of a judgement about the shape of the immune response, in the middle of a surge.

Not every country followed, and the ones that did not had reasonable grounds for staying with the schedule the trial had used. Neighbouring health services, looking at the same data in the same week, reached different answers, and more than one of those answers was defensible. That is worth holding on to whenever a pandemic decision is described afterwards as having been obvious. Almost none of them were obvious at the time. They were judgements made quickly with incomplete information by people who knew they

would be judged slowly afterwards with complete information, which is close to the definition of an unfair position and is also, unavoidably, the job.

Where the speed came from

It is worth gathering the answer in one place, because “the vaccine was rushed” became a settled belief for a substantial minority of people and the honest reply is more interesting than a denial.

Four things were genuinely different. The design step was effectively instant, because the platform was already built and the spike modification was already known. The stages overlapped instead of queueing. Manufacturing was paid for before anyone knew whether it would be needed. And the trials recruited at a speed that no trial has managed before or since, into an epidemic that was producing cases at a rate that made the statistics resolve quickly.

Two things were not different. The trials were full-size, controlled, randomised and blinded — the phase 3 trials for the Covid vaccines were larger than those for most routine childhood vaccines. And the safety monitoring did not stop at authorisation; it got substantially more intense, because tens of millions of people were being vaccinated in a matter of

months under more scrutiny than any medical intervention in history.

There is a third thing, harder to classify, which is that an enormous number of people worked without stopping. That is not a scientific explanation and it does not belong in a list of mechanisms, but it is part of the answer and leaving it out would be a kind of lie. Trial sites ran at weekends. Regulators read through Christmas. Manufacturing staff who could not work from home went in every day for a year, in 2020, before there were any vaccines for them. The speed had a technical cause and a financial cause and it also had a cost that was paid in people's lives rather than in money, and a good deal of it was paid by people whose names are not in any of the papers.

That safety point is the hinge of the next chapter. The reason we know about the rare clotting syndrome associated with the adenovirus vaccines, and about myocarditis in young men after mRNA vaccines, is that the surveillance worked. It found signals affecting a few people per hundred thousand, in months, in the middle of a pandemic. A system that can do that is not a system that was cutting corners. But it is also a system that produced, continuously and publicly, a stream of bad news about the thing everyone was being urged to take

— which turned out to be a very difficult thing to communicate, and is the subject of chapter 18.

The part that did not go fast

In the first year of vaccination, from 8 December 2020 to 8 December 2021, 8.33 billion doses were given to 4.36 billion people.

THE NUMBER — 19.8 million

A modelling study published in *The Lancet Infectious Diseases* in June 2022 estimated how many deaths vaccination prevented in its first year, across 185 countries and territories.

Measured against excess mortality, the estimate was 19.8 million deaths averted, with a credible interval of 19.1 to 20.4 million. Without vaccination the model puts the toll for that year at 31.4 million. The vaccines, on this estimate, cut it by 63 per cent.

It is a modelled figure and it carries all the caveats chapter 16 attached to modelled figures. It is also the largest number in this book.

> Of the 19.8 million, 12.2 million were in high- and upper-middle-income countries.

That last line is the whole of the problem.

The same study estimated what the gap cost. In low-income countries, had the 20 per cent coverage target set by the COVAX facility been met by the end of 2021, an additional 45 per cent of deaths could have been averted. Had the 40 per cent target set by the World Health Organization been met, an additional 111 per cent — more, that is, than were actually averted there at all.

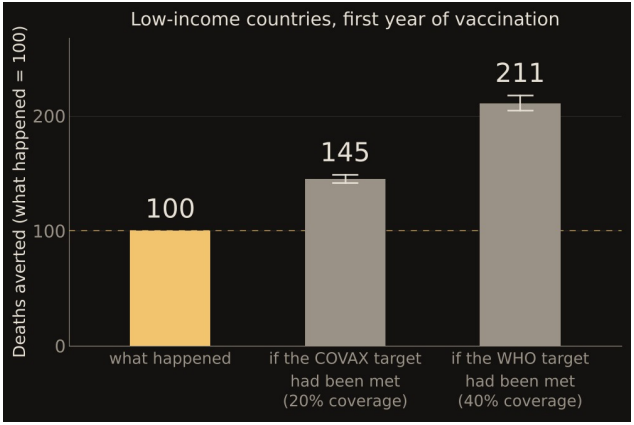


Figure F17-06. *Low-income countries: what vaccination achieved, and what two unmet targets would have added.*

Those targets were not ambitious. Twenty per cent coverage means four people in five unvaccinated. They were missed because the doses went where the money was, through advance purchase agreements signed in 2020 by countries that could afford to gamble, which is exactly the mechanism that made the vaccines fast in the first place.

That is not a comfortable sentence and it is not meant to be dismissed with one. The advance purchase model worked: it is why there were vaccines at all by December 2020. It also allocated the first year's supply by purchasing power. Both

of those are consequences of the same decision, and any serious proposal for doing better next time has to keep the first while fixing the second. Nobody has yet demonstrated how.

So the three hundred days are two stories. One is the fastest and most consequential piece of applied biology in the history of medicine, resting on a decade of unglamorous structural work that nobody could have justified at the time. The other is that having done it, the world distributed the result in the order of ability to pay, and a number somewhere north of ten million people died in the eighteen months afterwards who need not have.

Both happened. The chapter would be dishonest if it ended on either one alone.

Rare, Real and Rumoured

> *"Acting on the precautionary principle, and pending receipt of further information…" — Dr Ronan Glynn, Deputy Chief Medical Officer, Dublin, 14 March 2021*

Written in November 2026.

The fourteenth of March

On the morning of Sunday 14 March 2021, Ireland stopped giving the AstraZeneca vaccine.

The decision had been taken overnight by the National Immunisation Advisory Committee and was announced by the deputy chief medical officer. The reason was a report from the Norwegian Medicines Agency the previous day of four new cases of serious blood clotting in adults after vaccination, in people in their thirties and forties. The Irish statement was careful: it had not been concluded that there was any link between the vaccine and the cases. The committee

was acting on the precautionary principle, pending further information.

The AstraZeneca vaccine was about a fifth of Ireland's supply. The pause took thirty thousand doses out of that week's target.



Figure F18-01. *An empty chair. The pause lasted a little over a week and the argument about it has lasted five years.*

A hundred and fifty kilometres up the road, Northern Ireland carried on vaccinating. The British regulator said it was aware of what Ireland had done and that the evidence did not suggest the vaccine was causing the clots.

Two regulators, one island, the same four Norwegian reports, opposite conclusions, on the same Sunday.

Neither was being stupid. That is the thing to hold on to through this chapter, because the easy version of this story — one side was cautious and right, the other was reckless and wrong, or the reverse — is not what happened. They were weighing the same thin evidence against different judgements about what a regulator is for, and there is still no settled answer about which of them got it right.

Ireland resumed on 22 March, with the vaccine restricted by age. The clotting syndrome turned out to be real.

What they were looking at

The syndrome now has a name: vaccine-induced immune thrombotic thrombocytopenia, or VITT. It was described in April 2021 by a group in Germany led by Andreas Greinacher, from five healthcare workers aged between thirty-two and fifty-four who had become ill seven to ten days after a first dose.

It is a strange illness and it was identified because it is strange. Patients had clots, often in unusual places such as the veins draining the brain, and at the same time a very low platelet count. Clots and a shortage of the cells that make

clots do not normally go together, and that combination is the clinical signature of a condition doctors already knew: heparin-induced thrombocytopenia, in which an antibody response to the blood thinner heparin paradoxically triggers clotting. The German group found the same kind of antibody, against a protein called platelet factor 4, in patients who had never had heparin.

That is why VITT was accepted as real so quickly, and it is worth being precise about why, because it is the difference between this and most of what gets reported after a vaccine.

HOW WE KNOW — From a signal to a finding

Somebody reporting an illness after a vaccination is a signal. Signals are cheap: millions of people were vaccinated in a short period, so every illness that happens to anyone happened to somebody within days of a vaccination.

A signal becomes a finding when it survives three questions.

Is it happening more often than it should? That means comparing the observed count against the background rate — how often it happens anyway, in people of the

same age and sex, in a period when nobody was being vaccinated.

Is there a pattern? VITT clustered at seven to fourteen days after a first dose, in younger adults, more often in women. A pattern that tight is hard to produce by coincidence.

Is there a mechanism? This is the one that settles arguments. The German group did not merely count cases; they found the antibody, identified its target and explained how it produced both the clotting and the low platelets. Once you can say what is happening in the body, the question of whether it is really happening stops being open.

Most reported associations fail the first question. A few pass it and fail the second. Very few pass all three. VITT passed all three inside six weeks.

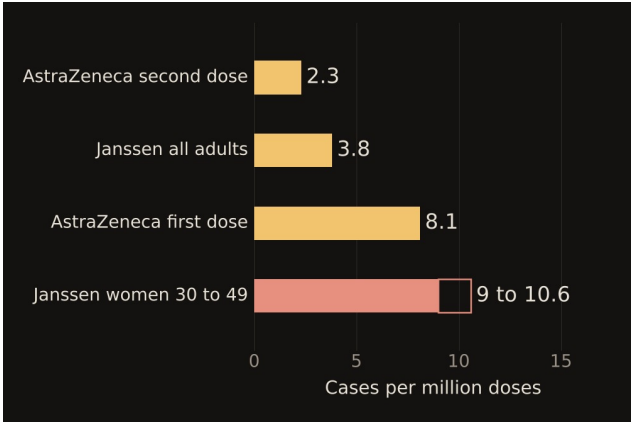


Figure F18-02. *How rare it was, once there was enough data to say. Rates per million doses.*

How rare is rare

Once there were enough doses to count against, the numbers came out like this.

A review of the AstraZeneca safety database covering Europe and the United Kingdom found 8.1 cases of VITT per million first doses, and 2.3 per million second doses. In the United States, the risk after the Janssen vaccine — a different adenovirus vaccine, same family of problem — was

estimated at 3.8 cases per million doses overall, and between nine and 10.6 per million in women aged thirty to forty-nine.

KEY CONCEPT — The background rate

"There were two hundred cases of X after vaccination" is not information. It is half of a fraction.

Give a hundred million people a vaccine and then wait a month, and some of them will have heart attacks, strokes, miscarriages and seizures, because a hundred million people have those things in any month. The question is never whether events happened after vaccination. It is whether more of them happened than would have happened anyway.

This cuts both ways, and the second direction is the one less often noticed. The background rate is also why a real effect can hide: if a vaccine caused a 10 per cent increase in something common, the extra cases would be invisible in the noise, which is why the things found fastest are the strange ones. VITT was found quickly precisely because its background rate is close to zero. Nobody gets clots and low platelets at once for no reason, so two dozen cases were enough to notice.

It is also why "I knew someone who" cannot settle anything, however sincerely it is offered, and why dismiss-

> ing those accounts is not an adequate reply either. Some of those people were describing something real. The only way to find out which ones is to count.

Eight per million means that for every person who developed VITT after a first dose of that vaccine, about a hundred and twenty-three thousand did not. It also means that in a country giving ten million first doses, around eighty people got it. Both of those sentences are true and they produce different feelings, which is most of the communication problem in this chapter.

The other one

The mRNA vaccines had a different rare effect: inflammation of the heart muscle, myocarditis, concentrated in young males and mostly after the second dose, usually appearing within a week.

The measured rates vary a good deal between studies, which is itself worth saying rather than smoothing over. American surveillance covering December 2020 to August 2021 found the highest rates after a second dose in adolescent males: about 70.7 cases per million second doses of the Pfizer-BioNTech vaccine in boys aged twelve to fifteen, and 105.9 per million in boys aged sixteen to seventeen. A population-

based study in Ontario found a rate as high as 299.5 cases per million second doses of the Moderna vaccine in males aged eighteen to twenty-four. A review of the whole literature put adolescents aged twelve to seventeen somewhere between 50 and 139 cases per million, with low certainty, and young adults aged eighteen to twenty-nine between 28 and 147, with moderate certainty. Across all adults aged twelve to thirty-nine the American figure was about 12.6 per million second doses.

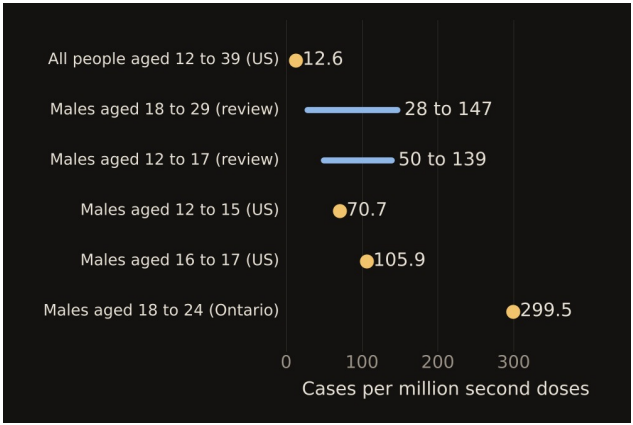


Figure F18-03. *Myocarditis after a second mRNA dose, by age and sex. The spread between studies is shown because it is part of what is known.*

Those ranges are wide. Some of the spread is real — the two mRNA vaccines differ, and the Moderna dose was larger — and some of it is an artefact of how cases were found, since a country running enhanced surveillance for a thing will find more of that thing than a country relying on doctors to report it unprompted. Anyone quoting a single figure for myocarditis after mRNA vaccination without saying which study, which vaccine, which age band and which sex is giving you a number with the uncertainty stripped out.

Most cases were mild by the standards of myocarditis, which is not the same as mild. Most patients were admitted to hospital, most recovered within days to weeks, and the great majority had normal heart function on follow-up. A small number did not. Myocarditis also occurs after Covid-19 itself, and several studies found it to be more common after infection than after vaccination, though the comparison is harder than it sounds, because infections are detected unevenly and the populations are not the same.

What a safety system is for

Step back from the particular numbers for a moment.

Both of these effects were found. VITT was described in a journal within about three months of the first mass vaccinations. The myocarditis signal was reported in May 2021 and

quantified within weeks. In both cases, the system detected something affecting a few people per hundred thousand, or fewer, in the middle of the fastest and largest vaccination campaign in history, and said so publicly.

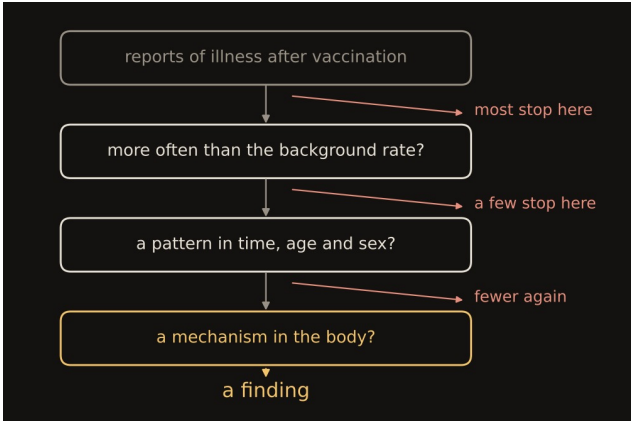


Figure F18-04. *The three gates a reported association has to pass. Most stop at the first.*

That is the system working. It is worth stating plainly because the same facts were used, at the time and since, as evidence that the system had failed — that the vaccines had been rushed and that these were the consequences coming to light. The logic does not hold. A programme that was concealing harms would not have produced a stream of pub-

lished papers quantifying them. The reason anyone knows about VITT and about myocarditis is pharmacovigilance, and pharmacovigilance is the part of the process that was running at full stretch.

But here is the difficulty, and it is a real one rather than a rhetorical concession.

A system of that sensitivity, running in public, during a campaign in which everyone is being urged to take something, produces a continuous stream of alarming headlines about the thing everyone is being urged to take. Each headline is accurate. Each represents the system functioning. And the cumulative effect on somebody who is not following the denominators is a strong impression that worrying discoveries keep being made.

There is no clever communication fix for this. Telling people the risks are small is true and does not help, because people are not actually bad at small probabilities when the stakes are personal; they are doing a different calculation, which includes how much they trust the person telling them. Chapter 10 made this point about measles and it applies here with more force.

Did the pauses help?

This is where the genuine argument is, and it is an argument about policy rather than about evidence.

> **THE NUMBER — 8.1 per million**
>
> That is the rate of vaccine-induced immune thrombotic thrombocytopenia after a first dose of the AstraZeneca vaccine, from a review of the manufacturer's safety database covering Europe and the United Kingdom. After a second dose it was 2.3 per million.
>
> Put beside it the fact that in the week Ireland suspended the vaccine, Covid-19 was killing several thousand people a week across Europe.
>
> Those two numbers are the whole of the argument in this section, and setting them side by side does not resolve it. A regulator is not permitted to kill eight people in a million in order to save more, at least not without saying so out loud and being held to it. That is an ethical constraint rather than an arithmetic one, and people disagree in good faith about how absolute it should be.
>
> Note also what the figure is not. It is the rate established months afterwards, from tens of millions of doses. On

> 14 March 2021 nobody had it. What the committee in Dublin had was four reports and a weekend.

The case for the pauses is the one Ireland made on 14 March. A regulator's job is to protect the public from products it has licensed. When a credible report arrives of a serious, unexplained harm, the honest response is to stop, look, and restart when you know more. Doing otherwise means deciding to keep administering something while you suspect it might be hurting people, on the grounds that stopping would look bad — which is precisely the behaviour that has destroyed trust in medicines regulators in the past. On this view the pause was not a cost to be weighed against lives; it was the thing that makes the whole system worth trusting.

The case against is that the pauses were an enormous, measurable harm imposed to resolve a question that could have been investigated without stopping. The vaccine was being paused across Europe at a moment when the virus was killing thousands of people a week on the same continent. Every week of delay had a body count that could be estimated, and the estimate was not small. The clotting syndrome, once quantified, turned out to be rarer than the risks people accept from many licensed medicines. And the pauses did

lasting damage to confidence in one specific vaccine — the cheap one, the one that did not need a freezer, the one intended for the rest of the world — in a way that outlasted the investigation by years.

Both of those are serious arguments made by serious people, and the second one has, in my reading, grown stronger as the evidence has come in rather than weaker. But it rests on a counterfactual: it assumes that a regulator that had not paused would have retained the public's confidence afterwards, and nobody can demonstrate that. It is also easier to make in 2026, knowing the rate was eight per million, than it was on 14 March 2021, when the entire evidence base was four reports from Norway and nobody knew whether the next week would bring four hundred.

What can be said without much hedging is narrower. The decision to restrict the adenovirus vaccines by age, once the rates were known, was correct and is not seriously disputed: VITT was concentrated in younger adults, who were also at the lowest risk from Covid-19, so for them the balance genuinely was different. The same reasoning led several countries to prefer one mRNA vaccine over the other for young men, which also looks right. Targeting a restriction at the group that carries the risk is a different act from suspending a vac-

cine for everyone, and the two should not be defended or attacked together.

What was rumoured

Alongside the two real findings ran a much larger body of claims that were not.

The vaccines did not alter anyone's DNA; chapter 13 explains why the mechanism forbids it. They did not contain tracking devices. They did not cause infertility — the claim originated in a misreading of a partial similarity between a short stretch of the spike protein and a placental protein, a similarity too short to mean anything, and large studies of pregnancy and fertility outcomes found nothing. They did not shed from vaccinated people to unvaccinated ones. They did not contain foetal tissue. The lipid nanoparticles did not accumulate in the ovaries, a claim built by taking a line out of a Japanese distribution study and reporting a figure that was a percentage of a very small total as though it were a large absolute quantity.

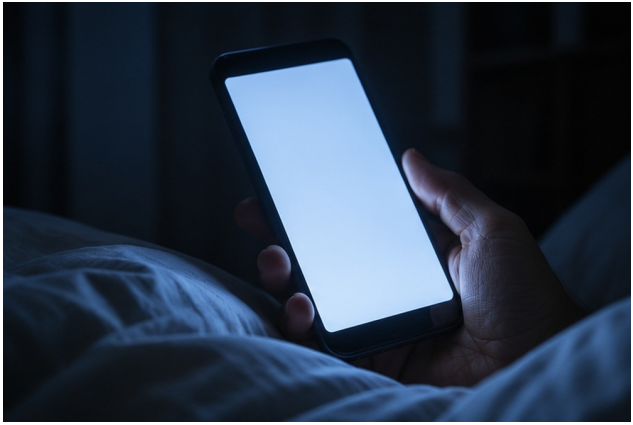


Figure F18-05. *Most of what circulated arrived this way, at night, from someone trusted.*

The interesting question is not why some people believed these things. It is why the real findings and the invented ones were so hard to tell apart from the outside, and the answer is uncomfortable: from the outside they looked identical. Both were reports of serious harm. Both were denied at first by officials who said the evidence did not support a link. Both were attributed to the vaccine by people who had experienced something frightening. In the case of VITT the deniers were wrong and the worriers were right. In the case of infertility it was the other way round.

Nothing on the surface of the two stories distinguishes them. The only thing that does is the work described in the HOW WE KNOW box above, which takes weeks, requires access to data most people cannot get, and produces an answer in a form — rates, intervals, denominators — that is actively unpersuasive to a frightened person. We ask the public to distinguish real from rumoured harms using a method they have no means of applying, and then express surprise when some of them use the method they do have, which is to decide who they believe.

And yet Ireland rolled up its sleeve

Chapter 10 left Ireland with measles coverage below the threshold, outbreaks in Dublin, and a vaccine the public had been taught to distrust by a fraud. This chapter has Ireland with one of the highest Covid-19 vaccination rates in Europe, achieved during a period in which the government had just suspended one of the vaccines on live television.



Figure F18-06. *The same country, the same decade, two very different answers to the same kind of question.*

The explanation is not that Irish attitudes to vaccination improved. It is that the two situations were not alike in the way the word “vaccine” suggests.

Measles, to an Irish parent in 2010, was an abstraction. Nobody they knew had seen it. The disease had been removed by the thing they were being asked to accept a small risk for, which is the central cruelty of a successful vaccination programme: it destroys the evidence of its own necessity. Covid-19 in early 2021 was the opposite. Everybody had seen it. The nursing home figures from chapter 16 were

not statistics to the Irish public; they were, in a country of five million, somebody's mother.

And the MMR scare had a named author making a specific, checkable, false claim. The Covid vaccine fears were diffuse, anonymous and unfalsifiable, which makes them more durable but much less able to move people who were frightened of the disease.

There is one more factor, and it reflects well on the pause. A government that stops a vaccine on a Sunday morning because four people in Norway got sick is demonstrating, expensively and in public, that it is willing to act against its own programme. Several commentators argued at the time that the suspension would destroy confidence. It is at least arguable that in Ireland it did the opposite. That is not a demonstrated claim and I will not pretend it is one — the counterfactual is unavailable, as ever — but it should be in the account alongside the costs.

Where this leaves us

The findings are settled. VITT is real, rare, mechanistically understood, and concentrated in younger adults after a first dose of an adenovirus vaccine. Myocarditis after mRNA vaccination is real, rare, concentrated in young males after a second dose, usually mild and usually resolving, with rates

that vary between studies by a factor of ten. For every group in which the vaccines were recommended after these findings were known, the benefits exceeded the risks by a margin that is not close and is not seriously disputed among people who work with the data.

What is not settled is the policy. Whether the European pauses of March 2021 saved more than they cost is genuinely open and probably unanswerable. Whether second doses should have been offered to teenage boys at the time they were is argued about by cardiologists and epidemiologists in good faith, and different countries reached different answers. Whether regulators should have been more forthcoming earlier about what they did not yet know is argued about by almost everyone, usually without specifying what the alternative announcement would have said.

And one thing is worse than it was. A large minority of people in several countries now hold, as a settled background belief, that the vaccines caused widespread serious harm that was covered up. That belief is false in its specifics. It is also, in its origins, a reasonable inference drawn by people who watched a safety system work exactly as designed and could not tell the difference between that and a scandal.

Fixing that is not a communications problem. Chapter 23 comes back to it, because it is the problem the next pandemic will start with.

The Long Tail

> *“Post COVID-19 condition occurs in individuals with a history of probable or confirmed SARS-CoV-2 infection.” — World Health Organization, Delphi consensus definition, October 2021*

Written in November 2026.

The people who did not get better

By the summer of 2020, while governments were arguing about lockdowns and companies were building vaccine factories, a third thing was happening that had no official status and almost no attention.

People who had been ill in March were still ill in July.

Not ill in the sense the textbooks allowed for. They had not been in intensive care. Many had not been in hospital at all, and a good number had never had a positive test, because in the spring of 2020 there were no tests to be had outside hospitals. They had had what they understood to be a moderate

illness, and then, weeks after it should have finished, they were exhausted in a way that sleep did not touch, breathless on a flight of stairs, unable to find words, with hearts that raced when they stood up.



Figure F19-01. *The commonest setting of this chapter.*

The specific quality that people kept describing was not tiredness. Everybody knows tiredness. What they described was a hard ceiling: a fixed and small allowance of effort, after which the body simply stopped cooperating, often not at the time but a day or two later. Cook a meal on Tuesday and be unable to get out of bed on Thursday. Walk to the shop and lose the following week. The delay is what made it

so hard to explain and so easy to doubt, including by the people experiencing it, many of whom spent months assuming that if they pushed a little harder they would break through.

They found each other online, because there was nowhere else to go. The name they are usually credited with coining — long Covid — came from that community rather than from medicine, and it arrived in general use before any health service had a code to record it under. That order of events matters, and this chapter comes back to it, because almost everything difficult about long Covid follows from the fact that the patients named it first.

Giving it a name

Medicine did eventually name it, several times over, which is part of the problem.

The World Health Organization settled on *post COVID-19 condition*. The British guidance body distinguishes *ongoing symptomatic COVID-19*, from five to twelve weeks, from *post-COVID-19 syndrome*, at twelve weeks and beyond. American agencies have used *post-COVID conditions* and researchers often write *post-acute sequelae of SARS-CoV-2 infection*, abbreviated to PASC. Patients, meanwhile, say long Covid, and so does nearly everyone else.

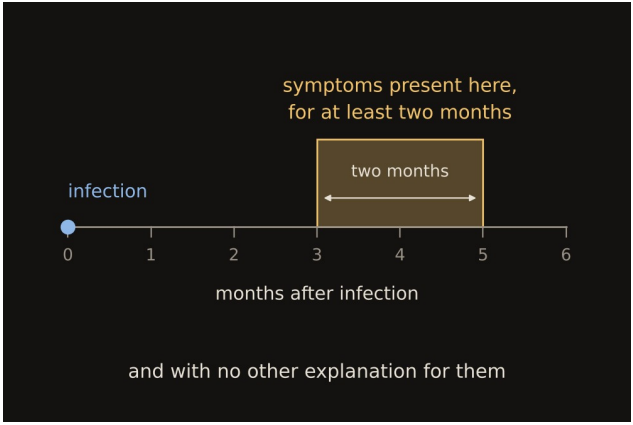


Figure F19-02. *The WHO definition as a clock: three months after infection, symptoms lasting at least two more, nothing else explaining them.*

The WHO definition, agreed by a structured expert consensus in October 2021, is the most widely used. It requires symptoms continuing or developing three months after the initial infection, lasting at least two months, with no alternative explanation.

> **KEY CONCEPT — A syndrome is not a disease**
>
> A disease is defined by its cause or its mechanism: tuberculosis is what *Mycobacterium tuberculosis* does to

you. A syndrome is defined by a pattern of symptoms occurring together, with the cause left open.

Long Covid is a syndrome. It is a label for a set of complaints that follow an infection, and it almost certainly covers several different things happening in different people for different reasons: damage left by the acute illness, a persistent immune response, small-vessel or clotting problems, disturbance of the autonomic nervous system, perhaps virus persisting somewhere out of reach.

This is not a criticism of the term. Every disease in this book started as a syndrome. Hepatitis was a syndrome until the viruses behind it were found one by one, and it turned out to be at least five separate diseases. AIDS was a syndrome for three years.

But it has a consequence that is often missed in arguments about long Covid. Two studies can both be done properly, both be honest, and disagree — because they are not studying the same group of people. Until a syndrome is resolved into its component diseases, every finding about it is a finding about a mixture.

How many, and why the number moves

Published estimates of how many people have had long Covid range from about 2 per cent of adults to about 20 per cent. That is a tenfold spread, and it is the first thing anyone looking into this runs into.

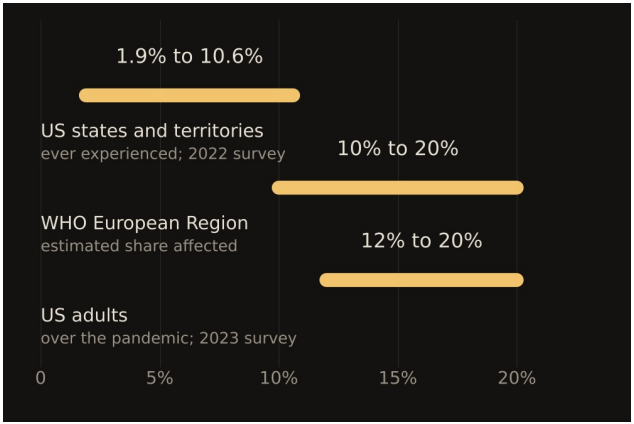


Figure F19-03. *Four estimates, four questions. The spread is mostly definitional.*

American state-level surveillance in 2022 found the proportion of adults who had ever experienced long Covid ranging from 1.9 per cent in one territory to 10.6 per cent in West Virginia, with seven states above 8.8 per cent. A later analysis of the same American survey system described 12 to 20

per cent of adults as having had the condition over the course of the pandemic. The WHO's European office has estimated that more than seventeen million people in its region may have experienced it in the first two years alone, and that the proportion affected should fall somewhere between 10 and 20 per cent.

HOW WE KNOW — Why the figure moves so much

Four choices, each defensible, each moving the answer by a lot.

Which definition. Twelve weeks or four? Must the symptoms be disabling, or merely present? Must another explanation have been ruled out, and by whom? The British five-to-twelve-week category and the WHO's three-month threshold will never produce the same number.

Which denominator. The proportion of infected people who develop it is a different quantity from the proportion of the whole population who have ever had it, and both get reported as "the prevalence of long Covid".

Ever or now. "Ever experienced" counts people who recovered in 2021. "Currently experiencing" counts people still unwell. Most recover; the two figures diverge further every year.

Who is asked, and how. A survey that asks people whether they think they have long Covid measures something different from a study that applies diagnostic criteria to medical records. Neither is wrong. They are different questions.

None of this is a reason to dismiss the estimates. It is a reason to ask, every single time, which of the four questions a number is answering before comparing it with another number.

Three things are not in dispute. The condition is real. It is commoner after severe infection but it happens after mild infection too, which is why the absolute burden is dominated by people who were never in hospital. And vaccination reduces the risk of developing it.

What might be going on

Several mechanisms are under serious investigation, and they are not alternatives: more than one is probably true, in different people.

There is the possibility of viral persistence — fragments of the virus, or whole virus, surviving in tissue reservoirs long after the acute infection has cleared from the airway, keeping

an immune response running. There is autoimmunity, in which the response to the infection produces antibodies that attack the body's own tissues; this is a known consequence of other infections. There is microvascular damage and abnormal clotting, which would fit the breathlessness and the exercise intolerance. There is dysautonomia, disturbance of the automatic nervous system that governs heart rate and blood pressure, which fits the racing heart on standing. And there is the possibility of reactivation of other viruses that most people carry quietly, Epstein-Barr among them — which, if it holds up, loops this chapter back to chapter 9.

It is worth dwelling on viral persistence for a moment, because it is the hypothesis that would change the most if it were confirmed. The assumption built into most of medicine is that an acute viral infection ends: the immune system clears the virus, the symptoms resolve, and whatever is left is damage rather than infection. That assumption is already known to be wrong for several viruses in this book. Chapter 8's HIV hides in resting cells for a lifetime. Chapter 9's Epstein-Barr never leaves. Measles, which most people think of as a childhood fortnight, can persist in the brain and produce a fatal condition years later.

If SARS-CoV-2 turns out to persist in gut tissue or elsewhere in a meaningful fraction of people, then long Covid is partly

an ongoing infection, and the thing to try is antivirals. If it does not, antivirals will do nothing and the effort should go elsewhere. That is not a subtle difference in emphasis; it is a fork in what to spend the next decade on, and it has not been settled.

What there is not, at the time of writing, is a diagnostic test, an agreed mechanism, or a treatment that addresses the cause rather than the symptoms. Trials are running. Some have reported. None has produced the thing patients most need, which is a blood test that comes back positive.

The absence of that test is not a small inconvenience. It determines whether a person gets a sick note, keeps a job, qualifies for benefits, and is believed by their family. A great deal of what makes this illness unbearable is administrative rather than medical, and it all follows from the fact that nothing shows up.

Why it was so hard to hear

Post-viral illness is not new, and this book has walked past it several times without naming it.

After the 1918 influenza pandemic came a wave of encephalitis lethargica, which left thousands with a Parkinson’s-like state for decades. Polio survivors developed post-polio syn-

drome, sometimes thirty years after their original illness. Ross River virus, Epstein-Barr, Q fever and Lyme disease all have well-documented post-infectious syndromes. The pattern of a lasting illness following an acute infection is one of the oldest observations in medicine.

It is also one of the most consistently disbelieved.



Figure F19-04. *The clinics that did not exist in 2019.*

Myalgic encephalomyelitis, also called chronic fatigue syndrome, is the clearest case. For decades, patients — disproportionately women — reported profound exhaustion made worse by exertion, and were told with great confidence by

respectable physicians that their illness was psychological, that they were deconditioned, that the cure was graded exercise. Patients said the exercise made them worse. They were not believed. It took until the last decade for the major guidance to change, and it changed because patients organised, collected evidence and would not stop.

It is worth being precise about the failure, because “doctors didn’t believe them” is too easy and lets the wrong lesson be drawn. The physicians involved were not inventing their position. They had a model — that a vicious cycle of fear, inactivity and physical deconditioning maintained the symptoms — that was coherent, that fitted some of what they saw, and that implied a treatment. They ran trials. The trials appeared to support them.

What went wrong was downstream of that. The outcome measures were largely subjective and the trials were not blinded, which in a treatment whose entire theory is about belief is a serious problem. The patients who got worse tended to drop out, which removes exactly the signal you need. And when patients said the treatment was harming them, the model itself explained that away: of course they say that, the illness makes them avoidant. A theory that converts its own contrary evidence into further confirmation has stopped being a scientific theory, and that is what happened.

That history is why long Covid patients were so quick to organise, and why they arrived at the medical profession already braced for an argument. They had watched what happened to the previous group. Some of them *were* the previous group.

To its credit, medicine moved faster this time. Clinics opened within a year. Research funding followed. The condition was recognised, named and coded while the pandemic was still running, which did not happen for ME. The reason it moved faster is uncomfortable but worth saying: a large number of the early patients were doctors, nurses and researchers, who had been exposed early and in quantity, and who could not be told that they did not understand their own symptoms.

It should not take that. The lesson from this is not that the system worked. It is that the system works when the patient is a colleague.

What helps

The honest summary of treatment is thin, and a chapter that pretended otherwise would be doing the reader harm.

There is no cure. What exists is management: pacing, which means staying within an energy envelope rather than pushing

through, now the standard advice and the direct opposite of what was told to ME patients for thirty years; specific treatments for specific components, such as drugs and compression for the circulatory problems; rehabilitation that is carefully graded and stopped if it makes things worse; and, for the substantial fraction who gradually improve, time.

Pacing deserves a word, because it sounds like doing nothing and is not. The thing being managed is post-exertional malaise: the delayed crash described at the start of this chapter. Because the crash arrives a day or two after the effort that caused it, the ordinary feedback loop a person uses to judge their own limits is broken. Pacing is the discipline of staying well inside a boundary you cannot feel, on the evidence of a diary rather than on how you feel at the time, and stopping while you still have something left. It is psychologically brutal — it asks people to decline things they are, in that moment, capable of doing — and it is the only intervention with broad agreement behind it.

Most people do improve. That is a real and underreported fact, and it belongs beside the rest. A minority do not, and for them this is now a chronic illness of indefinite duration that arrived from a two-week infection they may not even have tested positive for.

Prevention is the strongest tool available, and it is the familiar one: not getting infected, and being vaccinated if you are.

The emergency ends

On 4 May 2023, the WHO's emergency committee determined that Covid-19 was an established and ongoing health issue that no longer constituted a public health emergency of international concern. The declaration had been made on 30 January 2020. It had lasted three years and three months.



Figure F19-05. *A testing centre, afterwards.*

THE NUMBER — 7.1 million

By the end of November 2024, more than 776.1 million cumulative cases and about 7.1 million confirmed deaths had been reported to the WHO.

The second figure is the one to be careful with. It counts deaths that were confirmed as Covid-19 and reported by a national authority, which requires a test, a death certificate and a functioning reporting system. Chapter 16 described what happens to the count when any of those is missing, and the estimates of excess mortality are substantially higher than the confirmed toll — not by a few per cent, but by a multiple in some regions.

Both numbers are useful and they answer different questions. The confirmed count asks how many deaths we can document. Excess mortality asks how many more people died than should have. The first is a floor. The second is an estimate with wide intervals attached.

Nobody will ever know the real figure. That should be stated once, plainly, rather than smoothed over by picking whichever number suits the argument.

The end of the emergency was an administrative act, not a biological one. The virus did not go anywhere. It continues

to circulate, to evolve, and to kill people, mostly old and mostly in winter, in much the way influenza does. What ended was the legal and institutional state of exception: the emergency powers, the special procurement, the daily figures, the press conferences.

For most people that end came earlier and privately, at different moments through 2022 — the last time they wore a mask, the first crowded room that did not feel like a decision. There was no day on which it finished. Historians of the 1918 pandemic report the same thing, which is one of the reasons it vanished so completely from memory: a catastrophe that kills fifty million people and then simply stops being mentioned, for sixty years, because it had no ending to commemorate.

There is a tidier way to tell the end of this story, and it is worth naming so as not to tell it. In the tidy version the vaccines arrived, the virus weakened, and the world walked out. The first part is true and chapter 17 told it. The second is partly true and partly an artefact: later variants did appear to cause milder illness on average, but by the time they arrived almost nobody was meeting the virus for the first time, and a population with immunity experiences a milder disease whether or not the virus has changed. Separating those two is genuinely hard and is still argued about. The third part,

that the world walked out, is true for most people and not for the ones this chapter is about.

What Part III leaves open

Six chapters ago this part began in a hospital in Wuhan with a handful of pneumonia cases of unknown cause. It ends with a virus that is now one more thing in circulation, a body count that cannot be established, a set of institutional reports still being written, and several million people managing an illness that did not have a name in 2019.

Four things are genuinely unresolved and the rest of this book does not pretend to settle them.

Where the virus came from remains undetermined, and chapter 15 set out why both remaining explanations are consistent with the evidence and neither is directly evidenced.

Whether the restrictions were worth their cost is not answerable as a single question, and chapter 16 argued that treating it as one is the main way the argument goes wrong.

What the vaccine safety episode did to public trust is not measured, and chapter 18 declined to invent a measurement.

And what long Covid actually is remains open, with several mechanisms plausible and none established.

There is a fifth thing, less a question than a worry, and it belongs here rather than in a conclusion. Every institutional response described in these six chapters was built, more or less, out of nothing, by people improvising against a deadline. Some of it worked remarkably well: the sequencing networks, the trial infrastructure, the surveillance that found a syndrome at eight cases per million. Nearly all of it has since been dismantled, defunded or allowed to lapse, because the emergency powers that paid for it expired with the emergency.

The chair of Ireland's evaluation said, as chapter 16 recorded, that six years on we are probably no better prepared than we were in 2020. That is a remark about this specifically: not that nothing was learned, but that what was learned was embodied in temporary structures, and temporary structures go away.

Which is where Part IV begins. It looks forward rather than back: at the virus most likely to do this again and what has been built against it, at the vaccines now arriving for diseases this book has described as unbeatable, and at the one problem that outlasts every technical success in these pages. The first of those is currently in American dairy cattle.

Part Four

The future

The Bird in the Dairy

> *“A vast undercount.” — United Farm Workers on the official H5N1 case numbers, 11 November 2026*

Written in November 2026.

Yakima County, November 2026

A resident of Yakima County in south-central Washington died on 1 November 2026, having been hospitalised with severe avian influenza after exposure to infected dairy cattle. They had underlying health conditions and have not been publicly identified.

It was the first death from bird flu reported in the United States in 2026, and the third since February 2024, when the virus first got into dairy cattle. Two other people in the state had been confirmed infected, both since mid-October, both after exposure to infected cows.



Figure F20-01. *A milking parlour. The least likely setting in this book for the next pandemic, which is rather the point.*

This is the chapter that cannot be written with hindsight, so it will be written with a date on it instead. Everything in it is as it stood in the third week of November 2026. By the time you read it, some of what follows will have been overtaken, and the most likely direction of overtaking is that nothing much happened — which is also the most likely outcome and the hardest one to write about honestly.

What makes H5N1 worth a chapter is not that it is about to do anything. It is that it is the clearest available example of what a pandemic looks like in the years before it is a pan-

demic, if it ever becomes one: a virus doing something new in an animal population, a handful of people infected at the edges, a surveillance system that is better than it was and still not good enough, and an international framework that exists on paper.

A virus in the milk

Influenza A of the H5N1 subtype has been a worry since 1997. What changed in this decade is its range.

Since the current outbreak began in 2022, this virus has killed more than 168 million commercial and backyard poultry birds in the United States alone, through infection and through the culling that follows it. It has spread across wild bird populations on every continent including Antarctica, killing seabirds in numbers that have no precedent, and it has moved into mammals repeatedly: foxes, seals, sea lions, domestic cats, mink.



Figure F20-02. *Seabird colonies have been among the worst affected. A panzootic is a pandemic in animals, and this is one.*

Then, in February 2024, it got into American dairy cattle.

Two details of that spread are worth pausing on, because they are the kind of thing that only shows up when a virus does something nobody planned for.

The first is that the outbreak was established before anyone recognised what it was. Sick cattle do not look like a flu story. A herd goes off its feed, the milk changes, and a vet reasons through the list of things that usually do that, on which influenza does not appear. Nothing finds a disease that

nobody is looking for in that species, and nothing in the surveillance system was looking.

The second is that it then moved at the speed of cattle trucks rather than the speed of birds. A dairy herd is a traded commodity. Animals move constantly between farms and between states, which is a transmission route no bird virus has ever had before, and one that runs on a timetable.

Nobody expected any of it. Cows had not been a significant influenza host. The virus turned out to replicate extremely well in the udder, which means it is present in raw milk at very high concentrations, which in turn means that the ordinary business of a dairy — milking machines, shared equipment, the movement of animals between farms — is an efficient way of spreading it. By late 2026 it had been detected in dairy cattle in twenty-two states.

> **HOW WE KNOW — Looking in the milk**
>
> One of the few genuinely encouraging things in this chapter is how the spread in cattle was tracked, because for much of it nobody was testing the cattle.
>
> The virus in milk is detectable by the same kind of test used on people. Bulk milk — the pooled output of a whole herd, or of a whole processing plant — can there-

> fore be tested as a single sample, and a positive result says that the virus is somewhere in that pool without requiring anyone to test a thousand animals.
>
> The same principle, applied to sewage, gives wastewater surveillance. A city's wastewater contains material from everyone in it. Testing it finds a signal that would otherwise require testing the population, and it finds it without anyone having to feel ill enough to see a doctor.
>
> Both are examples of a general and underrated idea: when individual testing is too expensive, test the thing that everyone's output is mixed into. It tells you much less about any individual and far more, far faster, about the population.

The method has a limit, and the limit is the point of this chapter's epigraph. Bulk-tank testing finds the virus in cows. It does not find it in people. For people you still need a person who feels ill, who can afford to stop working, who is willing to be tested, and who is not afraid of what else a visit to a clinic might bring.

In Yakima Valley in November 2026, the farmworkers' union said publicly that cases among dairy workers were being substantially underreported, partly because many of those

workers are immigrants and there was an active immigration enforcement surge in the area at the time. The union did not produce a number, so the size of the gap is unknown. The state and county dashboards did not agree with each other either: the county listed three confirmed human cases while the state table listed one dairy-associated case for the year.

None of that means a hidden epidemic. It means that the figure in the next section is a count of detections, and that detections are made where someone is looking.

Why it has not jumped

Three people in Washington caught this virus from cows. None of them gave it to anybody else. There has never been a documented case of sustained person-to-person transmission of H5N1 anywhere, in twenty-three years.

That is not luck, and understanding why is the most useful thing in this chapter.

> **KEY CONCEPT — What a virus needs to spread between people**
>
> An influenza virus gets into a cell by grabbing a sugar molecule on the cell's surface. The sugars come in two slightly different arrangements. Birds have one kind lin-

ing their gut; humans have the other kind lining the nose and throat.

Avian influenza is built to grab the bird arrangement. Human influenza is built to grab the human one. That mismatch is the main reason bird flu does not spread in people: the virus arrives in a human nose and finds the wrong doorknobs.

Humans do have some of the bird-type sugar, but deep in the lungs. That explains the pattern of human H5N1 cases exactly. A person who inhales a very large dose — a poultry worker in an enclosed shed, someone milking an infected cow — can get the virus down far enough to infect those cells. The resulting illness is severe, because it starts in the lungs rather than the nose. And it does not transmit, because a virus that is replicating in the depths of the lungs is not being coughed into the air in the way that a virus replicating in the nose and throat is.

Severe and untransmissible are, in this case, two consequences of the same fact. What would change both at once is a change in what the virus grabs. That is a small number of specific changes in one protein, and it is the thing everyone is watching for.

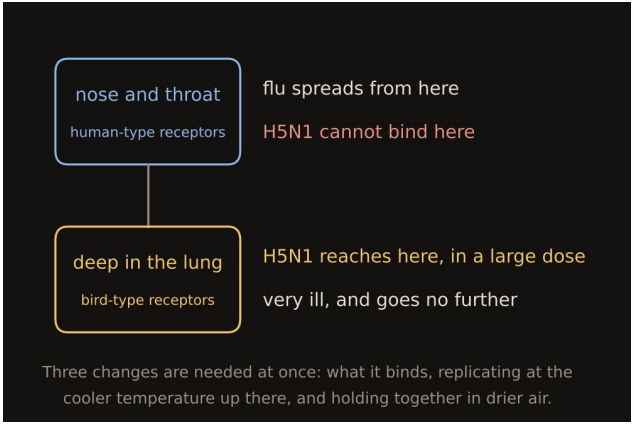


Figure F20-03. *What would have to change, and what has not.*

There is a second requirement, less often mentioned. The virus must also be able to replicate efficiently at the temperature of the human upper airway, which is cooler than a bird's gut or a human lung, and it must hold itself together in the drier conditions there. These are separate adaptations from the receptor change, and a virus needs all of them at once.

So the honest position is this. The changes required are known, there are not very many of them, and every additional mammal this virus infects is another opportunity to make them. The changes have not happened. There is no way to

say how likely they are, and anyone who gives you a probability is giving you an opinion dressed as a number.

What can be said is that the risk is higher than it was in 2020, because there is far more of this virus in far more mammals than there was then.

A thousand cases in twenty years

Here is the count, and here is why it needs handling with care.

From 1 January 2003 to 7 October 2026, the World Health Organization recorded 1,005 human infections with H5N1 from twenty-five countries. Of those, 479 were fatal.

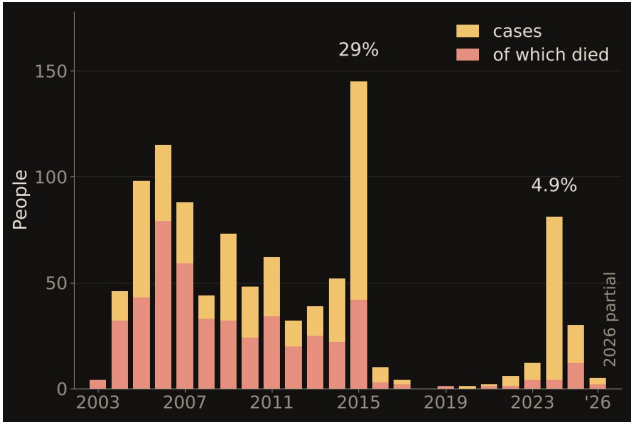


Figure F20-04. *Human H5N1 cases and deaths reported to the WHO, by year. The 2026 bar is partial.*

THE NUMBER — Forty-eight per cent, and why not to believe it

Divide 479 by 1,005 and you get a case fatality rate of 47.6 per cent. Roughly half of everyone known to have been infected with H5N1 has died of it.

That number is real and it is also one of the most misleading figures in infectious disease.

It is the proportion of *detected* cases that were fatal. For twenty years, almost everyone detected was someone

sick enough to be hospitalised in a country where H5N1 was suspected, usually after heavy exposure to poultry. Anyone with a mild infection who stayed at home was never counted.

The evidence that this matters is in the chart. In 2015 there were 145 reported cases and 42 deaths, a fatality rate of 29 per cent. In 2024, when the United States began actively looking for infections among dairy and poultry workers rather than waiting for them to arrive in hospital, there were 81 reported cases and 4 deaths — 4.9 per cent. The virus did not become twelve times less lethal in nine years. The denominator changed.

The United States has now recorded about 74 human cases since 2024 with three deaths. The Western Pacific region, over the same twenty-three years, reports a fatality rate near 66 per cent. Same virus, same subtype, different observers.

None of which makes H5N1 mild. The people who have died of it died of a genuinely terrible illness. It means only that nobody knows the true fatality rate, and that if this virus ever did learn to spread between people, the first thing that would happen to that figure is that it would fall a great deal — not because the virus had

> changed but because, at last, we would be counting everybody.

Vaccines on the shelf

Several countries hold stockpiles of H5N1 vaccine. They are made the old way — grown in eggs, which for an influenza virus that kills chicken embryos is an obvious difficulty — and they are made against a particular H5N1 strain chosen some years ago.



Figure F20-05. *A stockpile is a bet on which virus will arrive.*

The problem with a stockpile is the one chapter 11 described for seasonal flu, in sharper form. A vaccine made against the 2005 strain may or may not work against the 2026 one. Influenza drifts continuously, and H5N1 has been drifting for two decades across a vastly larger host range than it had before.

So the stockpile is best understood not as a solution but as a head start: a few million doses of something probably partly effective, available immediately, to protect the people most exposed while a matched vaccine is made. That is worth having. It is not the plan.

The plan, insofar as there is one, is the thing chapter 17 described. We now know that a matched vaccine against a new respiratory virus can be designed in days and tested in under a year, because it was done. The mRNA platforms exist, the manufacturing exists, and candidate H5N1 mRNA vaccines have been in trials. If H5N1 did start transmitting, the design step would not be the bottleneck.

The bottleneck would be the thing that was solved in 2020 by spending unlimited public money very fast, and which has not been solved since: manufacturing at scale, for everybody, in time. Chapter 17 ended on the fact that the world distributed the first year’s vaccine supply in order of ability

to pay. The next two sections are about the attempt to fix that, and about how it is going.

An agreement without a signature

In May 2025 the World Health Assembly adopted the WHO Pandemic Agreement. It took three years of negotiation and it was, by the standards of international health law, a substantial achievement.

Its central provision, Article 12, establishes a system called Pathogen Access and Benefit Sharing. The bargain is straightforward. Countries that find a new pathogen share the samples and the genetic sequence data promptly with the WHO. In exchange, manufacturers who use that material commit to making a fifth of their real-time production of vaccines, treatments and tests available to the WHO for distribution where it is needed rather than where it is paid for.



Figure F20-06. *The room where it has not happened yet.*

The reason that bargain exists is in chapter 15 and chapter 17. Sequence data from an outbreak is the raw material of every vaccine in Part III of this book, and the countries most likely to produce it are the least likely to receive what gets made from it. A country that watches its own outbreak sequence turn into a vaccine it cannot buy will, reasonably, think twice about sharing the next one. Indonesia did exactly that with H5N1 samples in 2007, and the dispute that followed is why this article exists at all.

The operational details of that system — what counts as a pathogen, how the sequence data is handled, how the twenty

per cent is enforced — were too contentious to settle in 2025 and were deferred to a separate, legally binding annex.

Article 31 of the agreement states that it opens for signature only after that annex is adopted. Article 33 says it enters into force thirty days after the sixtieth ratification.

The annex has not been adopted. It was due at the World Health Assembly in May 2026 and was not ready. Negotiations were extended, with a revised deadline of May 2027. Successive rounds through 2026 failed to produce a text; negotiators returned for a ninth round in the first fortnight of November 2026, as this was being written. In October 2026 the United Nations General Assembly adopted a political declaration on pandemics — not by consensus — urging states to finish the annex so that the agreement could open for signature at all.

Meanwhile the United States withdrew from the World Health Organization, with effect from 22 January 2026, and an executive order halted American participation in the negotiations.

So the position in November 2026 is that the world has a pandemic treaty that no country has signed, that cannot be signed, that would then need sixty ratifications, and which by one legal officer’s estimate would take a year to eighteen

months after that to come into force. The largest funder of global health has left the organisation that would administer it.

It would be easy to write that paragraph as a failure, and parts of it are. But the deadlock is not bureaucratic obstinacy. The disagreement is about whether access to pathogen samples should be legally conditional on benefit-sharing, and that is a real question with real stakes on both sides: tie them too tightly and manufacturers may walk away, leave them untied and the promise is worthless. Everyone involved watched what happened in 2021 and has drawn different conclusions about how to prevent it.

The hundred-day mission

Alongside the treaty runs a more concrete and more cheerful idea.

The hundred-day mission proposes that, from the moment a pandemic threat is identified, safe and effective vaccines should be available within a hundred days. The figure is deliberately aggressive. Covid-19 took about three hundred and thirty days, which was itself a record by a factor of ten.

Getting from three hundred to a hundred requires compressing the parts of chapter 17 that were not already compressed:

prototype vaccines for whole virus families designed and tested before any outbreak, so that the response is an adaptation rather than a start; trial protocols and sites agreed and sitting ready; regulatory review arrangements negotiated in advance rather than improvised; and manufacturing capacity distributed across more of the world rather than concentrated in a few rich countries.

Progress on each of those is real and partial. Prototype work on priority virus families is funded and running. Standing clinical trial networks exist in more places than they did. Regional manufacturing capacity, particularly in Africa, is being built, slowly and against difficult economics.

What is missing is the money and the political arrangement that would make all of it work in an emergency, which is what the treaty was for.

It is worth being clear about what a hundred days would actually buy, because the figure is often quoted as though speed were valuable in itself. A hundred days from identification is, roughly, the point at which a respiratory virus with the transmissibility of SARS-CoV-2 has finished its first wave in the countries it reached first. Vaccines at day one hundred would not have prevented the spring of 2020 anywhere. What they would have done is prevent the second

wave, and the second wave is where most of the deaths were.

That is the honest case for the hundred-day mission: not that it stops a pandemic, but that it cuts one short. Chapter 16's whole argument about lockdowns was that restrictions buy time at enormous cost and the time has to be spent on something. The hundred-day mission is a proposal about what to spend it on, and it is the strongest answer anyone has to the question that chapter could not resolve.

Where this leaves us

H5N1 may never do this. Most viruses that frighten epidemiologists do not. Chapter 12 listed several that could have and did not, and the base rate for "worrying animal virus becomes human pandemic" is reassuringly low.

But the shape of the thing is now visible in a way it was not before 2020, and three things about that shape are worth taking away.

The science is no longer the slow part. A matched vaccine for an H5N1 that had learned to transmit could be designed over a weekend, as the spike was in January 2020, and tested within the year. That is an extraordinary change and this book has spent six chapters earning the right to say it.

The counting is still the weak part. A fatality rate that moves from 29 per cent to 5 per cent because somebody changed who gets tested is not a measurement. Surveillance that stops at the farm gate because the workers on the other side of it are afraid of the authorities is not surveillance. Both of those are solvable and neither is expensive.

And the arrangement for sharing what gets made is the part that has not been solved at all. It was not solved in 2021, the treaty written to solve it cannot yet be signed, and the deadlock is genuine rather than procedural.

Which is to say that the next pandemic, if there is one, will be met with better tools, worse institutions and the same argument. The following chapter looks at the tools, because they are remarkable and this book has been gloomy for long enough.

Vaccines Against Everything

> *“Dose-series completion, not delivery alone, is now a key implementation challenge.” — The Malaria Vaccine Implementation Cascade, 2026*

Written in November 2026.

Four hundred and thirty-eight thousand

That is roughly how many African children died of malaria in 2024.

It is a number worth sitting with for a moment before the rest of this chapter, because the rest of this chapter is cheerful and the reason it is cheerful is that figures like that one have started to move.



Figure F21-01. *Most of what follows happens in rooms like this rather than in the laboratories where it was invented.*

Malaria is not a virus. It is a parasite, a single-celled organism with a life cycle so convoluted that it spends part of it in a mosquito's gut and part of it inside human red blood cells, changing its surface coat at every stage. That complexity is why a malaria vaccine took a century longer than a measles vaccine, and it is why this chapter opens with one: if a vaccine can be made against *that*, the question of what else is possible looks different.

It is worth being specific about why it took so long, because the reasons are instructive about everything else in this

chapter. A virus has, at most, a few dozen genes. The malaria parasite has thousands, and it uses them to present a different face at each stage of its life: the form injected by a mosquito is not the form that infects liver cells, which is not the form that bursts out of red blood cells. A vaccine that trains the immune system against one of those forms is irrelevant to the others. Worse, the parasite carries a family of genes whose entire function is to vary the surface of an infected red cell, so that the immune response is always a step behind. Natural infection does not produce lasting immunity either; people in high-transmission areas acquire partial protection only after repeated infections through childhood, which is another way of saying that many of them do not survive to acquire it.

The vaccines that now exist aim at the first stage only — the parasite in the few minutes between the mosquito bite and reaching the liver. That is a narrow window, which is why the efficacy figure in the next section is what it is.

What follows is a survey of the frontier as it stood in November 2026, with a rule applied throughout. For each thing on it, I have tried to say plainly which of three categories it falls into: proven in a trial and now in use; proven in a trial and not yet in use; or promising, and still a press release. The categories matter more than the enthusiasm.

Malaria, at last

Two malaria vaccines are now recommended by the World Health Organization: RTS,S/AS01, which completed its phase 3 trials in 2014 and showed about 36 per cent efficacy against clinical malaria over four years on a four-dose schedule, and R21/Matrix-M, which came later.

Thirty-six per cent sounds disappointing beside the ninety-five of chapter 17, and for years that figure was used to argue the vaccine was not worth deploying. The argument was wrong, and the way it was shown to be wrong is the best piece of evidence in this chapter.

> **HOW WE KNOW — Introducing a vaccine as an experiment**
>
> When RTS,S was first rolled out in Ghana, Kenya and Malawi, the rollout was designed as a study.
>
> A hundred and fifty-eight areas were randomly assigned either to begin giving the vaccine immediately or to begin later. The seventy-nine areas where introduction was delayed became the comparison group. Nobody was denied a vaccine that would otherwise have been available — the delayed areas got it, just afterwards — and in the meantime the two groups could be compared.

> This design answers a question that a conventional trial cannot. A trial measures what a vaccine does under trial conditions: enrolled participants, reminders, follow-up, trained staff. An introduction study measures what it does in a real immunisation programme, with the coverage you actually achieve and the doses people actually miss.
>
> Over the forty-six months after introduction, RTS,S prevented an estimated 993 deaths. All-cause child mortality among eligible children fell by 13.2 per cent — roughly one death in eight — with no difference between girls and boys. The results were published in *The Lancet* in May 2026.
>
> The crucial detail is the one the chapter's epigraph is about. Only 71 per cent of children received the third dose. That reduction was achieved at imperfect coverage, which means the achievable figure is higher, and it means the next gain is not a better vaccine but a better schedule.

Note what was measured: deaths from all causes, not malaria cases. A vaccine with 36 per cent efficacy against clinical malaria cut overall child mortality by an eighth. That is possible because malaria is such a large share of what kills chil-

dren in those places, and because a child who does not get malaria also does not get the anaemia, the secondary infections and the malnutrition that follow it.

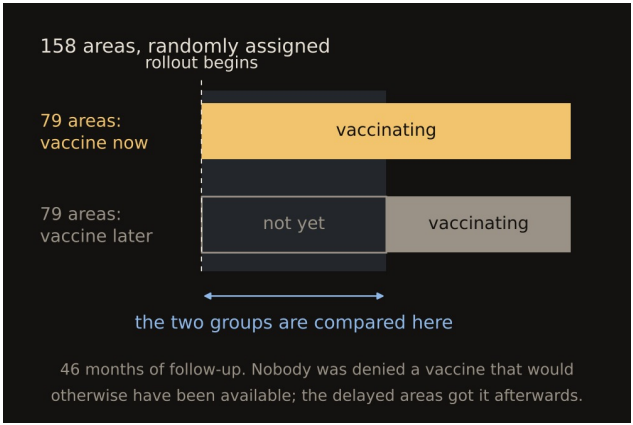


Figure F21-02. *How you test a vaccine that everyone is going to get anyway.*

As of 2026, twenty-five sub-Saharan African countries have added a malaria vaccine to their immunisation schedules, between them delivering more than 52 million doses since 2023, with more than ten million children targeted each year.

That is the fastest introduction of a new vaccine into African immunisation programmes that has ever happened, and it is

being done for a disease that most of the people reading this book have never personally feared.

The old and the very young

Respiratory syncytial virus is the most consequential virus most people have never heard of. Almost every child catches it by the age of two. For most it is a bad cold. For some, particularly in the first months of life, it fills the small airways and puts them in hospital, and it is one of the commonest reasons an infant is admitted anywhere in the world. It also kills older adults, in numbers that were substantially underestimated until recently because it was rarely tested for.

Chapter 13 described the structural work that made RSV tractable: the discovery that the virus's fusion protein springs into a different shape, and the mutations that hold it in the shape the immune system needs to see. That work was done on RSV before it was borrowed for coronaviruses. The RSV products arrived after the Covid vaccines, which is a strange order of events — the technique's original target was the last to benefit.

There are now three approaches, and only one of them is a vaccine in the ordinary sense.

Older adults get a vaccine. Pregnant women can be given a vaccine so that the antibodies cross the placenta and the baby is born protected. And infants can be given a long-acting monoclonal antibody, nirsevimab, which is not a vaccine at all — it supplies the antibodies directly rather than teaching the body to make them, which is why it works immediately and why it wears off.

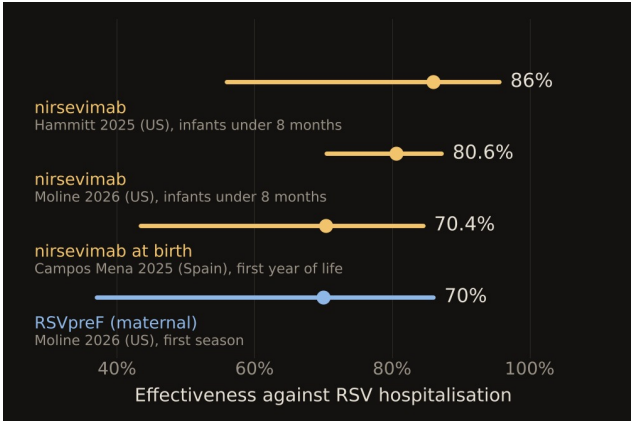


Figure F21-03. *Four estimates against RSV hospitalisation in infants, with their intervals. The studies differ in design and population.*

A systematic review for the 2025–26 season found that maternal vaccination, nirsevimab in infants and RSV vaccines in adults aged sixty and over all showed effectiveness of 68

per cent or more against hospitalisation. Individual studies of nirsevimab in infants under eight months have reported effectiveness against RSV hospitalisation of 86 per cent in one American series and 80.6 per cent in another, with a Spanish population study finding 70.4 per cent across the first year of life, and maternal vaccination giving 70 per cent against infant RSV hospitalisation in the first season.

Those are not trial numbers. They are effectiveness in use, in whole populations, in the first seasons of deployment. They are, by the standards of this book, extraordinarily good.

The nirsevimab case is worth pausing on, because it breaks a category this book has used throughout. Everything else in these chapters has worked by teaching the immune system to make its own antibodies, which takes weeks and lasts years. A monoclonal antibody does the opposite: it hands over finished antibodies, which work within hours and are gone within months. For a newborn facing a single winter that is precisely the right trade. A vaccine given at birth would not protect the child until the most dangerous weeks were over.

It also illustrates the economics that run under this whole chapter. Monoclonal antibodies are expensive to make and the dose scales with body weight, which is why giving them to every infant is affordable and giving them to every adult is not. The maternal vaccine exists partly because it solves the

same problem for a fraction of the cost, by using the mother's immune system as the factory and the placenta as the delivery route. Which of the two a country offers is a question about money rather than about biology.

The older-adult side has been slower to settle. The vaccines work, the effectiveness figures are good, and the arguments have been about who should get them and how often — at what age the first dose, whether a further dose is needed, and whether the cost per hospitalisation prevented justifies a universal offer. Those are the arguments a successful vaccine generates. They are a better class of problem than not having one.

Teaching the body to see a tumour

Chapter 9 described viruses that cause cancer and the vaccines that prevent those cancers by preventing the infection. This is something else entirely: a vaccine against a cancer that has already happened, made for one person, from that person's own tumour.

KEY CONCEPT — A vaccine against your own cancer

A tumour is your own cells with the brakes off, which is the central difficulty: the immune system is built not to attack you.

But a tumour's cells are not quite you any more. They have accumulated mutations, and some of those mutations change proteins in ways that produce fragments the immune system has never seen. Those fragments are called neoantigens — new antigens, belonging to no healthy cell in the body.

The procedure follows from that. Take a piece of the tumour. Sequence it, and sequence healthy tissue from the same person for comparison. Find the mutations unique to the tumour. Predict, computationally, which of the resulting fragments the immune system is most likely to recognise. Then build an mRNA vaccine encoding a few dozen of them and give it to the patient.

Every dose is different, because every tumour is different. This is not a product in the usual sense; it is a manufacturing process applied once per person, and the fact that it can be done at all inside a few weeks is a consequence of the mRNA platform described in chapter 13.

> No other technology could make a bespoke vaccine in that time.

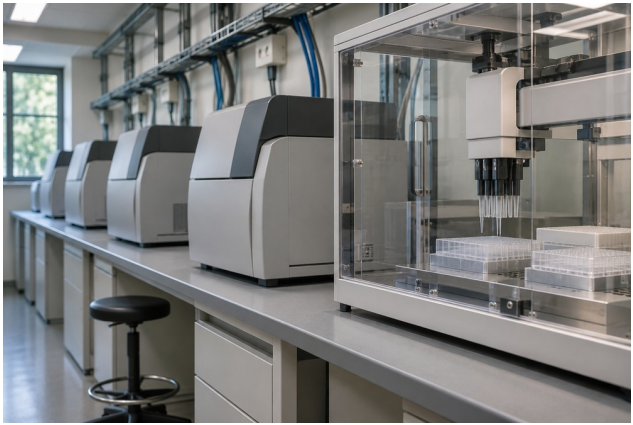


Figure F21-04. *Sequencing is the step that makes the rest possible, and it is the step that was impossibly expensive twenty years ago.*

In August 2026 Merck and Moderna announced that a phase 3 trial of this approach had met both its efficacy endpoints, and the full results were presented at the European oncology congress in October.

The trial, INTerpath-001, enrolled 1,137 patients with completely resected stage IIB to IV melanoma — that is, melanoma that had been surgically removed, in people at high risk

of it coming back. All of them received pembrolizumab, an established immunotherapy and the current standard treatment. Two thirds also received the individualised mRNA therapy.

THE NUMBER — Forty-two per cent of what

The headline was that the combination reduced the risk of recurrence or death by 42 per cent. The hazard ratio was 0.58, with a confidence interval from 0.45 to 0.75.

That is a *relative* reduction, and relative reductions sound larger than they are. Here is the absolute version. At two years, 82.2 per cent of patients on the combination were alive and free of recurrence, against 70.0 per cent on pembrolizumab alone. The difference is 12.2 percentage points.

Both numbers describe the same result. The 42 per cent is the proper statistical summary and the 12.2 points is what it means for a person deciding whether to have the treatment. A chapter that quoted only the first would be doing what press releases do.

One further detail, and it is the most interesting thing in the trial. For the first six months the hazard ratio was 0.99 — no benefit whatsoever. The separation appeared

only after that, and in the period after six months the reduction was 58 per cent.

That is exactly what you would expect if the thing were working as designed. It takes months to sequence a tumour, design a vaccine, make it and complete the course. The curves should not separate until the immune system has been taught, and they did not. A treatment whose benefit appears on the right schedule for its proposed mechanism is more believable than one whose benefit simply appears.

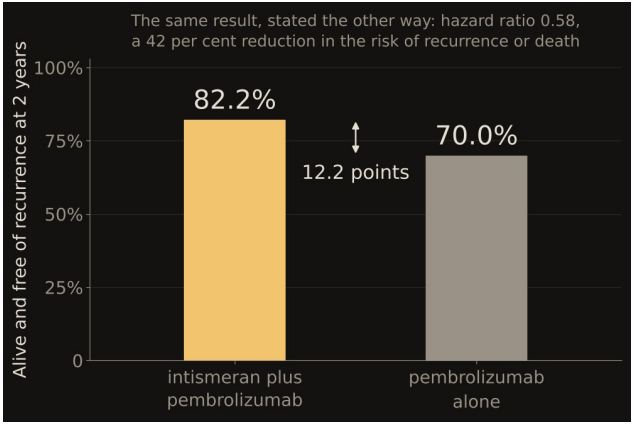


Figure F21-05. *Recurrence-free survival at two years. The absolute difference is the one to hold on to.*

Two cautions belong here, neither of which undoes the result.

This is an addition to an existing treatment, not a replacement for one. Everybody in the trial received standard immunotherapy; the question was what adding the individualised vaccine did. And the follow-up is short — a median of about two years. Recurrence-free survival is a real outcome that patients care about, but whether it translates into people living longer takes longer to find out.

What the trial establishes is the principle. A vaccine can be made from one person's tumour and it can change what happens to them. Trials in other cancers are running.

Melanoma was chosen first for a reason that will not generalise easily. It carries an unusually heavy burden of mutations, because the thing that causes it — ultraviolet light — damages DNA indiscriminately, so a melanoma offers an unusually rich supply of neoantigens to aim at. Lung cancers in smokers are similar for the same kind of reason. Many other tumours are not: a cancer driven by a handful of mutations presents the immune system with very little that looks foreign, and for those the approach may simply have less to work with.

There is also the question of what this costs and what it requires. Every dose involves a biopsy, sequencing of both tumour and healthy tissue, computational prediction, and a bespoke manufacturing run, inside a window short enough to matter clinically. That is a supply chain, not a product, and it exists today in a small number of places. Chapter 17's closing argument applies here before this treatment has even been licensed: a therapy that must be made individually will reach the people who can be reached, and the question of who those people are is being decided now rather than later.

One flu shot for life

Everything in chapter 11 about influenza vaccines follows from one fact: the virus changes its surface so fast that the vaccine has to be reformulated every year, against strains chosen months in advance by people making an educated guess.

A universal influenza vaccine would end that.

The idea rests on an asymmetry. The influenza surface protein is shaped something like a mushroom. The head changes constantly, and the head is what current vaccines aim at. The stem beneath it changes far more slowly, because it has a mechanical job to do and most changes to it break the virus. An immune response aimed at the stem should therefore work against strains the vaccine was never made against — including, in principle, the H5N1 of the previous chapter.

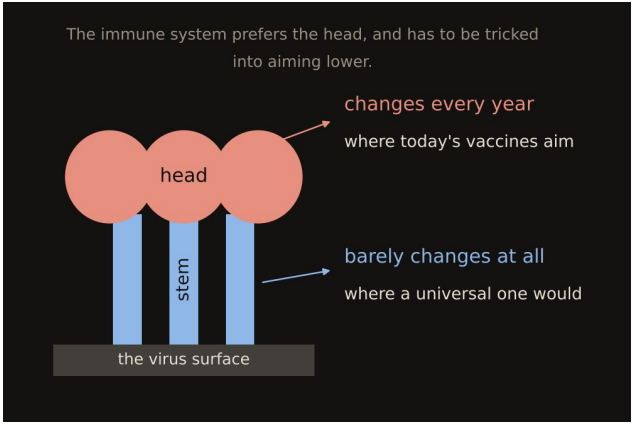


Figure F21-06. *The head changes every year. The stem barely changes at all. Current vaccines aim at the first; a universal one would aim at the second.*

The difficulty is that the immune system does not naturally aim at the stem. It goes for the head, which is more accessible and more immunogenic, and it has to be deliberately tricked into doing otherwise. Several approaches exist: stabilised stem constructs with the head removed entirely, mosaic designs presenting heads from many strains at once so that the only consistent target is the stem, and whole inactivated preparations given intranasally.

Where this stands, stated carefully. One stabilised-stem design has been through a phase 1 trial as an mRNA vaccine. A mosaic quadrivalent candidate has been in first-in-human testing. An intranasal whole-virus candidate has reached later-stage trials and is designed to block transmission rather than only disease, which is a different and more ambitious goal. In 2025 the American government announced a programme to develop universal influenza and coronavirus vaccines in-house, with trials beginning in 2026 and regulatory approval targeted for 2029.

That target should be read as an aspiration rather than a forecast. Universal influenza vaccines have been three to five years away for approximately twenty years. What has changed is that there are now several distinct candidates in human trials at once, and that the mRNA platform makes each iteration much faster to test than it used to be.

Separately, and much nearer, mRNA versions of the ordinary seasonal flu vaccine have reached late-stage trials. Those would not be universal. They would simply let the annual strain selection happen later in the year — closer to the season, with better information — which would make the guess less of a guess.

The ones still missing

Four failures, listed plainly, because a chapter like this is dishonest without them.

There is no HIV vaccine. Chapter 8 explained why: the virus mutates faster than the immune response can follow, it hides its vulnerable surfaces behind sugars, and it integrates into the genome within days of infection, so there is almost no window in which an antibody response can win. Trials have failed repeatedly. The most effective prevention now available is not a vaccine but a long-acting drug, and the arrival of injections that protect for months at a time has changed HIV prevention more than any vaccine candidate has. It is a magnificent result and it is not what this book is about.

There is no vaccine for hepatitis C, which is unusual, because there is an excellent one for hepatitis B. The difference is the same one as with HIV: an RNA virus with enormous diversity and a talent for escaping antibodies. There is now a cure — a course of tablets that clears the infection in most people — which has taken much of the urgency out of vaccine development, and which only helps people who know they are infected.

There is no cure for herpes simplex, which establishes itself in nerve cells and stays there for life, beyond the reach of

anything circulating in the blood. The same is true of the other herpesviruses, including the Epstein-Barr virus of chapter 9. A vaccine against Epstein-Barr would be a vaccine against several cancers and possibly against multiple sclerosis, and candidates exist, and none has yet worked.

And there is nothing for the common cold, which is not one disease but a few hundred, most of them rhinoviruses with no shared target worth aiming at. This is the failure that annoys people most and matters least.

Where this leaves us

The pattern across all of this is worth naming, because it is not the pattern anyone would have predicted twenty years ago.

The hard part has stopped being the science and started being everything else. A malaria vaccine took a century to make and then achieved a thirteen per cent cut in child mortality with only 71 per cent of children getting their third dose. The limiting factor there is not immunology; it is the fourth clinic visit. The RSV products work; the question is who is offered them and who pays. The melanoma result is real; whether a bespoke vaccine costing a great deal per patient can be delivered at population scale is not a scientific question.

That said, something has genuinely changed, and chapter 13 explains it. mRNA turned vaccine-making from a craft into something closer to software. Each new target is a different sequence through the same process. That is why a cancer vaccine can be made for one person in a few weeks, why a universal flu candidate can be tested and retested quickly, and why chapter 20 could claim that an H5N1 vaccine would not be the bottleneck.

The remaining failures are all cases where the virus has found a way to make the target itself the problem: too variable, too hidden, or already inside the cell before the immune system is told. Those are not engineering problems and they will not yield to a faster platform.

Which brings the book to its last real chapter, and to the thing that none of this solves. Every vaccine described in these pages only works on people who take it.

The Trust Problem

> *"…depends primarily on immunity levels in the affected community." — European Centre for Disease Prevention and Control, October 2026*

Written in November 2026.

The twenty-sixth of January

On 26 January 2026 the World Health Organization announced that six countries in its European region had lost their measles elimination status: Armenia, Austria, Azerbaijan, Spain, the United Kingdom and Uzbekistan.

Losing elimination status means something specific. It means that measles has been transmitting continuously in a defined area of that country for twelve months or more — that the virus is no longer being imported and stamped out, but living there.



Figure F22-01. *This is the mechanism. A disease that spreads in classrooms needs only a few unprotected children in the same one.*

Chapter 7 ended with the eradication of smallpox, the only human disease ever deliberately driven to extinction. Chapter 10 described how a fraudulent paper about the measles vaccine lowered uptake across the English-speaking world and what that cost in Dublin in the year 2000. This chapter is about the fact that measles, in western Europe, is now going backwards, and that the reason has nothing to do with the virus and nothing to do with the vaccine.

The vaccine is the same vaccine. It works as well as it did. What has changed is the number of people taking it.

Why measles comes back first

Measles is always the first thing to return, everywhere, and the reason is arithmetic.

> **THE NUMBER — Ninety-five per cent**
>
> Every infectious disease has a threshold: the proportion of a population that must be immune before the disease cannot sustain itself. Below that line each case causes, on average, more than one further case, and an outbreak grows. Above it, each case causes less than one, and chains of transmission die out on their own.
>
> Where the line falls depends on how contagious the disease is. For most of the diseases in this book it is somewhere between seventy and ninety per cent. For measles it is about ninety-five.
>
> Measles sets the highest threshold of any common human infection because it is the most contagious. It hangs in the air of a room for up to two hours after the infected person has left it. Somebody who has it is infectious for days before the rash appears, so nobody — including them — knows to be careful.
>
> Ninety-five per cent is a brutal number. It means that in a class of thirty children, one unvaccinated child is about

> the limit. It means a programme that reaches nineteen children in twenty is still failing. And it means measles will find the gap in a population where every other vaccine-preventable disease is still comfortably contained.
>
> This is why measles is the canary. When coverage slips, measles is what comes back first, and its return says that the slippage is real rather than statistical.

That threshold explains something else, which the rest of this chapter depends on. A disease needing ninety-five per cent coverage is a disease with no tolerance for a national average. The average is not what matters. What matters is whether there is any community, anywhere, where coverage is low enough to sustain a chain — and in a country at eighty-seven per cent on average, there certainly is.

Six countries, one year

The European figures need handling carefully, because the honest version is less dramatic than the headline and more worrying.

Across the WHO European Region, 33,998 measles cases were reported in 2025. Within the European Union and European Economic Area specifically, thirty countries repor-

ted 7,655 cases in the calendar year 2025. The rolling twelve-month total then fell steadily: 4,623 in the year to the end of February 2026, 3,779 in the year to the end of April, and 2,928 in the year to the end of May.

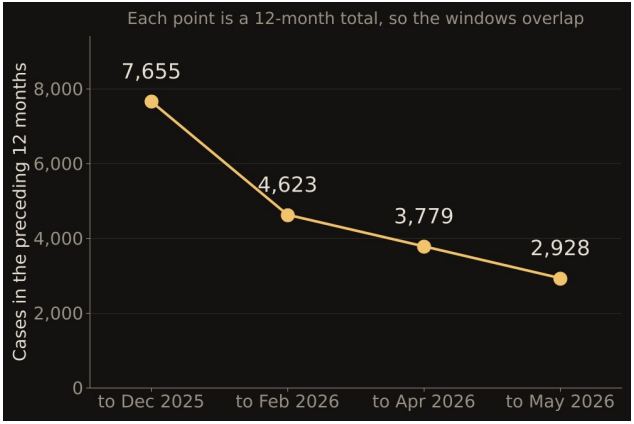


Figure F22-02. *Cases are falling. Each point is a rolling twelve-month total, so the measure is the same at every point.*

So cases are falling. That is real and it should be said first.

It is also almost beside the point, and the reason is in the word “elimination”. The goal is not a low number of cases. The goal is zero chains of transmission that start and continue inside the country, because every one of those chains is a failure that was preventable and every one of them is evid-

ence that somewhere there is a community below the line. Judged against a target of fewer cases than last year, European measles looks like a success story. Judged against the elimination goal that every one of these countries signed up to, it is not.

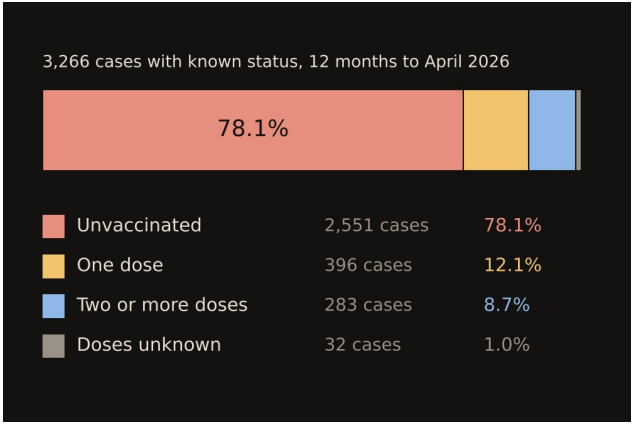


Figure F22-03. *Who is getting measles. Of cases with known status, roughly four in five had never been vaccinated.*

The breakdown of who gets it is the most useful figure in this chapter. In the twelve months to the end of April 2026, of the 3,266 cases with known age and vaccination status, 2,551 — 78.1 per cent — were unvaccinated. Another 12.1

per cent had received one dose. Only 8.7 per cent had received two or more.

In the 2025 calendar year the pattern was the same: of 7,210 cases with known status, 79.9 per cent were unvaccinated.

Two things follow. The vaccine works, which the eight per cent demonstrates rather than undermines — no vaccine is perfect, two doses of this one protect about ninety-seven people in a hundred, and in a large outbreak a small failure rate produces a small number of cases in vaccinated people. And the outbreaks are not happening to populations. They are happening to a particular, identifiable group of people within those populations, and that group is defined by a decision rather than by bad luck.

The United States is the sharper case, and it is sharper because the change there has been in policy rather than only in behaviour. American measles cases have surged, and the country is now at risk of losing the elimination status it has held since 2000. Vaccine exemptions claimed on personal or religious grounds are at their highest recorded level, having risen by 18 per cent between the 2024–25 and 2025–26 school years.

An exemption rate is a different kind of measurement from a coverage rate, and the difference matters. Coverage can fall

because people are busy. An exemption is a form that somebody filled in. It is a decision, recorded, and it is the cleanest available measure of the thing this chapter is about.

I am not going to adjudicate American vaccine politics in a book written in Ireland, and the arguments there are live and contested in ways that will have moved by the time this is read. The structural observation is the one that travels. When the dispute moves from what individuals believe to what the state requires and what its agencies say, the thing being argued about stops being evidence and becomes authority, and authority is much harder to restore than a statistic.

> **KEY CONCEPT — Why the average lies**
>
> Chapter 7 introduced herd immunity as a property of a population. This chapter has to revise that, because it is not quite true.
>
> Herd immunity is a property of a *contact network*. What protects an unvaccinated child is not the national coverage figure; it is the coverage among the people that child actually meets — their class, their street, their congregation, their family's circle.
>
> People do not distribute themselves randomly with respect to their beliefs. Families who decline vaccination

tend to know other families who decline vaccination: they share schools, neighbourhoods, communities, and increasingly the same information sources. So the unvaccinated are clustered, and clustering is what makes an outbreak possible in a country whose average looks adequate.

Two countries can report identical national coverage of eighty-seven per cent. In one, the missing thirteen per cent are scattered evenly and no community falls below the threshold; outbreaks fizzle. In the other, the missing thirteen per cent are concentrated in a handful of communities where coverage is fifty per cent; those communities will have a measles outbreak, and it is only a matter of when somebody brings it home from a holiday.

This is why the national figure, which is the only number anyone ever quotes, is close to useless for predicting outbreaks, and why public health people ask for coverage by small area rather than by country.

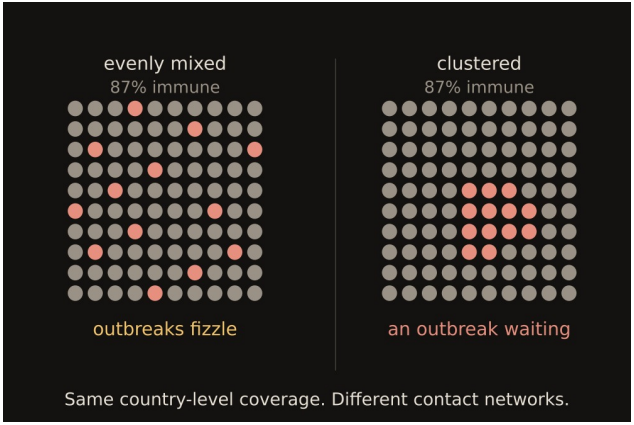


Figure F22-04. *The same national coverage, two different countries.*

Ireland’s eighty-seven per cent

Ireland achieved measles elimination status in 2017.

Its MMR coverage at twenty-four months was 87.6 per cent in the second quarter of 2025. In the third quarter it was 86.3 per cent. Both figures come from the national surveillance centre, both are below the ninety-five per cent threshold, and the second is lower than the first.

That is the position in a wealthy western European country with universal free childhood vaccination, no supply prob-

lem, no access problem that affects most of the shortfall, and a health service that has been running a catch-up campaign.

In September 2026 the surveillance centre reported that uptake was lower than anticipated despite those efforts, and set out the difficulties its own teams had encountered — among them limitations in public health staffing. In April 2026 the chief clinical officer of the health service told a parliamentary committee that measles and whooping cough outbreaks were the consequence of a continued decline in childhood vaccination, that rates had been decreasing and remained below ninety-five per cent, and that the reasons were, in his words, manifold and complex. He suggested one of them might be that people no longer understand how serious these diseases are.

Chapter 10 left Ireland with the measles vaccine damaged by a fraud with a named author and a specific, checkable, false claim. That fraud is twenty-eight years old. The paper is retracted, its author struck off, the claim refuted by a cohort of 657,000 children. And Irish coverage is lower now than it was then.

So whatever is happening is no longer about Wakefield. Something else is going on, and chapter 18 described the shape of it: a country that took the Covid vaccines in very high numbers while that disease was visible and frightening,

and will not quite complete a two-dose childhood schedule for a disease nobody has seen.

What hesitation actually is

The word “anti-vaxxer” does more harm than any other term in this subject, because it describes about two per cent of the problem and is used to explain all of it.

A committed opponent of vaccination — someone with a worked-out position who will not be moved — is rare. Most of the missing thirteen per cent are not that. They are, in rough order of size:

People who meant to and did not. The appointment was at an awkward time, the child was ill that week, they moved house, the letter went in a pile. This is the largest group almost everywhere, and it is not an attitude at all. It is friction.



Figure F22-05. *The room where most of this is actually decided.*

People who are not sure and have not decided. They have heard things. They have not sought out the sources of those things and could not tell you what they were. They are not refusing; they are deferring, and they will keep deferring unless someone makes the decision easy.

People with a specific worry, usually a real one that has been generalised. A reaction a relative had. A bad experience with a different medicine. Something in the family history. These people have a question and are often treated as though they had a position.

People who distrust the institution rather than the product. This group has grown and it is the one chapter 18 was about. They are not making a claim about immunology. They are making a claim about whether they are being told everything, and for several of them that claim has a specific grievance behind it.

And then the committed opponents, who are loudest, who are the ones anyone thinks of, and who are a small fraction of the shortfall.

Those five groups need five completely different responses, and only the last one is a debate. The first needs a better appointment system. The second needs a clear recommendation from someone they trust. The third needs a conversation. The fourth needs institutional behaviour over years rather than a leaflet.

Treating all five as the last one is the standard error, and it is worse than useless, because being argued at by someone who has assumed the worst of you is itself a reason to distrust them.

There is a further reason the groups get collapsed, and it is not laziness. The fifth group is the only one that produces anything to respond to. The first four are silent: a parent who missed an appointment writes nothing, a parent who is

quietly unsure says nothing in public, and a parent with a specific worry asks their own doctor rather than the internet. Everything visible — the posts, the videos, the comment sections, the person shouting at a public meeting — comes from the smallest group. So the visible evidence of hesitancy is drawn almost entirely from the part of it that is least representative and least reachable, and both the people designing responses and the people reading the news end up with a picture of the problem that is dominated by its rarest form.

That is worth sitting with, because it has a practical consequence. If you believe the shortfall is made of committed opponents, you will spend your effort on argument, and argument does not work on the four groups that make up most of it. If you believe it is mostly friction and deferral, you will spend it on clinic hours and reminders, which do work, and which nobody will ever thank you for because nothing visible will happen.

What rebuilds trust, and what does not

Here the honest answer is thin, and a chapter that pretended otherwise would be doing the thing this book has spent twenty-one chapters refusing to do.

HOW WE KNOW — The evidence on changing minds is weak

There is a substantial research literature on interventions to reduce vaccine hesitancy, and the most reliable summary of it is disappointing.

A review of fifteen earlier reviews and meta-analyses concluded that there is no strong evidence to recommend any specific intervention. A systematic review and meta-analysis comparing educational with motivational approaches found both improved uptake, that the effect for education was weak, and that the effect for motivational approaches was uncertain because so little research had been done. A 2026 systematic review of fifty-nine experimental studies across Europe and North America found information- and motivation-oriented interventions generally more effective than incentives or mandates, with digital interventions giving promising but heterogeneous results, and effectiveness improving when interventions were tailored to the local context. Several reviews report that heterogeneity between studies is so great that no single conclusion is possible.

The clearest effects in the whole literature come from the least interesting interventions. Two large randomised

> trials of text message reminders, with more than 160,000 participants between them, produced an absolute increase in uptake of 3.57 per cent for the first reminder and 1.06 per cent for the second.
>
> That is worth stating plainly. The best-measured, most reliably replicated intervention in this field is a text message, and it moves uptake by about three percentage points. Nothing in the literature reliably changes a settled mind.
>
> What the text message result actually shows is the first group from the previous section. Three per cent of people vaccinated their child because they were reminded, which means three per cent had not decided against it at all.

The practical conclusion is unglamorous. The reachable gains are in friction rather than in persuasion: appointments that are easy to make and easy to keep, reminders that arrive, catch-up offered without judgement, clinics open at times when working parents can attend, and a clear recommendation from a clinician the family already knows.

The unreachable gains are in trust, and trust is not an intervention. It is the residue of how an institution has behaved.

Chapter 18 argued that the pharmacovigilance system of 2021 worked exactly as designed and was widely read as a scandal, and that nothing about the way it was communicated would have fixed that. The same applies here. You cannot run a campaign to be trusted. You can only be the kind of organisation that, over years, turns out to have been telling the truth, including when the truth was unwelcome.

Almost alive

This book is called *Almost Alive* because a virus sits on the boundary of the category. It has genes and it evolves and it does nothing whatever on its own. Outside a cell it is an inert parcel of chemistry; inside one it is the most successful thing on the planet. Chapter 3 introduced that puzzle and the book has not resolved it, because it does not resolve.

There is a second sense in which the title applies, and it belongs at the end rather than the beginning.

Every disease in these pages was defeated by an argument as much as by a molecule. Lady Mary Wortley Montagu had to persuade a London that thought her mad. Jenner had to persuade the Royal Society, which declined. Semmelweis was destroyed by the persuading. Henderson's eradication programme was unpopular inside his own organisation before it became the greatest thing public health has ever done. The

HPV vaccine was reasoned back into Irish schools by a dying woman in her twenties who made the argument better than the state had.

The science is the easy half. It was always the easy half. What this book has actually been about is the other half: the business of getting a true thing accepted by people who have every right to ask why they should believe you, and who will not be moved by being told that the question is stupid.

Where this leaves us

So the position, at the end of all this.

We can design a vaccine against a new virus in two days. We can test one in under a year. We can make a vaccine against a parasite with a thousand genes, and one against a single person's own tumour, and we are some way towards one that would work against every strain of influenza at once.

And measles is back in six European countries, and Irish coverage fell between one quarter and the next, and the thing that stops a child being vaccinated is usually not a belief but an appointment on a Tuesday afternoon.



Figure F22-06. *The decision that all of this comes down to.*

I will not end this chapter by telling you what to do, because that would be exactly the mistake the chapter is about, and because I do not know your circumstances, your family's history or what your doctor has told you. Those decisions belong with you and with the clinician who knows you, and that is not a disclaimer, it is the actual answer.

What I will say is what this book has tried to earn the right to say. The evidence on these vaccines is better than the evidence on almost anything else in medicine, because more people have taken them than have taken anything else, and they have been watched more closely than anything else, by

more people with more reasons to find a problem. That is not a reason to stop asking questions. It is the reason the answers are as good as they are.

One short chapter remains, about the enemy this book has mostly ignored.

The Other Enemy

> *“Stability or small improvements.” — World Health Organization, Global Tuberculosis Report 2026, on the year 2025*

Written in November 2026.

The size of the enemy

Four days before this was written, the World Health Organization published its annual tuberculosis report.

In 2025, an estimated 10.6 million people fell ill with tuberculosis and 1.13 million died of it. It remains the world’s leading cause of death from a single infectious agent, and it ranks thirteenth among all causes of death of any kind.



Figure F23-01. *Mycobacterium tuberculosis. The first bacterium in this book, twenty-three chapters in.*

This book has been about viruses. Tuberculosis is not a virus. It is a bacterium, and the difference is the subject of this short closing chapter.

A virus, as chapter 3 established, is genes in a coat. It cannot do anything on its own: no metabolism, no growth, no reproduction except by getting inside one of your cells and using the machinery there. That is why the defences against it are what they are. You cannot poison a virus, because there is almost nothing in it to poison that is not also in you. What you

can do is teach the immune system to recognise it before it arrives, which is what every vaccine in this book does.

A bacterium is a cell. It eats, it builds, it divides, it has its own machinery for doing all of that — machinery which is not the same as yours. Which means it can be poisoned selectively. That is what an antibiotic is: a molecule that wrecks something a bacterium needs and a human cell does not.

That single difference has shaped the whole of the last century of medicine, and it explains why this book is about viruses and the next one would have to be about something else.

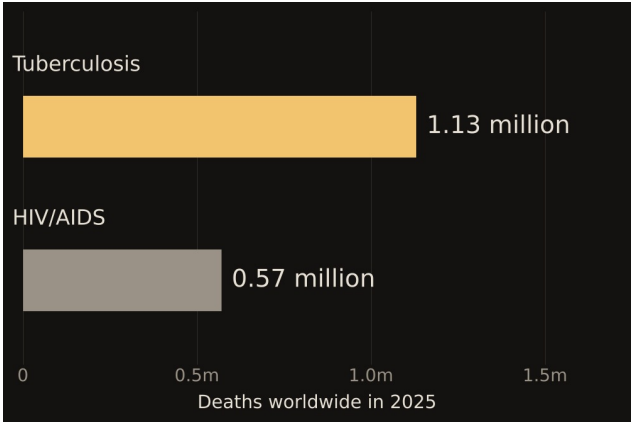


Figure F23-02. *Two infectious killers in the same year. One of them has a vaccine in development; the other has had one since 1921 and it does not work well enough.*

A century of antibiotics

When a disease can be cured, the pressure to prevent it falls away.

That is the plain reason the twentieth century built so many vaccines against viruses and so few against bacteria. Chapter 6 described the extraordinary run of viral vaccines from the 1950s to the 1980s. Over the same decades, bacterial disease was being handled a different way: you caught it, you took a course of tablets, you got better. A vaccine against something

curable is a harder commercial and public health case to make, and in most places it was not made.

There were exceptions, and chapter 6 covered them: diphtheria and tetanus, where the damage is done by a toxin that a vaccine can neutralise, and whooping cough. But the pattern held. For most bacterial disease, the answer was treatment rather than prevention, and for seventy years that answer worked.



Figure F23-03. *Before antibiotics, this was the treatment: rest, altitude and air. It is what tuberculosis care looked like within living memory.*

There is a second reason, less often stated, and it is about what bacteria do to us. Many of the bacteria that cause seri-

ous disease are organisms we already carry harmlessly — in the nose, the throat, the gut — and that only become dangerous when they get somewhere they should not be, or when the person carrying them is weakened. A vaccine against such an organism is a strange proposition: you are immunising against something that is already living in most of the population and causing nothing. The target is not the organism but a particular thing it occasionally does, and that is a harder brief than “teach the body to recognise this virus”.

It is worth remembering how complete the change was. Before antibiotics, tuberculosis was treated with mountain air and bed rest in sanatoria, for years, with outcomes that were largely a matter of luck. A scratch could kill you. Childbirth was dangerous in a way it no longer is anywhere with a functioning hospital. Antibiotics did not improve that situation; they abolished it, inside a single generation, and they did it so thoroughly that most people alive today have no working intuition for what the world was like beforehand.

Learning not to die

The trouble with antibiotics is in their nature, and chapter 2’s account of natural selection is all the explanation needed.

KEY CONCEPT — Why antibiotics stop working and vaccines do not

Give an antibiotic to a person with a bacterial infection and you kill the bacteria. Almost all of them. In a population of billions, a few will carry some chance variation that lets them survive — a pump that expels the drug, an enzyme that breaks it down, a slightly different shape at the place the drug attaches.

Those few are now the only ones left, and they have the patient to themselves. Every course of antibiotics is therefore an experiment in natural selection with the outcome built in. Use the drug enough times, in enough people, and the survivors become the norm.

Bacteria make this worse by swapping genes directly. Resistance genes travel on small loops of DNA that can be passed from one bacterium to another, including across species. A resistance gene that arises in a harmless gut bacterium can end up in a dangerous one without either of them having had to evolve it.

A vaccine does not work this way, and the reason is worth being exact about.

An antibiotic is a single molecule attacking a single target, so a single change can defeat it. A vaccine trains an

immune response against many different parts of the organism at once, in many different people, each of whom responds slightly differently. There is no single change that escapes all of that.

A vaccine also acts before the infection rather than during it. The selection pressure an antibiotic applies falls on a vast population of bacteria multiplying inside a sick person. A vaccinated person does not become sick, so that population never exists.

This is not a claim that vaccine escape is impossible — chapter 19 described influenza and SARS-CoV-2 doing exactly that. It is a claim about rate. Pathogens escape vaccines over decades. They escape antibiotics in years.

There is one more asymmetry, and it is the cruellest. An antibiotic becomes less effective the more it is used, so every doctor who prescribes one correctly is spending down a resource that belongs to everybody. A vaccine becomes *more* effective the more it is used, because each additional vaccinated person removes a link from the chains of transmission that threaten the rest. The two technologies have opposite relationships with their own popularity, and almost all of the economics of the last century has been built around the first one.

Resistance is now rising across the bacterial world, and tuberculosis is among the worst affected. The WHO has estimated that about three per cent of new cases and eighteen per cent of previously treated cases had multidrug-resistant disease, which translates into close to half a million new cases a year of a disease whose standard treatment no longer works. Treating it takes months to years, with drugs that are toxic, and the results are substantially worse.

Which brings the argument round. If the drugs are failing, prevention stops being the less urgent option and becomes the main one — which is where this book came in.

The vaccines that waited

The good news is that the bacterial vaccines were never impossible. They were simply not the priority.

The conjugate vaccines of the 1990s and 2000s are the demonstration. Several of the worst bacterial killers of small children — the organisms that cause most bacterial meningitis and much pneumonia — hide behind a sugar coat that young immune systems cannot respond to properly. The trick, worked out in those decades, was to chemically link that sugar to a protein, which lets the immune system treat it as a protein and respond to it properly. It worked. Diseases

that filled children's wards in living memory became rare in every country that introduced them.

THE NUMBER — 1.13 million, every year

Tuberculosis killed 1.13 million people in 2025, including 137,000 people living with HIV.

Set that beside the figure from chapter 19: about 7.1 million deaths from Covid-19 confirmed and reported to the WHO over the whole course of that pandemic, to the end of November 2024.

Tuberculosis kills, every year, roughly a sixth of what the defining catastrophe of our lifetimes killed in total. It has been doing so for every year any of us has been alive, and for most years before that.

The 2025 figure represents a four per cent fall from 2024 and a thirty-five per cent fall from 2015, and the report describes it as a milestone reached five years behind schedule. Both of those things are true: real progress, and a disease that still kills more people than any other single infectious agent.

The reason you have not been thinking about it is that it is no longer a disease of rich countries. That is the whole of the reason.

The conjugate vaccines did something else worth noticing, and it is the strongest argument in this chapter. Because they stopped people carrying the organism in the throat as well as stopping disease, they protected the unvaccinated too — the herd immunity of chapter 7, working against a bacterium. And because they removed the organism rather than killing it in sick people, they took away the occasions on which antibiotics would have been prescribed.

That is the quiet argument for bacterial vaccines over bacterial drugs. A vaccine does not merely avoid creating resistance. It reduces the number of times anyone reaches for the thing that does.

Malaria belongs in this chapter too, awkwardly, because it is neither a virus nor a bacterium but a parasite — a eukaryote, a cell more like one of yours than like a bacterium. Chapter 21 described the vaccines now reaching twenty-five African countries, and the evaluation showing one child death in eight averted. That result matters here as a proof of concept: if a vaccine can be made to work against an organism that complicated, the argument that bacterial targets are too difficult loses most of its force.

And tuberculosis itself now has candidates. The vaccine in use, BCG, dates from 1921; it protects small children against the most severe forms and does not reliably protect adoles-

cents or adults against the pulmonary disease that spreads. As of August 2026 three candidates were in phase 3 trials. The one most watched, M72/AS01E, is being tested in South Africa, and the answer on whether it protects older children and adults is expected by 2028.

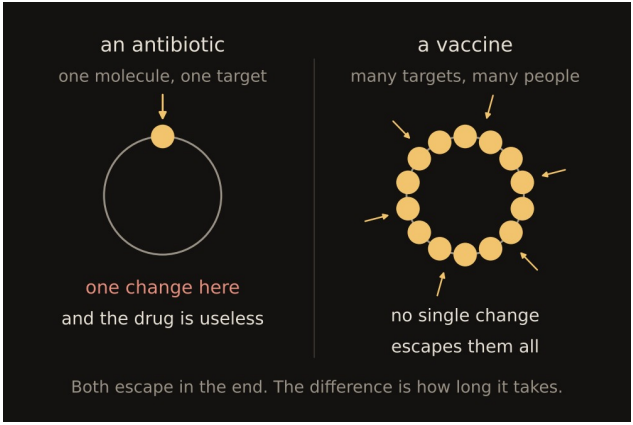


Figure F23-04. *Why one change defeats a drug and does not defeat a vaccine.*

A tuberculosis vaccine that worked in adults would be the largest single advance in global health available this decade. It is not a certainty. It is a trial with a result due.

It is also, if it arrives, a vaccine that will meet every problem chapter 22 described and several of its own. It will have to

be given to adolescents and adults rather than to infants, which means reaching people who are not already in a vaccination system. It will be needed most in the countries least able to pay for it, which is chapter 17's argument again. And it will arrive into a world that has just spent six years arguing about vaccines. None of that is a reason not to want it. All of it is a reason to notice that the trial result, whenever it comes, will be the easy part.

The next book

Everything in this chapter is a gesture at a subject rather than an account of one.

The history of bacteriology has its own cast — Koch, Ehrlich, Fleming, Florey and Chain, Hodgkin, Franklin's unacknowledged crystallography, the long argument about whether Semmelweis's hands mattered. It has its own disasters, its own frauds, its own scandals of access and price. It has the sanatorium century, the discovery of penicillin and the thirteen years it took to make it usable, the astonishing decades afterwards, and the present slow unravelling.

That is a book, and it is not this one.

What makes it urgent rather than merely interesting is the thing this chapter has been building towards. For seventy

years, bacterial disease was the solved problem and viral disease was the hard one. That has begun to reverse. We can now design a vaccine against a new virus in two days, as chapter 17 described, while the drugs that made bacteria manageable are being steadily disarmed by the organisms themselves.

The tools this book has spent twenty-three chapters describing are about to be needed on the other side of the line.

And they transfer less well than you might hope. mRNA, the technology that made Part III possible, encodes a protein — which is exactly what you want against a virus, where the target is usually a protein on the surface. A great deal of what the immune system needs to recognise on a bacterium is not protein at all; it is sugar, which is why the conjugate vaccines had to be invented in the first place and why they are made by chemistry rather than by genetics. The platform that can produce a new vaccine over a weekend cannot, as things stand, produce those.

So the next book is not a sequel in which the same methods are pointed at a new enemy. Some of them will transfer. Some will not, and new ones will be needed, and that work is largely not funded, for the same reason the bacterial vaccines waited the first time: the drugs still mostly work, and "still mostly" has been good enough for seventy years.

Where this leaves us

A last word about the thing the title has been pointing at.

A virus is almost alive. It has the information and none of the apparatus. It is, on its own, a crystal — chapter 3 described Wendell Stanley proving exactly that, and the uproar it caused. A bacterium is simply alive, with no qualification, and has been for three and a half billion years, which is roughly a thousand times longer than anything resembling a human has existed.

We have spent two hundred and thirty years, since a country doctor scratched a boy's arm, learning to defend ourselves against the first of those. We got good at it faster than anyone expected and more completely than anyone in 1796 could have imagined. The remaining problem, as chapter 22 argued at length, is not the science.

The second one is older, larger, and quietly taking back ground.

That is the next book.

Colophon

Almost Alive: How viruses were found, fought and sometimes forgiven was written between 15 and 16 November 2026.

Twenty-three chapters, about 75,200 words, in four parts. One hundred and thirty-one figures: sixty-eight photographic images, forty charts drawn from published data, and twenty-three labelled diagrams.

Every chart in this book is built by a script that reads a data file, and every data file carries a comment naming the source it came from. Both are kept with the manuscript, so any figure here can be rebuilt from scratch and checked against the record it claims to represent. Where a planned chart could not be built that way, it was replaced or dropped, and the chapter says which.

The text is set in Liberation Serif at 9.6 point on a 105 by 148 millimetre page. Chapter titles are small capitals between rules. The three kinds of box are set on a tinted panel. The jacket is set in Times.

The book was composed in Markdown, one file to a chapter, and built to PDF by a set of scripts kept alongside the sources. Nothing in the finished book is edited by hand after the build; the Markdown is the only original.

Written by Keith McLoughlin in collaboration with Claude, an AI assistant made by Anthropic. The preface says what that meant in practice.

For most of human history the things that made us ill were invisible, and nobody knew whether they were alive. A virus is genes in a coat. It does nothing on its own. Inside a cell it is the most successful thing on the planet.

This is how we learned to fight something we could not see, from a milkmaid's hands in Gloucestershire to a vaccine designed in two days in January 2020.

The science turned out to be the easy half.

