THE YEAR THE WORLD WENT MAD

THE YEAR THE WORLD WENT MAD

MARK WOOLHOUSE

Foreword by Matt Ridley

SANDSTONE PRESS

First published in Great Britain in 2022
Sandstone Press Ltd
PO Box 41
Muir of Ord
IV6 7YX
Scotland

www.sandstonepress.com

Editor: Robert Davidson

ISBN: 978-1-913207-95-3
ISBNe: 978-1-913207-96-0

Sandstone Press is committed to a sustainable future.
This book is made from Forest Stewardship Council ® certified paper.

Cover design by Heike Schüssler
Typeset by Iolaire Typesetting, Newtonmore
Printed and bound by TJ Books Limited, Padstow, Cornwall

This book is for my wife and daughter, the best
lockdown companions I could ever wish for.

Francisca, thank you for listening. You must have
heard everything in this book a hundred times before.

Nyasha, thank you for bringing joy to our
lockdown lives. I'm sorry that your generation
has been so badly served by mine.

CONTENTS

CHAPTER 1

SOUNDING THE ALARM

Early in the New Year of 2020 I was in my office in Edinburgh reading through media accounts of a puzzling respiratory disease – possibly a viral pneumonia – that had surfaced in the city of Wuhan in eastern China. I was concerned and tried to find out more.

On January 7th a helpful journalist sent me a copy of a report written by the Wuhan Municipal Health Committee. The report said that fifty-nine patients were suspected to have caught the mystery disease, seven of whom had become critically ill. The first cases dated back to mid-December 2019 and many were linked to the South China Seafood wholesale market in the north of Wuhan.

The same day the Chinese authorities announced that they had identified the cause of Wuhan pneumonia: it was a coronavirus.

A pandemic begins

As I'd been studying the emergence of new viruses for more than twenty years, I knew what to look out for. New human viruses usually come from animals, and most of them don't spread well between humans. Some coronaviruses can do though, which meant they were high on the list of viruses to worry about.

The fact that there were already fifty-nine cases in a single outbreak told me that this coronavirus probably did spread from person to person, so it was potentially a pandemic virus. If it was, it was probably already too late to stop it. I knew that a respiratory virus could spread around our highly interconnected world and seed a pandemic in a matter of days and I had just learned that this one had been spreading for weeks. There were still plenty of unknowns, but I was now very worried.

On January 12th Chinese scientists published the new virus's genome sequence – its genetic code. The genome sequence confirmed that it was a coronavirus and told us that it is closely related to the SARS coronavirus. This was more bad news. SARS is an extremely dangerous disease that killed more than seven hundred people in 2003. Only a prompt and vigorous international response had averted a full-scale pandemic. The possibility that SARS might one day re-emerge had been a concern ever since. Now we were facing a SARS-like coronavirus with unknown potential.

On January 21st the World Health Organization reported that there had been over two hundred cases and six deaths from the new virus in China, with further cases in Japan, South Korea and Thailand. We call a large but localised outbreak of disease an epidemic, but when an epidemic spreads to multiple countries across a wide region we call it a pandemic. This was not yet officially a pandemic but there could be little doubt that it was going to become one. The virus had already spread beyond China to three other countries and was most likely present in others that we didn't know about yet.

That same day I sent an e-mail to Catherine Calderwood, the Chief Medical Officer (CMO) of Scotland. Even though no cases had yet been reported in the UK, in that e-mail I said that we needed to start preparing for an epidemic that would affect the whole country, Scotland included.

Four days later, new data were published by the World Health Organization that prompted me to write again, with even greater urgency. I explained that – based on these new data – I estimated that the novel virus was capable of infecting more than half the population, tripling the mortality rate and overwhelming our National Health Service (NHS) within two months. I acknowledged that there was a lot of uncertainty but stressed how serious this could turn out to be.

I received a polite reply to my e-mails telling me that everything was under control. It wasn't, as we were all going to find out in the coming weeks.

Warning the Scottish government's most senior medical advisor that we were facing an unprecedented public health emergency wasn't something I did lightly. Before sending those e-mails to Catherine Calderwood I'd consulted with two colleagues who, like me, have been studying epidemics for many years. One was Jeremy Farrar, Director of the Wellcome Trust, and the other was Neil Ferguson, Director of the Centre for Outbreak Analysis at Imperial College, London. The three of us were in complete agreement and Jeremy and Neil were already in touch with the CMO England, Chris Whitty, and the Chief Scientific Advisor (CSA), Patrick Vallance.

In that flurry of communications in January 2020, we set out how we expected the next few months to unfold.

First, the pandemic would be fuelled by mild cases but with significant mortality in vulnerable groups.

Second, there was little prospect of a vaccine against a novel coronavirus becoming widely available in less than twelve months. Third, the prospects for effective therapies were not much better.

Fourth, case isolation, infection control and contact tracing would be crucial, but capacity to deliver them could be overwhelmed if case numbers rose too high. Fifth, social distancing measures such as restrictions on public gatherings and closures of workplaces and schools would then be needed.

Time was of the essence; every day it looked more certain that a pandemic was heading our way. I continued to share information with the CMO Scotland and brought Sheila Rowan – CSA Scotland – and Anne Glover – President of the Royal Society of Edinburgh and former CSA to the European Commission – into the conversation. I was told that the Scottish government was now working 'to address preparedness', though not what that meant in practice. I wasn't convinced. I'd have expected a lot more urgency if the government had been doing the same calculations that I'd been doing.

Three crucial numbers

To get a preliminary idea of how an epidemic could play out we need three numbers.

The first is called the basic reproduction number – this is a measure of how transmissible the infection is and allows us to estimate the fraction of the population who will be infected.

Next is the generation time, the interval between a person getting infected and infecting others – this sets the timescale.

Last is the infection fatality rate – this tells us how many people will die.

These three numbers drove my initial estimates of the potential scale, speed and severity of a novel coronavirus epidemic in the UK. None of these crucial numbers was known precisely in January 2020, but the most uncertain of all was the infection fatality rate. The infection fatality rate is a ratio of deaths to infections. In the early stages of the epidemic in China it was entirely possible that some deaths were being missed. However, there had to be far greater uncertainty about the number of infections. At that time, only clinical cases were being counted, not milder cases or those with few or no symptoms at all. The more of those there were, the more the infection fatality rate would be overestimated.

On January 25th the World Health Organization published an infection fatality rate estimate of almost 5% – meaning that one in twenty of those infected would die. If the true value was anywhere near that figure we were facing a catastrophe. Over the next few months, as surveillance of mild cases improved, the estimate would eventually fall to around 1%. That was bad enough, much higher than influenza which typically has an infection fatality rate less than 0.1%.

All of this told me that this pandemic could be much more severe and much harder to control than the flu pandemic the UK had spent years preparing for. That is why I wasn't convinced by the reassurances I was receiving. I felt we already knew enough to take even stronger action. Giving advice was the easy part; getting anything to happen proved a lot more difficult.

Gearing up

There was a lot more talking than action in the weeks following that first flurry of e-mails.

The UK government's Scientific Advisory Group for Emergencies (SAGE) met on January 22nd. I heard that it was a surprisingly low-key meeting given the gravity of the situation.

I was hurriedly appointed to a SAGE sub-committee called the Scientific Pandemic Influenza Group on Modelling (SPI-M). SPI-M met on January 27th to discuss the need for mathematical modelling of a novel coronavirus epidemic in the UK.

By the end of the month questions about the novel coronavirus had been tabled in both the UK and Scottish parliaments, so the new virus was on our politicians' radar too.

I finally met with the CMO Scotland on February 28th and at that meeting gave an update on possible scenarios for the UK's novel coronavirus epidemic. The main messages hadn't changed since January, but we now had more data to support them.

The Scottish government subsequently formed its own expert committee, the Scottish Covid-19 Advisory Group, informally known as SAGE-for-Scotland. I was asked to join the group but we didn't hold our first meeting until March 26th, three days after the UK went into lockdown, by which time the course of the epidemic in Scotland and the UK as a whole was pretty much set.

The lack of urgency was troubling. Scientists were warning in mid-January 2020 of an imminent epidemic on a scale we had not seen since Spanish flu more than a hundred years ago. This was an extraordinary event that demanded an extraordinary response and got a very ordinary one.

It was a different story in Taiwan. The Taiwanese government introduced rigorous health checks for arrivals from Wuhan on December 31st 2019, long before most people in the UK had heard about the new virus. By the end of January, Taiwan was screening all international arrivals for signs of infection. At the same time, isolation of cases and their contacts plus quarantining of anyone deemed at high risk was made mandatory and was strictly enforced.

The UK is not Taiwan but, even so, when the alarm was first raised here we could have reacted with the same sense of seriousness and urgency. We didn't. I can think of several reasons why.

For a start, there was a lack of global leadership. In this kind of situation, that is the role of the World Health Organization, but it failed to live up to one of its prime responsibilities. It didn't even declare a Public Health Emergency of International Concern – a precursor to declaring a pandemic – until January

30th. It didn't declare a pandemic until well into March. This badly undermined the case being made for early action in Scotland, the UK or anywhere else.

Another factor was complacency. Since the swine flu pandemic in 2009, the emergence of a new strain of influenza had been on the UK government's national risk register and we had detailed and rehearsed plans for responding to it. It was assumed that those plans would work for the novel coronavirus, but this wasn't flu and an even more vigorous response was required.

There was also a sense of déjà vu. In the early stages of the 2009 swine flu pandemic some scientists had confidently predicted a crisis on a scale far beyond what transpired, largely thanks to an initial overestimate of the infection fatality rate. Policy-makers might reasonably ask whether scientists were crying wolf again.

Those are all plausible explanations for the lack of action, but I think there was something more: sheer disbelief. We were asking officials and politicians to engage with a scenario lifted from a science fiction movie. They simply couldn't take it in.

Looking forward

For those who could take it in, it was already apparent that 2020 was going to be a traumatic year. Many people were going to die and life was going to change for all of us. The policy-makers may have been slow off the mark but at least the scientists did respond during those early weeks.

My research group at the University of Edinburgh was ideally positioned to study this new threat because we do research on the epidemiology of emerging viruses, meaning that we study their origins, distribution and spread. Accordingly, I begged and borrowed the funding to expand and reconfigure my team to work on the novel coronavirus. The same thing was happening in thousands of research groups in universities, government laboratories and industry around the world. Scientists working on diagnostics, drugs and vaccines got to work as soon as the new virus's genome was published in January.

Beyond scientific research, I had a good idea of what the imminent pandemic would ask of me and my team. I'd been deeply involved in the UK's response to previous epidemics: BSE (mad cow disease) in 1996; foot-and-mouth disease in 2001; swine flu in 2009. I was familiar with the sometimes fraught relationship between science and policy. I'd worked with the media on communicating the science to a public anxious to understand what was happening. I'd experienced the pressures that come when the stakes are high and there is heated debate about how best to proceed.

I knew in January 2020 that I was about to go through all that again and more. I knew it would be the same for many of my colleagues in the UK and around the world. I hoped that science would rise to the occasion and that we scientists could make a difference. What I did not expect was that elementary principles of epidemiology – my own subject – would be misunderstood and ignored, that tried-and-trusted approaches to public health would be pushed aside, that so many scientists would abandon their objectivity, or that plain common sense would be a casualty of the crisis.

I did not expect the world to go mad, but it did.

CHAPTER 3

MODELS

Scientists use mathematical models to study complex biological processes at every conceivable scale from single molecules to ecosystems. An epidemic has fewer moving parts than an entire ecosystem but it's still extremely difficult to predict how it will play out, even if you're an expert in public health. Models can help.

Epidemic maths

An epidemic is the consequence of infection being transmitted from person to person. The crucial characteristic of this process is that the more infections there are in the population the greater the risk that an uninfected member of that population gets infected. The technical term for this characteristic is positive feedback and it is why epidemics grow exponentially.

Exponential growth is multiplicative (for example, 1, 2, 4, 8, 16, . . .) rather than linear (1, 2, 3, 4, 5, . . .) – it will be a recurring theme in this book. An epidemic can't grow exponentially for ever – sooner or later it has to slow down because there are fewer people left to infect – but it can do so in the early stages, as we would soon see with novel coronavirus.

Numbers of cases of non-communicable diseases like stroke or diabetes do not behave this way. For those diseases we sometimes find clusters of cases within families or in a local area but there is no person-to-person spread, so no positive feedback and no exponential growth.

Positive feedback makes it hard to say what will happen if you put in place an intervention designed to reduce the rate of transmission, such as lockdown. If, for example, a lockdown halves the transmission rate (by halving the number of people we come in contact with) it's not obvious what that would do to the size of the epidemic. We'd expect it to be smaller, but would it be half the size, or more, or less? Models are good for answering that kind of question. (The answer, by the way, is that sometimes halving the transmission rate has little impact on epidemic size, sometimes it has a big impact and sometimes it stops the epidemic from taking off altogether – it depends where you're starting from. Epidemics are complicated.)

This is much more than an academic nicety. If those commenting or advising on the response to novel coronavirus do not understand the dynamics of epidemics – and many clearly did not – then their comments or advice can be misleading. Throughout 2020 we saw one example of this after another in public discussions of the R number, herd immunity, elimination, travel bans and the second wave. The resulting confusion clouded those important and necessary debates about what to do next and increased the likelihood of getting crucial decisions wrong.

My favourite example of epidemiological models outperforming expert opinion involves the work of a pioneer in the field, Roy Anderson of Imperial College. Roy and his colleagues modelled the future scale of the HIV/AIDS epidemic in the late 1980s when it was a new and still relatively rare disease. Their predictions of millions of deaths globally were ridiculed by public health 'experts' who failed to grasp that – unlike more familiar infections such as flu – this would be a long, drawn-out epidemic and difficult to stop. As the world now knows, Roy's models were right and the critics were wrong.

SPI-M

For all these reasons, SPI-M – the modelling sub-committee of SAGE – played a pivotal role throughout the UK coronavirus epidemic. SPI-M was co-chaired by Graham Medley – who had invited me to join the committee – and Angela McLean, Deputy CSA. I knew Graham and Angela from when we all worked together in Roy Anderson's team at Imperial College in the 1980s. SPI-M is a large committee and the meetings were always lively. The discussions were well-informed and uninhibited, ideas were put forward and challenged, every session was a crash course in our rapidly expanding knowledge of coronavirus epidemiology.

In happier circumstances, working with SPI-M would have been a joy, but the meetings during 2020 always had an undercurrent of deeply felt concern. The whole point of the modelling was to help us glimpse the future, and none of us liked what we saw there.

Since the models were informing literally life-or-death decisions it was vital that the outputs and the recommendations coming from SPI-M were as robust as possible, and that any uncertainty was reported alongside the headline results. To achieve this, SPI-M uses 'ensemble' modelling.

Ensemble modelling works by consolidating the outputs of multiple, independent models rather than reporting only the output of a single 'best' model. If the models agree, we are more confident in our conclusions. If they disagree, we look for the reasons why and, in the process, may learn something about questions we still need to resolve. Ensemble modelling is a tried-and-trusted approach to modelling complex problems where there are many uncertainties. Climate change modelling works the same way.

SPI-M was a tremendous resource for the UK government in 2020, an assemblage of top scientists in the field, with a structure ready and waiting to deliver modelling outputs in real time as needed, working under the direction of the Department of Health and Social Care (DHSC).

That said, SPI-M frequently found itself at the centre of the storm. The models were (quite rightly) intensely scrutinised and strongly challenged. When, for reasons we'll come to later, trust in the models declined, there was plenty of negative publicity, not least in the November 2020 BBC2 television documentary *Lockdown 1.0 – Following the Science?*

SPI-M had to contend with one challenge right away; it was set up to tackle the wrong disease. That's apparent from its full name: the Scientific Pandemic Influenza Group on Modelling. The committee was created in the wake of the 2009 swine flu pandemic to model the much-anticipated next influenza pandemic (something we may experience in coming years).

For the most part, the influenza expertise in SPI-M was an asset; many of the epidemiological drivers – particularly human behaviour – are similar for novel coronavirus and for flu. Nonetheless, adjustments were needed and some of the models had to be re-written from scratch.

Even then, the new, bespoke coronavirus models betrayed their influenza pedigree. A good example is that the models included schools but not care homes. Schools were the major driver of the 2009 swine flu epidemic in the UK but they were not the major driver of novel coronavirus – this is one of many features of its epidemiology that are more like SARS than flu. Care homes, on the other hand, were crucial. A large proportion of deaths due to novel coronavirus occurred in care homes. As we were to learn, the course of the epidemic in care homes was distinct from that in the wider community – that needed to be modelled explicitly.

Another missing feature was shielding of the vulnerable. Shielding featured prominently in the government's response from March 2020 onwards, but most of the modelling of this intervention was done by independent researchers. Shielding surely had some effect on transmission to people most likely to die – that was what it was intended to do, after all. It was also likely that many people at risk were taking precautions of their own – independent of government advice. The influenza models didn't capture any of this.

A striking omission from the original influenza models was lockdown. A full lockdown was not part of the UK's planning for an influenza pandemic. The models included options for some elements of social distancing – notably travel restrictions and closing schools and workplaces – but not an instruction for most of the population to stay at home. Over the past ten years, SPI-M had built up an evidence base for the interventions we expected to use to control a pandemic. Now we were contemplating doing something different. We'd done our homework, but we'd prepared for the wrong exam.

The influenza legacy was a weakness, but the novel coronavirus models were still useful tools. They were flexible enough to answer the many different questions being asked by government, such as when cases would peak, how much intensive care unit capacity would be needed or what difference more testing would make.

I told a House of Commons Select Committee in November 2020 that decisions on how to respond to an epidemic like novel coronavirus should not be taken without input from models. As I've explained, they provide insights into how the epidemic might unfold that couldn't be obtained in any other way, particularly in the early stages.

That said, I wouldn't want decision-making to be over-reliant on models either. In a Good Practice Guide for modellers I co-authored in 2011 we stressed that models should only be one of the lines of evidence informing policy-makers, never the only one.

In March 2020, however, you could easily get the impression that the UK government's mantra of 'following the science' boiled down to following the models. That's how it looked and that's how the media presented it.

The R number

One of the most prominent outputs of the models was an estimate of the R number. The formal term for R is the 'case reproduction number' and it has a simple definition: R is the average number of cases generated by a single case. The single case is referred to as the index case, and the cases that it generates are known as secondary cases.

The R number is a measure of how well the infection is spreading in the population. R greater than one means that, on average, an index case generates more than one secondary case and the number of cases will grow. R less than one means that, on average, each case generates less than one case and the number of cases will decline. R equals one means that the incidence of new cases stays the same. The phrase 'bringing the epidemic under control' implies reducing the R number from above one to below one.

For a given infection in a given population, the maximum possible value of R is called the basic reproduction number, written as R0 and pronounced 'R nought'. For novel coronavirus in the UK in March 2020 R0 was about three. At that stage of the epidemic, the case reproduction number, R, was equivalent to the basic reproduction number, R0, but once steps were being taken to reduce transmission rates the R number fell and it would fluctuate between 0.7 and 1.5 for the rest of the year.

SPI-M put a lot of effort into those estimates of the value of R. The outputs of these calculations were reported by DHSC as the government's official weekly R number. The R estimate was never a single value, always a range, the size of the range reflecting the degree of uncertainty in the true value. There is inevitably some uncertainty because, in the absence of detailed and systematic tracing data telling us directly how many secondary cases are generated by an index case, we can only measure R indirectly, estimating it as best we can from the epidemiological data available.

Though R is important to infectious disease epidemiologists, it isn't that useful as an operational public health tool.

R doesn't tell you about the number of infections in the population – that's the prevalence – or the number of new infections – that's the incidence.

R only refers to infections, so it doesn't tell you about the numbers of hospitalisations and deaths, which matter most to the NHS and government.

Nor does R help anyone understand their individual risk, which matters most to you and me.

R does tell you about the trajectory of the epidemic but not how fast the epidemic is growing or shrinking. You get that from the doubling time, which I've long argued is a more useful number to communicate than R.

And on top of all that, the methods SPI-M used gave an estimate of R over the past two weeks or so rather than its current value, so it was out of date by the time it was published.

Given these limitations, I was concerned when 'keeping R below one' became a policy objective. SPI-M's weekly estimate of the R number was hugely influential, widely reported in the media and often discussed by politicians and commentators, most of whom had never heard of R before the pandemic began and clearly didn't understand it. As one of my colleagues put it: we've created a monster.

The R monster turned out to be quite dangerous. The policy objective of keeping R below one was often expressed as 'suppressing the virus' and became the central justification for lockdown. The problem is that suppressing the virus is not the only nor the most direct way of minimising the public health burden. The relentless focus on the R number detracted from the usual public health priorities of saving lives and preventing illness.

Herd immunity

Another feature of the models that was widely discussed but poorly understood was herd immunity. Here's how it works. When we are infected by a virus our immune system responds, clearing the infection and creating an immunological memory that protects us against reinfection. Immunity is important to individuals, and it is important to the population as a whole as well. If some people are immune to infection then it is more difficult for a virus to spread through the population, which means the R number will be lower.

The fraction of people that are immune tells us the level of herd immunity. If a large enough fraction is immune then the R number is reduced to one without any additional interventions. That fraction is called the herd immunity threshold.

I've explained what herd immunity is, but we need to be equally clear about what it is not. Herd immunity does not mean that everyone in the population is immune to infection. Nor does it mean that the virus cannot transmit at all in the population (which is why many find the alternative term 'population immunity' confusing, so I will stick with 'herd immunity' in this book). If the herd immunity threshold is passed, a large epidemic is not possible but there may still be outbreaks. The good news is that – because R is less than one – any outbreaks will be self-limiting, meaning that they die out of their own accord. Most of these self-limiting outbreaks will be small, though if R is not much below one they could stretch to hundreds of cases.

The fraction of people that must be immune to achieve herd immunity is related to the basic reproduction number R0, which we can think of as the maximum possible value of R. A crude but useful estimate of the herd immunity threshold is given by the formula: 1 minus 1/R0. If novel coronavirus has an R0 of around three then the formula tells us that the herd immunity threshold is about two-thirds, or 67%. Measles has an R0 of around ten and so the herd immunity threshold is about 90%. Swine flu had an R0 of around 1.5, giving a herd immunity threshold of about 33%.

I must stress – because some commentators were confused on this point – that the herd immunity threshold is *not* an upper limit to the number of people who would be infected during an uncontrolled epidemic. That number – the attack rate – could be considerably higher.

With this in mind we can turn our attention to the heated debate about the role of herd immunity in the UK's epidemic response. In February and March 2020, there was lively argument about the pros and cons of a 'herd immunity strategy' for tackling novel coronavirus. It was a difficult debate to follow for a couple of reasons.

First, the concept of a 'herd immunity strategy' was completely new – the term didn't exist before 2020, so there was no consensus about what it meant.

Second, no-one could argue against herd immunity – that would be irrational, herd immunity works in our favour. Opponents of a herd immunity strategy were really arguing for the strongest possible suppression of the virus. The conflict arises because suppressing the virus has the effect of slowing the build-up of herd immunity by preventing people becoming infected in the first place.

This tension between the short-term benefits of suppressing the virus and the long-term benefits of herd immunity was unavoidable, whatever we did. Epidemiologist John Edmunds put it neatly when he said that 'herd immunity is how this will end, sooner or later'. He meant that the epidemic would continue until enough people had either been infected or vaccinated.

This is why Patrick Vallance and other government advisors – myself included – were talking about herd immunity in March 2020. It was important that we did so. Admittedly, some public pronouncements on the topic were poorly phrased, allowing critics to claim that one of the policies being considered was not to suppress the virus at all. That extreme 'do-nothing' scenario was modelled but – as we shall see shortly – it was never a policy option. Quite the reverse, the do-nothing scenario was modelled to make the case for intervention.

Regardless, Richard Horton – editor of *The Lancet*, a leading medical journal – and others continued to rail against their straw man version of a herd immunity strategy. One of my colleagues was struck by the irony of this, making reference to *The Lancet*'s publication back in 1998 of a notorious paper that raised doubts about the safety of the MMR vaccine. The paper was deeply flawed, and was eventually retracted, but the alarm it caused led to reduced uptake of MMR and the UK almost losing its herd immunity to measles, a public health disaster only narrowly averted.

The debate descended into farce. Some said they disagreed with the idea of a herd immunity strategy because they disapproved of people being likened to a 'herd' – a nadir for the level of public discussion during the pandemic. On one occasion, one of my colleagues was clearly taken aback when asked on BBC Radio 4's *Today* programme if he 'believed' in herd immunity. That's like being asked if you believe in tides. Like the tide, herd immunity happens whether you believe in it or not.

One of the ironies of this debate was that most epidemiologists thought that herd immunity was almost irrelevant in March 2020. That's because too few people had been infected – and so might now be immune – to make any discernible difference to the trajectory of the epidemic.

There was a minority view that this was wrong, that many more people had been infected than we thought because the virus had been circulating for longer than we were assuming. This idea was fuelled by a trickle of anecdotal reports from people who claimed to have had Covid-19-like symptoms in January 2020 or even earlier.

Serological surveys – which test for antibodies against the virus and so indicate how many people have had the infection – eventually put paid to this theory. Data published from April 2020 onwards indicated that no more than 10% of the population had been infected in London, and only around 5% in other parts of the UK. This tallied with expectations from my team's mathematical models and didn't support the idea of large, hidden epidemic.

The UK never adopted a herd immunity strategy, but the debate left a legacy of misunderstanding. Months later, I was taken aback when some research done by my team was criticised as being based on a 'herd immunity model'. Of course our model included a representation of herd immunity. That didn't make it controversial, it made it a standard, text book epidemiological model of a human viral infection. All the novel coronavirus models incorporated herd immunity – they would be deficient if they didn't.

Using models to make predictions

Epidemiological models have various uses. They are often used to interpret epidemiological data; for example, helping to assess in retrospect the impact of interventions. They can also be used to make predictions.

Prediction was challenging for novel coronavirus in the first half of March 2020. Accurate prediction of an epidemic is hard enough even for infections we know a lot about, and there were still many gaps in our knowledge of this new virus. The models were not data-free – there was information from China and a few other countries – but we were at too early a stage of the UK epidemic to calibrate the models against UK data. That left plenty of scope for uncertainty.

The role of herd immunity was one area of uncertainty. I said earlier that herd immunity is an integral feature of virus epidemiology, but there's a complication. For some infections – including some coronaviruses – immunity is not 100% protective and wanes over time. Until we knew whether either of these applied to novel coronavirus we couldn't make any useful predictions of the long-term course of the epidemic. Fortunately, this wasn't such a problem for short-term predictions – as we saw earlier, herd immunity didn't play a significant role in the early stages of the epidemic.

The biggest challenge to making short-term predictions was the role of human behaviour. The spread of any respiratory virus depends on how people behave, especially the number and nature of the contacts they make with other people.

We know a lot about people's contact behaviour in normal circumstances. As part of the preparations for pandemic influenza there had been several large-scale studies of volunteers recording who met with whom and how often. These data were an essential input into the epidemiological models. However, we were now expecting those behaviours to change; that was the whole point of social distancing measures.

The problem was that we didn't know exactly *how* those behaviours would change. That would depend on how well people complied with advice and regulations, on what they did instead of doing the things they'd been told not to do, and on how they changed their behaviour of their own accord, independently of government directives. Basically, we were being asked to predict how people would respond to a once-in-a-lifetime crisis and there was no sound basis for doing so. The best we could do was guess.

If I had to pick anyone to guess how people's behaviour might change during a pandemic then I would choose a group of scientists who have been studying this kind of problem for many years – I would choose SPI-M. At the same time, it's important not to overstate how much confidence we should have in those guesses. The best we can do is try out different assumptions – assumption is a more scientific-sounding term for guess – about how people might behave and proceed from there.

How people will behave is not the only uncertain aspect of the epidemic. We might also want to explore different assumptions about immunity, the transmission rate, the generation time, the infection fatality rate and every other component of our model. On top of that, we will often want to compare the expected effects of different interventions or mix of interventions.

I shall use to the word 'scenario' to refer to a single combination of assumptions and modelled interventions. A scenario is a list of 'if' statements: if the virus does this . . . and if people behave like this . . . and if the government does this . . . then we expect...

You will immediately see that at any point in the epidemic there could be lots of possible future scenarios to consider.

With modern computers that's not a problem – we can look at a huge number of scenarios extremely quickly – my team sometimes looked at a million in a single modelling study. Even then, we didn't necessarily expect any of the million to turn out to be precisely correct. We were much more interested in what the exercise told us about the range of possible outcomes. Where there was consistency we could be reasonably confident that we knew (at least roughly) what was likely to happen. On the other hand, if there was a wide range of possible outcomes and we had no good basis for preferring some scenarios over others then the wisest course was to say that we don't know what will happen and wait for more data.

Use and misuse of the reasonable worst case scenario

Mind you, even if the future is uncertain, governments and health care providers still have to plan for it. One useful planning tool is a scenario called the reasonable worst case. This scenario isn't the most likely outcome – it's at the extreme end of the plausible range where cases, hospitalisations and deaths are highest – but you can't categorically rule it out. The reasonable worst case boils down to 'hope for the best but plan for the worst'. If you plan for it – for example, in terms of building extra hospital capacity to deal with a possible surge in severe cases – then you shouldn't get caught out.

One problem with the reasonable worst case scenario is that the media cannot resist treating it as a prediction. There was an outcry in 2019 when 'Operation Yellowhammer' was made public. This was a planning exercise based on a particularly gloomy scenario for the state of the UK in the weeks following a 'no-deal' departure from the European Union. It was treated by the media as though it were a prediction and got plenty of attention. The reasonable worst case is not a prediction – we'd be (unpleasantly) surprised if it happened – it's a planning tool.

That brings us to Imperial College's Covid-19 Response Team Report 9, published on March 16th 2020. The model used to produce this report generated a worst case scenario of over half a million deaths due to novel coronavirus in the UK by the end of July. This wasn't exceptional: all the models suggested that novel coronavirus was capable of causing a huge epidemic. I'd shared such a scenario with the Scottish CMO back in January. The problem was that these worst case scenarios weren't realistic and weren't intended to be.

The worst case epidemic is one in which the virus is allowed to spread in the absence of any countermeasures at all. This is the do-nothing scenario we encountered earlier. You might call it the *un*reasonable worst case scenario. It's useful to know how big that epidemic would be but it would never happen in practice. It is inconceivable that we would all carry on as normal while hundreds of thousands of people were dying, even in the purely hypothetical event that government did nothing at all.

Report 9 went on to look at the effects of different combinations of interventions – self-isolation of cases and household contacts, school and university closures, and full lockdown – over a period of two years. I was extremely sceptical about this. It was perfectly obvious that no-one could predict the course of this epidemic over such a long timescale, so what was the point of publishing these outputs?

What's more, the outputs the study wasere based upon a single set of assumptions – or, if you prefer, guesses – about the impact of each intervention. That was a concern because no-one could be sure what those impacts would be in practice, introducing massive uncertainty into the analysis.

Even if Report 9 wasn't a good justification for a full lockdown, it did contain one important result: if we tried to tackle the novel coronavirus epidemic mainly by social distancing then we were in for a torrid time. The report should have been a wake-up call that we needed to invest quickly and heavily in other ways to control novel coronavirus or – according to the model – we'd end up in lockdown. This implication was barely mentioned – lockdown was accepted as a necessity the first time it was proposed.

When Report 9 was published the details of the scenarios modelled were quickly forgotten, as were any mentions of the assumptions, caveats and uncertainties of the analysis. Report 9 was condensed to the simple but misleading message that, if the government didn't impose full lockdown immediately, over half a million people would die.

My last on-camera TV interview before lockdown was an attempt to alleviate some of the huge public anxiety created by Report 9. The BBC interviewer seemed genuinely rattled by the frightening 'prediction' of so many deaths. I tried my best to explain that this was the kind of exercise that mathematical modellers undertook to understand the epidemic better, it wasn't realistic and it wasn't a prediction of what was going to happen come what may. I don't think I succeeded.

Defining the strategic objective

Half a million deaths may have been implausible but we were still facing a public health emergency. The challenge facing SPI-M in February and March was to work out how best to meet the UK government's stated policy objectives: to save lives and protect the NHS. This wasn't as straightforward as it might sound.

Let's start with the first objective, saving lives. If that objective is interpreted as trying to minimise deaths due to novel coronavirus while ignoring deaths from other causes, and if social distancing is the intervention of choice, then we don't really need a complex computer model to tell us what to do. The best solution is to lock down immediately, as tightly as possible, and wait until some alternative – such as a vaccine or effective treatment – presents itself, regardless of how long that might take. There was no shortage of people advocating this approach, but none of the UK administrations pushed the 'saving lives' objective to its logical conclusion.

The second objective was to prevent the NHS from being overwhelmed. This implies that an R number above one might be manageable as long as the resulting epidemic was not so large that the NHS couldn't cope, which was usually interpreted as not running out of intensive care capacity. Prime Minister Boris Johnson characterised this strategy as 'flattening the curve' and 'squashing the sombrero'. Unfortunately, there were two problems with the approach.

First, we come up against one of the paradoxes of the dynamics of epidemics. As you'd expect, a higher R number results in a larger epidemic with a higher peak number of cases and, therefore, a higher peak number of hospitalisations and deaths. An R number that is lower (but still above one) results in a smaller epidemic with a lower peak. That seems straightforward enough. Unfortunately, the smaller epidemic is also more prolonged. That's not obvious at all – as we saw at the beginning of this chapter, intuition is often a poor guide to how an epidemic will play out. If the epidemic was going to be a long-drawn-out affair then we needed to think not only about the NHS's surge capacity but also about sustainability.

That brings us to the second problem: the NHS does not have a great deal of spare capacity at the best of times. Whether there was a huge surge in demand for a limited period or a more modest increase spread out over many months, the NHS couldn't cope with either.

This meant that flattening the curve was only ever a realistic strategy if it were linked to the provision of both additional infrastructure – more beds and more intensive care units – and many more trained staff. We got the infrastructure – in the form of the Nightingale and Louisa Jordan hospitals – but staff numbers remained limiting.

Looking back, it seems perverse that the UK tried to combat a major public health emergency without making a more determined effort to boost health service capacity. This was one of many instances where we were handicapped by a failure to appreciate the seriousness of the situation we were facing and a reluctance to accept that the crisis would not be over in a few weeks.

By default, SPI-M was left with trying to identify a strategy which was more than flattening the curve but less than all out suppression – not an easy task. I recall a heated discussion of whether it was possible to maintain the R number at round about one. I described this as 'utterly unrealistic' on the grounds that we had no way of knowing what combination of countermeasures would be required. This was when I first realised that we couldn't try to control this epidemic simply by managing the rate of transmission. There was too little room for manoeuvre. We'd end up yo-yoing between intolerably severe restrictions and unsustainable pressure on the NHS (and that, of course, is exactly what happened).

My team started looking for alternatives right away, but it was to take a couple of months before we had a firm proposal to bring to the table. In the meantime, there was a good argument that it was better to err on the side of caution – if the NHS were truly overwhelmed the consequences would be horrendous, and not only for patients with Covid-19. I fully agreed with this – I had highlighted the danger of the NHS being overwhelmed back in January. The problem was the proposed solution: lockdown.

CHAPTER 6

DANGEROUS LIAISONS

One of the first things we want to know about any infection is how you catch it. Novel coronavirus is a respiratory virus, like influenza virus and like the coronaviruses that cause common colds. All these viruses replicate in the upper respiratory tract and virus particles are expelled in droplets or aerosols as we exhale, cough or vocalise. We catch novel coronavirus by inhaling these airborne virus particles.

The airborne route is believed to be by far the most important way that novel coronavirus spreads but there are others. The virus can persist on surfaces – hence the advice for frequent cleaning early on in the pandemic – but it is now thought that this is not a major route of transmission. It may also be present in faeces (many respiratory viruses are also found in the gut), though that doesn't seem to be a common transmission route either.

Defining a contact

There are echoes of the novel coronavirus's evolutionary history in how and where it transmits most easily between humans. Novel coronavirus is thought to be descended from a bat coronavirus and its ancestors will have spread among huge colonies of bats roosting in close proximity in caves or mine shafts. Knowing this gives us some clues.

First, the ideal environment for the transmission of novel coronavirus is a poorly ventilated, enclosed space where it is still and dark, away from direct sunlight – 'cave-like' would be a fair description, but indoors works just fine.

Second, the virus spreads most efficiently wherever people are gathered together in close proximity for a prolonged period of time – much like bats roosting.

Third, the virus transmits especially well between people who are talking, shouting or singing – bats are noisy too. In normal circumstances, most people spend a lot of time indoors, close to others, and vocalising. This explains why – having made the jump from bats – novel coronavirus was able to thrive in human populations; it was predisposed to do so.

Understanding the biology of novel coronavirus helps us understand the concept of a 'contact'. A contact is any interaction that provides an opportunity for the virus to spread from one person to another. The World Health Organization defined a contact in terms of distance apart and the time spent within that distance.

We don't have to be in direct contact with someone to catch novel coronavirus, but the further apart the better. There was early evidence that the risk of transmission was reduced by around 90% by being two metres apart, and 80% by just one. We don't have to be within that distance for long, but the fact that novel coronavirus spreads best where the virus is able to build up in still air implies that longer contact makes transmission more likely. The World Health Organization advised that a contact has to last at least fifteen minutes for there to be significant risk.

So, for practical purposes, a contact between people involves being no more than one or two metres apart for at least fifteen minutes. This is the basis for guidance to reduce the risk of transmission by physical distancing. The World

Health Organization initially recommended that people keep two metres apart wherever possible, later reducing this to one metre. The UK government settled on 'one metre plus'. The time dimension tended to get forgotten when guidance was issued, but it matters too.

Of course, not all encounters that satisfy the definition of a contact are equally risky. There will be a big difference in risk between spending fifteen minutes in a well-ventilated room wearing face coverings and sleeping in the same bed for an entire night – though both meet the criteria for a contact. Nor is anyone claiming that there is zero risk at five metres distance or from a fleeting encounter, just that the risk is too low to worry about. Not everyone accepted this, however.

The BBC News was fond of showing an animation illustrating how a single cough could infect everyone in a small restaurant. I complained to David Shukman – BBC Science Editor – about this. He replied politely enough but persisted with the animation. These are impossible arguments to win. I couldn't categorically state that such a thing could never happen, only that it had to be exceptionally rare. If novel coronavirus were so infectious that a single cough could routinely infect a roomful of people – or a bus-full or shop-full – then it would be unstoppable.

Novel coronavirus spreads particularly well within households. The bigger the household, the greater the risk. Visitors – especially overnight visitors – are just as problematic. This is one of the features of novel coronavirus that makes it exceptionally difficult to control. Workplaces and public spaces can be made safer, but there is little that governments can do to manage spread within our homes; that's up to us. What government can do is put an outright ban on social contact between households, which they did.

By far the safest place for people to get together is outdoors. As far back as March 2020 there was evidence from China that outdoor transmission of novel transmission was extremely rare. One study traced over seven thousand transmission events; just two occurred outdoors. For whatever reason, this knowledge did not influence UK government policy when it came to lockdown; we were told to stay at home and not meet others, even outside. There was no need for this, as long as physical distancing is observed the risk of transmission outdoors is negligible. Yet not only were sports such as golf and tennis banned – despite the low risk – but at one stage the police were even harassing solo hill walkers. That has to be one of the safest activities imaginable.

As the summer of 2020 wore on there were repeated outcries about crowds on beaches. People were castigated for their 'irresponsible' behaviour by politicians, police officers and public health experts. We were told to expect surges in cases. These never materialised, which was no surprise to anyone familiar with the epidemiological data. To my knowledge, no novel coronavirus outbreak has been linked to a beach anywhere in the world, ever.

There are a couple of caveats though. Intimate contact still poses a risk outdoors. Outdoor mass gatherings are problematic too. Large numbers of concert goers or sports fans sit or stand close to one another for a prolonged period and are often noisy (as we have seen, novel coronavirus transmission is linked to vocalisations). Even outdoors this combination may pose a risk, although as mass gatherings were one of the first activities to be banned there is little hard evidence. We should also be worried about so-called pinch points at such events – toilet facilities, indoor refreshments, and travel to and from the venue by car or public transport. All of these provide opportunities for the virus to spread.

Non-pharmaceutical interventions

There are three ways of preventing the spread of novel coronavirus by managing person-to-person contact. These are collectively referred to as non-pharmaceutical interventions because they do not involve any kind of medication.

First, there is a suite of measures that are intended to help reduce the risk of transmission associated with day-to-day activities. I call these Covid-safe measures and I'll describe them in more detail shortly. The point for now is that these interventions are about making contacts safe rather than not making contacts at all.

Second, there is self-isolation of cases and their contacts, an intervention that only affects people who have or are likely to have a novel coronavirus infection, thus minimising disruption to everyone else.

In contrast, the third approach – social distancing – is about reducing the number of everybody's contacts. The ultimate social distancing measure is lockdown.

Covid-safe measures include basic respiratory hygiene and hand-washing or sanitising. These measures were promoted by the UK government from the outset. Physical distancing, including the use of screens, was also quickly implemented where it was possible to do so. (Confusingly, this was often referred to as 'social distancing' – a term better reserved for measures such as banning household mixing, which is how I use it in this book.) Good ventilation improves Covid-safety, but fresh air is even better. Screening people for a high temperature was not recommended by SAGE and was not widely used in the UK, though it was adopted by some other countries.

One peculiarly controversial Covid-safe measure is wearing face coverings. It took some time for the UK to embrace this idea, which was odd because masks were always required for health care workers (who were, of course, at increased risk of exposure to the virus). There was speculation – but no evidence that I know of – that face coverings might engender a false sense of security leading people to behave less responsibly in other ways. This was akin to the argument in the 1980s that people wearing seat belts might drive more dangerously.

By July 2020 the advice changed and we were told to wear face coverings in enclosed public spaces. The UK was a bit slower off the mark than the World Health Organization, but even they didn't advocate general use of face coverings until early June.

Covid-safe measures do two things at the same time: they help protect you from any of your contacts should they be infectious, and they help protect your contacts from you should you be infectious. This makes it more difficult for the virus to spread and so reduces the R number. If we could reduce the average risk of transmission per contact by two-thirds (enough to reduce the R number to one from its natural maximum of three) using Covid-safe measures then we wouldn't have to do anything else to keep the virus under control.

To put this another way: if you could make contacts safe enough then there would be no need to cut the number of contacts you make, no need for social distancing and no need for lockdown. Unfortunately, it's not as easy as that. Relying on hygiene, physical distancing and face coverings alone is not enough to reduce the R number for novel coronavirus below one. Fortunately, we don't need to turn to social distancing right away, there are other ways to stop the virus spreading. First, though, we have to find it.

Outbreaks and super-spreaders

'Outbreak' is a familiar word (there was even a 1995 Warner Bros film of that name) but it is a technical term too. Public Health England defines an outbreak as two or more epidemiologically linked cases. That's a strict definition – most people don't think of two cases as an outbreak – but it is a standard and widely used one. 'Epidemiologically linked' means that public health workers can identify a time and a place where infection could have been transmitted – sharing a house would be an easy example. Two cases are not many though and here I will focus on larger outbreaks.

In 2020 epidemiologists at the London School of Hygiene and Tropical Medicine assembled a database of novel coronavirus outbreaks reported from around the world. By July they had amassed a list of well over two hundred, the majority with fewer than twenty cases but a few with over a thousand.

One striking feature of the London School's list is that few outbreaks occur outdoors. The ones that do are mostly associated with places of work – large farms or building sites – which raises the possibility of other links between the cases, such as shared eating places or sleeping accommodation. Consistent with what we know about the conditions that favour transmission, the great majority of large novel coronavirus outbreaks occur indoors in settings such as churches, hotels, factories, ships and care homes.

Digging a little deeper, we find different kinds of outbreak in the list. Some are linked to groups of people that frequent the same location, such as a workers' dormitory, university residence, school or workplace. Others are linked to one-off events, such as a wedding, bus trip or conference. This tells us that an outbreak results from a combination of setting, behaviour and perhaps characteristics of the individuals involved. Which brings us to super-spreaders.

The notion of super-spreading isn't new. In 1997 I wrote a paper showing that – for a variety of infectious diseases – some people contribute much more to transmission than others. A useful rule of thumb is that 80% of transmission is attributable to 20% of the population. A paper published in the journal *Science* in November 2020 reported that this rule worked well for novel coronavirus too. By analysing contact tracing data the researchers found that 20% of index cases were responsible for even more than 80% of secondary cases. We would like to know what is special about that small minority of cases. There are several possibilities.

First, some people will be more infectious because they shed larger quantities of virus, perhaps over an extended period of time. This is reminiscent of the sad story of Typhoid Mary who – back in the early twentieth century in the days before antibiotics – could not be cured of a persistent infection with typhoid-causing bacteria and was forced to spend much of her life in isolation to protect others. Persistent infections with novel coronavirus can occur – especially in people who are immunocompromised – but they are not common and, thankfully, we surely wouldn't react in the same way. But some variation in infectiousness is bound to occur, as it does for typhoid, influenza and many other infections.

Second, some people may have more contacts to whom they could pass on infection. A high number of contacts might, for example, be linked to working in public transport, hospitality, education or retail. For this reason, it might be a good strategy to target interventions at people in high-risk occupations or with high-risk behaviours.

The final possibility is that there is nothing special about the 20% other than the case happens to be – from the virus's point of view – in the right place at the right time. This brings us back to those one-off events that are associated with novel coronavirus outbreaks; I mentioned weddings, bus trips and conferences but it could be any occasion where large numbers of people spend time together in a confined, poorly ventilated space. Anyone present who happens to be highly infectious at the time – say, in the twenty-four hours before the onset of symptoms – could become a super-spreader.

In practice, where super-spreading has occurred – in a French ski chalet, an Israeli school or even (as has been claimed) the inauguration ceremony for a US Supreme Court judge – it is likely that a combination of all these factors played a role. For this reason, I have always been reluctant to apply the term 'super-spreader' to an individual. Super-spreading isn't just about an individual, it's also about how the individual fits into the wider epidemiological picture, who they have contact with and in what circumstances. I prefer the term super-spreader event.

The role of contact tracing

Ultimately, whatever factors are responsible for an outbreak the public health imperative is to contain it. That involves a practice called 'outbreak investigation', a mainstay of the public health response to infectious diseases that can be dated back to John Snow's landmark study of a cholera outbreak in London in 1854. Outbreak investigation was renamed 'cluster analysis' during the novel coronavirus pandemic. I couldn't see any reason for this, unless someone thought a nineteenth-century public health tool sounded too old-fashioned for a twenty-first-century pandemic.

Outbreak investigation combines identifying the source of the outbreak – a procedure known as backward tracing – with trying to identify those who may have been exposed during the outbreak – forward tracing. An important feature of tracing is that it is much easier to do when numbers of cases are low. Once there is widespread community transmission every outbreak is quickly absorbed into the rest of the epidemic and outbreak investigation loses a lot of its value.

Contact tracing is another routine public health activity. It is equivalent to forward tracing in an outbreak investigation but applied to a single index case. The aim is to identify every person who might have been infected by that case. In the UK, contacts of novel coronavirus cases were asked to self-isolate for fourteen days (reduced to ten days towards the end of 2020) and to report if they developed symptoms, at which point their contacts will be traced too. Members of the same household are almost always designated contacts. That makes sense: transmission within households is one of the most important ways that novel coronavirus and many other kinds of infection spread.

Contact tracing of this kind has to be done manually and has the disadvantage that it will only identify contacts or locations beyond the household that the index case is able to recall and is willing to divulge. For obvious reasons, willingness to divulge can be a huge problem, notoriously so in the context of sexually transmitted diseases. There is a short-cut available if the index case has been to an establishment keeping records (as many were required to do during the pandemic) that provide a ready-made source of potential contacts.

Speed is essential if contact tracing is to be as effective as possible. For tracing purposes, the infectious period is taken to be from forty-eight hours before symptoms appeared to seven days after. In practice, most transmissions are thought to occur in the first half of that period, which is why there is general agreement that contact tracing needs to happen fast. There is a lot to be gained if contacts are traced within a day but much less if it takes a week. This needs to be taken into account when tracing teams are overstretched; it will often be best to prioritise the most recent cases rather than trying to clear a backlog.

Mobile phone apps offer a fully automated – and fast – way of identifying contacts. In the first half of 2020 apps were promoted by some as an essential tool for tackling novel coronavirus. There was a lot of criticism of the UK government's inability to deliver one that worked. Enthusiasm had waned by the time the apps were finally rolled out in September and take-up was modest: about sixteen million people across the UK, just over one-quarter of the population. It would be bad enough if that low take-up made the app only one-quarter as effective, but the impact is far worse. A rough estimate of an app's effectiveness is the square of that fraction who use it. That's because both the case and the contact have to have the app for it to work. On that measure the apps in the UK were less than 10% as effective as they could have been. Even so, one study has estimated that self-isolations resulting from the use of apps saved several thousand lives in the second half of 2020.

Contact tracing apps bring us back to the definition of a contact. We saw earlier that a contact has to be within two metres for at least fifteen minutes, but the context matters: was the contact indoors or outdoors, was the room well-ventilated, were face coverings worn, was there physical intimacy? To an app, all contacts are equal but in terms of likelihood of transmission they are not equal at all, some kinds of 'contact' are thousands of times riskier than others.

How self-isolation works

Finding cases and contacts are important, but it's what happens next that reduces the spread of the virus. Self-isolation of cases and their contacts are at the heart of our response to a novel coronavirus epidemic. If self-isolation is working well then we have gone a long way towards winning the battle. It can have a major impact on transmission rates and is precisely targeted where the risk lies, so keeping disruption to wider society to a minimum.

We saw in Chapter 3 that the maximum R number for novel coronavirus in the UK in March 2020 was about three. That's when people who are infected behave normally: they may stay home if they feel unwell, or they may have no choice but to go out, or they may have only minor symptoms that don't alter their behaviour at all. If everyone with symptoms stayed home that would limit their contacts to members of their household, and they could be made safer by taking all precautions possible. By how much might that reduce R?

This answer depends on what fraction of transmissions occur before symptoms have appeared. That fraction is so important that it has its own name, sigma. The smaller the value of sigma the more likely any intervention targeted at symptomatic cases is to succeed. SARS had a low sigma value of about 0.10 (or 10%) and I'll use this to illustrate the argument. Sigma equals 0.10 means that if everyone with symptoms immediately and completely isolated themselves then the R number would be reduced by 90%, easily enough to bring it below one and control the epidemic. That's the key to how SARS was eradicated in 2003.

We were not so lucky in 2020. Sigma for novel coronavirus is considerably higher, close to 0.5, meaning that almost 50% of transmissions occur before symptoms appear. That implies that self-isolation triggered by the appearance of symptoms can never be enough to reduce R below one (from its maximum value of three). It gets worse. We also need to pay attention to the small print in the paragraph above: the words 'everyone', 'symptoms', 'immediately' and 'completely'. Let's look at these one by one.

Everyone: for one reason or another not everyone will choose to or be able to self-isolate.

Symptoms: not everyone will recognise that they have symptoms and some infected people may not even develop symptoms at all.

Immediately: some people may wait 'to be sure', or wait until they have a positive test result.

Completely: anyone living in a multi-person household or needing care will find it difficult to isolate themselves fully from other household member, as will anyone requiring care. The upshot is that, in practice, only a modest reduction in R can be achieved by self-isolation of symptomatic novel coronavirus cases.

Fortunately, we can improve matters by contact tracing and having contacts of cases – including members of the same household – self-isolate too. In principle, self-isolation of contacts is more effective than self-isolation of cases. That's because if it begins promptly it can reduce pre-symptomatic transmission too. It will have to be prompt though. If the contact is not traced until four or five days after they were exposed – and remember there will be a delay while the index case is reported and confirmed – then they may already be infectious and infecting others.

There is still the question of how many of the contacts are traced in the first place. By September 2020 the UK government was publishing this statistic every week and it was taken as a key indicator of how well the entire system was working. The target set was 80% of contacts reached; this was often met and sometimes exceeded. That isn't quite as good as it sounds because it includes household contacts, and you don't need an elaborate tracing system to identify household contacts. The requirement for members of the same household as a case to self-isolate had been introduced even before we went into the March 23rd lockdown.

Putting all this together in a mathematical model, a report from the Royal Society of London estimated the combined effectiveness of contact tracing and self-isolation of cases and contacts in terms of reducing the R number. They came up with a figure that was not enough to bring R below one, but would get us well over half-way there. That is why I said earlier that if contact tracing is working well – and people do self-isolate when they need to – then we can make good progress in reducing the R number.

There is one rider, however. The 80% target for contact tracing ignores a critical element of contact tracing – the fraction of cases that are recognised and reported in the first place. As we shall see later, case finding was far from perfect and this had a significant impact.

This evidence tells us that we need to pay close attention indeed to how well self-isolation is working in practice. If it isn't working well we have two options: try to improve matters, or bring in additional restrictions.

That is why I was extremely concerned when a study called CORSAIR was published in September 2020. The CORSAIR study reported that, although most people said that they fully intended to adhere to the self-isolation rules, few managed to do so: the headline figure was less than 20%. There is an important caveat that not adhering meant not adhering perfectly. Many breaches may have been minor, so this does not mean that self-isolation was only 20% effective. It will have been better – perhaps much better – than that. Even so, I feel this report should have received far more attention than it did. Instead, when cases started to rise later that same month, the talk was not of improving adherence to self-isolation, it was all about lockdown.

Behaviour matters

As we've seen, the transmission of novel coronavirus is directly related to our behaviour – the number and nature of close contacts we make with others. During the March 23rd lockdown the R number fell to around 0.7, less than a quarter of its natural maximum. As various surveys showed, a large part of this reduction was due to people changing their behaviour.

Comix surveys showed that we had reduced our average number of contacts from over ten to less than three per day. Mobility data collected anonymously from people's phones showed that we moved around much less, spending much more of our time in residential areas, not going to work or travelling further afield. High percentages of people reported wearing face coverings and adhering to social distancing. As you'd expect, we also spend more time outdoors in the warmer months; that helps because novel coronavirus does not transmit well outdoors.

Stephen Reicher – a behavioural scientist at the University of St Andrews – argued that changes in people's behaviour were motivated not only by a desire to keep themselves safe but also to protect others, that this was a community enterprise. I have to say that many of the behavioural scientists I spoke to in 2020 had a more positive view of human nature than I'd be comfortable assuming, but I hoped Stephen was right. The battle with novel coronavirus had to be a community enterprise. The most important elements of our response – case reporting and self-isolation of cases and contacts, which involve considerable self-sacrifice – are not for the benefit of the cases themselves, they are for the benefit of everyone they might otherwise come in contact with.

UK administrations did not, however, stake everything on our community spirit, they also introduced sanctions and penalties. Many of my behavioural science colleagues felt that there was too much emphasis on penalising those who broke the rules and not enough on providing help and support to those that needed it.

For me, the main problem with the ways the rules were implemented was the danger of everyone concerned losing sight of what was most important. Hill walkers and beach goers are not a threat to public health; anyone with symptoms of Covid-19 who does not get tested and self-isolate most certainly is. That message was diluted by failing to focus on the behaviours that posed the greatest risk.

The situation wasn't helped by a series of incidents where high-profile figures were caught flouting the rules they themselves were responsible for. First there was Catherine Calderwood, CMO Scotland. She was followed by Neil Ferguson, SAGE member. Then it was Dominic Cummings, senior advisor to the PM, and finally Margaret Ferrier, Scottish National Party MP. A team at University College London found that the Cummings incident coincided with a marked fall-off in the public's confidence in the UK government and affected willingness to adhere to the rules as well.

I should add that I was accused of rule-breaking myself. Somebody – I don't know who – decided to tell a tabloid journalist that I'd moved my family to the Western Isles of Scotland after the March 23rd lockdown was imposed. This was easily disproved: we'd made a long-planned trip a few days before and – like many thousands of others – got caught by the lockdown. Admittedly, this wasn't my most impressive piece of pandemic forecasting, though – as one sympathetic colleague later remarked – one of the hardest tasks we had in 2020 was predicting what the government was going to do and when.

I was directed to some very capable solicitors who stopped the false claim being printed, but the tabloid published a nasty little story anyway implying we were not welcome in a community I'd been part of for forty years. Any ill will was in short supply. Instead, we received a stream of kind and greatly appreciated messages of support, accompanied by gifts of eggs, vegetables and jams. Our sheep-farming neighbour – one of my oldest friends – loaned my teenage daughter three pet lambs to look after and our week-long stay ended up lasting four-and-a-half months, by which time the lambs had grown a lot bigger and seemed to think they had the run of the house as well as the garden.

A problem we all faced during 2020 was the sheer complexity and ever-changing nature of the rules we were supposed to live by. Rules can only have the desired impact if people know what they are, understand what they mean and can follow them. I lived in constant fear of being asked in a media interview what FACTS stood for and failing to remember. FACTS was Scotland's equivalent of 'hands, face, space' in England and stood for face coverings, avoid crowds, clean hands, two metre distance and self-isolate. Polls showed that only a tiny minority could recall the whole set. There were repeated calls for clearer public health messaging but we never got it.

How we behave is, of course, influenced by what government tells us to do, but that will not be the only factor. Our personal circumstances inevitably affect what we are able or willing to do and, perhaps most important of all, our behaviour will be driven by our perception of the risk to us and those around us. To assess those risks properly we need to understand not only how novel coronavirus spreads – which was the subject of this chapter – but also the consequences of getting infected – which is the subject of the next.

CHAPTER 8

CHILDREN AND SCHOOLS

We have arrived at one of the most disappointing episodes in the UK's response to the novel coronavirus epidemic: the closing of schools. There was never any compelling evidence that school closures would have much public health benefit, but we did it anyway. To explain how this happened, I need to start with some recent history.

Schools and flu

The 2009–10 swine flu epidemic in the UK was driven by schools and schoolchildren. The virus arrived during the summer term and the epidemic grew rapidly until the beginning of the school holidays. The epidemic then declined until schools went back in the autumn, at which point it picked up again. School-aged children had the highest infection rates. Around seventy children died, with a case fatality rate of 0.02%.

There were two reasons why swine flu disproportionately affected children.

One is that many adults had a degree of cross-immunity having previously been infected with influenza strains similar to swine flu. Children are more prone to many respiratory viruses simply because they have had no previous exposure.

The second reason is that school-aged children tend to make more contacts than any other age group, which gives them more opportunities to get infected. Typically, they have approximately 50% more contacts per day than their parents' age group and at least double those of their grandparents'.

Schools provide opportunities for children to mix, and also teachers, non-teaching staff and parents. Schools have another impact too: if the children are not at home it is easier for the parents to go to work. For these reasons, schools are seen as important hubs in the network of contacts that allows an infection to spread through a community. Secondary schools play a bigger role than primaries as they are mostly larger and have bigger catchment areas.

The epidemic models built by SPI-M over the years included all these features. In 2008, a large data set on contacts within and between different age groups was used to model a theoretical epidemic caused by a completely novel respiratory virus. The study predicted that school-aged children would be most affected, a result apparently validated by the swine flu pandemic the following year. That is why the early novel coronavirus models all indicated that schools would make a substantial contribution to the R number and school closures would be an effective way to slow the pandemic, perhaps the most effective single measure we could take.

This was the backdrop to the UK's preparations for the next pandemic. Novel coronavirus wasn't flu but it took time to overturn the preconception that schools and school-aged children would be central to the way the epidemic unfolded.

Evidence versus expectation

The first clues that novel coronavirus might be different from our flu-laden preconceptions come from the original SARS coronavirus. We knew that during the SARS epidemic in 2003 children suffered less severe symptoms and none died, although some older children were hospitalised. We also knew that children with SARS were not highly infectious, even in school settings. One review of the evidence found only one instance of transmission by a child. We knew back in January 2020 that the novel coronavirus was a close relative of the SARS coronavirus; that should have made everyone wary about any preconceptions that schools would be important.

The early data from China supported the idea that the novel coronavirus was similar to SARS with respect to children. The disease was generally much milder in children than adults and there was a lower fatality rate even among symptomatic cases (and we now know that in children most infections are mild or asymptomatic). Contact tracing studies struggled to find any evidence of children passing on infection either. In June 2020, my team did a thorough search of the published literature and we only found eight reported instances of a child passing on infection from anywhere in the world. There had been a lot of contact tracing studies published by then so this was eight out of many thousands.

Only one of those eight was in a school, and we found it in a report of a study of five primary and ten high schools in New South Wales, Australia, that was published in late April 2020. In the fifteen schools in the study there were nine cases in students that were thought to have been acquired outside school. Just one secondary case was reported in a student – in one of the high schools – out of over five hundred students in the same class as a case. No staff members were infected by a student.

The New South Wales study does not give the impression of a virus that is easily transmitted in the school environment. The study captured all the key features of novel coronavirus transmission in schools and demonstrates that we had the evidence as early as April had we chosen to act on it.

Other data supported this assessment. No large outbreaks had been reported in boarding schools anywhere in the world, suggesting limited transmission even where children were living together. There was a steady trickle of scientific papers reporting that children were at low risk from novel coronavirus and that they were less likely to transmit the virus even if infected, describing the evidence in favour of school closures as 'equivocal' and warning that interventions aimed at schools may have 'a relatively small impact'.

Over the course of 2020 more research was published on the epidemiology of novel coronavirus in children, much of it based on further contact tracing studies. Several of these papers claimed to demonstrate that children were just as infectious as adults. On closer inspection, this turned out to be only half the story. It seems that children are just as infectious as adults *if* they are showing symptoms. The catch is that children rarely show symptoms.

There were occasional reports of children who did become severely ill with novel coronavirus – fewer than a hundred in the UK during the first wave, five of whom died. Though details are not routinely released on individual cases, where we do have information the children who died were, sad to say, already very poorly with underlying health conditions. It goes without saying that these are tragic cases, but they do not alter our understanding of the epidemiology of novel coronavirus.

We would like to know why healthy children are so unaffected by novel coronavirus. So far, no-one has come up with a convincing explanation. One idea is that children are protected by recent exposure to other coronaviruses which they experience as common colds. I'm not so sure.

Factors like previous exposure to an infection are termed 'environmental' to distinguish them from genetic traits. Severe illness in healthy children is so extraordinarily rare that it's difficult to see how environmental factors can be the explanation: it would have to apply to almost every child in the world, but not to every adult. It's more likely that whatever protects children must be hard-wired into the ageing process, perhaps to do with the ageing of the immune system, particularly around puberty.

Counterarguments

Despite the absence of any evidence that children were at significant risk, there was considerable resistance to the idea that it might be safe to keep schools open – as I've explained, it challenged firmly held preconceptions. I completely accept that a lot of the resistance was driven by a genuine and deeply felt concern for the safety of the children – I am, after all, a parent myself – and there was an understandable feeling that it was far, far better to be safe than sorry. That said, we still have to interpret the evidence as objectively and dispassionately as we can, or we may make poor decisions.

One important issue was that some children who had had novel coronavirus infections went on to develop MIS-C (multisystem inflammatory syndrome in children). The syndrome is rare – there had been fewer than forty reported cases in the UK by the end of April 2020 and though it can be serious no child had died.

Naturally, when these first MIS-C cases were reported this heightened concerns over children's safety and therefore the risk in schools, even though many children become infected with novel coronavirus at home where they come in closer and more prolonged contact with adults. Over time, these concerns at least partly abated, although MIS-C cases continued to accumulate during the winter of 2020–21.

Another concern was that Public Health England was reporting dozens of novel coronavirus outbreaks in schools during the summer term even though few children were in school at the time. These data were frequently cited by politicians as a reason to be cautious about re-opening schools.

You will recall that just two linked cases can be called an outbreak, so I suspected the problem was being exaggerated and asked my team to inspect the raw data in these reports closely. We found that not all the 'outbreaks' had any confirmed cases at all, just suspected ones. We also found that a large minority involved only staff, and most of the rest involved children who were likely to have acquired infection at home, leaving only a handful of cases acquired by children (or staff) in the school environment. There were cases in schools but these were playing a minor role in the epidemic and the risk to children remained extremely low.

Looking further afield, however, there were some large outbreaks in schools. There was one in New Zealand (apparently linked to a social event involving parents), one in France (which seemed to affect the entire community but was centred on a large school), and one in Israel (which seemed to involve a super-spreader event).

For me, these were the exceptions that prove the rule: the whole world was alerted to the possibility of large novel coronavirus outbreaks in schools and yet there was just a handful of examples. Nonetheless, though the risk is low it is not zero and that justifies some public health response, as long as it is a proportionate one.

Public Health England entered the debate again during the June outbreak in Leicester. This was widely reported – not least by Matt Hancock – as 'disproportionately' affecting children, a statement that caused a great deal of concern.

Again, my team took a close look at the report on the Leicester outbreak. It was true that cases were rising among school-aged children, but they were still considerably fewer than the cases in younger adults, and those were rising at least as quickly. There was no evidence that children were disproportionately affected in Leicester. The data were consistent with the by then typical pattern that if cases start to rise in adults there is spill-over into children, much of which is likely to occur in the home.

Another worry was that – even if children were at extremely low risk themselves – they might pick up an infection and pass it on to someone more vulnerable, such as an elderly grandparent or parent with an underlying health condition. It's a valid concern but, to decide just how we should react, we need to think about where the risk really lies. To do that, we need to consider the whole household. Some vulnerable adults will be the only adult in the household, but many will live in households with one or more other adults. Those adults are also a potential risk, perhaps more so than the children. In those circumstances, the argument for keeping children out of school is also an argument for keeping the rest of the household at home too.

This brings us back to the chain of trust. Every member of the household is part of that chain, children included. This kind of situation does need careful managing – we need to do all we can to protect vulnerable individuals – but there are ways to do that other than by closing schools.

The teaching unions also got involved in the debate. Their concerns seemed to reflect a misunderstanding of the risk to teachers. There was an entirely justified worry about teachers in the vulnerable category, and the same applied to any vulnerable person in the workplace. Overall, however, there was never any evidence that teachers were at elevated risk.

An Office for National Statistics survey published in November 2020 confirmed that teachers were no more likely to have had a novel coronavirus infection than members of any other profession. Primary school-teachers had the lowest levels of infection in the survey. When you think about it, these are telling observations. It's self-evident that teachers have a lot of contact with a lot of children – far more than most other people. If children are anywhere near as infectious as adults then you'd expect exceptionally high levels of infection in teachers. This isn't what the data show, which can only tell us that children in school are not that infectious.

The indirect contribution of opening schools to the R number was also widely cited as a reason to keep them closed. It's likely that there is a modest contribution but to this day it remains hard to find much sign of it in the case data. Before the start of the 2020–21 school year there was little direct evidence on this question anyway. The evidence we had came mostly from mathematical models, and the starting point for the models was that novel coronavirus is similar to influenza. With regard to schools, it isn't.

To recap, the arguments for closing schools were that the children were at some risk, the staff were at high risk and that schools would drive community transmission, thereby raising the R number. None of these concerns were supported by the epidemiological data and surely we should expect compelling evidence before we took such a serious step as closing schools. Sadly, this remained a minority view and none of the advisory committees were willing to recommend re-opening schools, so they stayed closed for the rest of the summer term.

Correcting misconceptions

In June 2020 I was asked by Moray House, the University of Edinburgh's School of Education, if I would give a virtual seminar on novel coronavirus in schools to a group of a few dozen head-teachers. I was happy to do so and presented much the same evidence that I have recounted here.

I was completely unprepared for the reaction I got. The head-teachers looked astonished; most of what I said was completely new to them and contradicted what they had understood about the threat to the children and the staff. I'm told that the recording of the talk was widely circulated. Yet that shouldn't have been necessary; the same information should have been available through the Department for Education or other government channels.

Not long afterwards I contributed to a briefing to Scottish cabinet ministers on the risks of re-opening schools. Given the low incidence of novel coronavirus infections in Scotland at the time, together with what we now knew about the huge influence of age on the severity of infection, my team did some simple estimates of the risk of catching a fatal infection at school.

Our calculations put the risk to adults in the school at below one in four million per day, comparable with the risk of dying from an accident in the home. The risk to the children was too low to estimate properly but, at less than one a billion per day, was far less than the risk of being killed in a traffic accident on the way to and from school. In fact, it was of the same order as being struck by lightning. The ministers seemed no less astonished than the head-teachers had been. They had a different impression, but I never found out where it came from; it certainly wasn't from any epidemiological evidence.

I also got drawn into the debate on re-opening schools raging in the press. Both *The Times* and *The Mail on Sunday* made a great deal of a statement of mine that there was no evidence of any school-teacher ever being infected with novel coronavirus by a student in a classroom. The statement was accurate at the time but, as usual, the headline wasn't the whole story. I was not claiming that such a thing was impossible only that – given the intense scrutiny on schools – the lack of published examples suggested it wasn't common.

Again, it's about relative risk. The evidence is that if a school staff member gets an infection then it is at least as likely to have come from another staff member than a pupil (and, of course, staff can get infected out of school as well). There's a good argument that the most dangerous room in a school is not the classroom, it's the staffroom.

I got a trickle of hate mail every time I suggested that schools are not nearly as risky – for staff or students – as many people have been all too willing to assume. Most of my mail on this topic, however, was straightforward scepticism. Many parents, teachers and teachers' spouses (who seemed to be the most concerned of all) wrote to tell me that either I was misinterpreting the data or the data were wrong. Apparently, everyone 'knows' that children are walking virus factories so how could I be claiming that they aren't. Well, for flu and many other respiratory infections they are, but not for novel coronavirus.

Closing schools was not much help in controlling the epidemic, but it was hugely damaging and disruptive. During the summer, several reports were published setting out the harms being caused to children by denying them face-to-face schooling. The quality of online learning was patchy and students from lower-income households were likely to be at a greater disadvantage. In July, the Royal Society of London published a report – I mentioned it in Chapter 5 – citing harm to the children's mental and physical health and to their long-term prospects, as well as significant long-term damage to the economy through having a less well-educated workforce.

The Royal Society report helped to turn the tide and the scientific consensus shifted quite quickly. Russell Viner – President of the Royal College of Paediatrics and Child Health and SAGE member – was an influential and vocal advocate in favour of re-opening schools from the outset. International experience added to the evidence: Sweden didn't close schools for under fifteens at all, but still brought its epidemic under control. Denmark re-opened its primary schools in April and secondary schools in May, with few problems.

Re-opening schools and its impact

The groundswell of evidence that schools were safe and keeping them closed was harming the children won the day. Ministers changed course and schools re-opened in full in Scotland in August and in England and Wales in September.

To my mind, this was one of the biggest successes of evidence-based policy making during the pandemic and I'm proud of my team's contribution to making that happen. The irony is that this was a triumph of science over scientists. There was never any evidence that children were at significant risk from novel coronavirus or that schools were driving the epidemic in the UK. The only 'evidence' was scientific opinion based on knowledge of pandemic influenza, a different infection. Yet schools across the UK were closed (except to a small minority of children) for more than three months in the first half of 2020, affecting around ten million children and their families.

When schools did re-open there were – as expected – minimal repercussions for public health. There were cases in children but serious illness remained extremely rare. There was also a rush of demand (naïvely unforeseen by the testing system) for testing of children with possible symptoms. Positive cases in children did rise in the autumn, though they rose faster in other groups, notably university students.

In schools, the rise in cases was most apparent among older teenagers. Some epidemiologists were concerned that this age group had a higher prevalence of infection than their parents' age group. Older teenagers are considerably more susceptible and more infectious than younger children, but not more than adults. The obvious explanation for the pattern we were seeing was that these children were living lives far closer to pre-pandemic normality than any other age group, so their numbers of contacts were closer to normal too. Even so, sixteen-, seventeen- and eighteen-year-olds make up only a small fraction of the population so it was hardly likely that they were driving a resurgence of the epidemic that was affecting everyone.

This is not to say the pandemic didn't cause problems in schools in 2020. It certainly did in some, though the majority of schools had few or no cases. Where there were outbreaks these usually involved staff, sometimes bypassing the students altogether. Any disarray was due more to the need for children and staff in the same class or bubble to self-isolate, some of them on multiple occasions during the course of the autumn term. We had anticipated this in the Scottish Advisory Group, warning before term started that schools needed to plan the steps they would take if cases did occur so as to minimise disruption.

I accept that it would have been difficult not to have closed schools when the UK locked down in March 2020 – if only as a precaution – but it took far too long to realise that there was no good reason to keep them closed. Some argued that it was too risky to re-open schools while levels of infection remained high. Yet we knew at the time that the main risk from novel coronavirus is to the elderly, frail and infirm, not to schoolchildren. If we'd re-opened schools – perhaps for all but the oldest children – there would have been limited impact on the epidemic and we would have avoided massive disruption to the education of millions.

CHAPTER 15

WORLD VIEW

This was a truly global pandemic. By the end of 2020 only a handful of countries – mostly small islands – had not reported any Covid-19 cases. The global death count reached one million in late September 2020, hit two million in January 2021 and three million just three months later. The true numbers will be higher still because of under-reporting – not every death due to novel coronavirus is identified as such.

This is a huge toll, though we need to keep it in perspective. According to the Institute for Health Metrics and Evaluation at the University of Washington, novel coronavirus was responsible for 4% of deaths globally during the first fifteen months of the pandemic, putting it fourth in their cause-of-death rankings, fractionally ahead of other respiratory infections combined. The numbers change to 9% and third in the rankings using estimates corrected for under-reporting, but novel coronavirus remains well behind heart disease and stroke.

Pitfalls of comparing countries

By some estimates, the UK suffered the highest per capita mortality rate due to novel coronavirus in the world during 2020, vying for top spot with the likes of Belgium, Italy, Spain, USA, Mexico and Peru. Rigorous comparisons are tricky because – as we saw in Chapter 5 – tallying deaths due to novel coronavirus is surprisingly difficult and different countries do it in different ways.

Excess deaths show the UK in a slightly better light, though still around the top ten. Of course, the pandemic did not end in 2020 and the rankings may well change over time, not least because the UK's vaccine roll-out was faster than most. Nevertheless, this poor outcome for the UK in the first year of the pandemic inevitably led to comparisons with other countries to try to work out what we did wrong.

This kind of comparison appears deceptively simple and it was certainly popular: the media, politicians, commentators and global health experts all chipped in. I don't object to comparative studies in principle – I use them in my own work – but they need to be done as rigorously as possible and interpreted cautiously. This wasn't the case for many of the comparisons between countries carried out during the pandemic. The result was that despite general agreement that the UK had lessons to learn there was a lot of disagreement on what those lessons might be.

It was easy enough to identify countries that had done 'better' than the UK and also had greater testing capacity, had stricter self-isolation rules, had closed their borders sooner, had made apps compulsory or had gone into lockdown earlier. This was almost invariably interpreted as evidence that if only the UK had done the same then we would have had a better outcome too and was often expressed as a desire that the UK be 'more like' South Korea, Iceland or whatever country happened to support the case being made. This is known as cherry picking. It is poor science and should be resisted.

The problem with these comparisons is that countries differ from each other in a multitude of ways, so identifying the meaningful differences isn't straightforward. There are good arguments that the UK was predisposed to suffer a severe epidemic whatever we did. Here are some suggestions: infection was seeded quickly and widely thanks to our high volume of international travel; our crowded cities and our lifestyles promoted the spread of the virus; and our ageing population and generally poor population health meant that more of us were at risk of developing severe disease. The more these factors contributed to the severity of the UK's novel coronavirus epidemic the less we can blame inadequacies in our response.

I will discuss three countries and one continent that offer different perspectives on novel coronavirus: Taiwan, New Zealand, Sweden and Africa. If I am guilty of my own charge of cherry picking my excuse is that I'm aiming to show that they all illustrate why care and caution are needed for this kind of analysis. There are two kinds of lesson we might learn. There are lessons for the future, such as it helps to have prepared for the right pandemic and to act early when it arrives. Then there are lessons that could have been applied during 2020, such as lockdowns may not be necessary and can do more harm than good.

Taiwan

Taiwan managed the novel coronavirus pandemic successfully with fewer than a thousand cases and just seven deaths by the end of 2020. There were no lockdowns. The huge difference in outcome between Taiwan and the UK is part of a wider pattern. Almost every east and south-east Asian country fared better than almost every European country.

There is an obvious explanation for this dichotomy: SARS. Taiwan and other countries in the region had been concerned about the possible re-emergence of SARS ever since that disease first appeared in 2003. This was reflected in their pandemic preparedness planning. Novel coronavirus is similar to SARS coronavirus – they are closely related – so it is no surprise that these countries were able to respond effectively. The UK and other countries in Europe were much more concerned about pandemic influenza. This left them at a significant disadvantage from which they never fully recovered.

There are many differences between SARS preparedness in Taiwan and pandemic influenza preparedness in the UK. I will highlight two.

First, as we saw in Chapter 1, Taiwan acted extremely quickly, starting at the end of December 2019. We should give full credit to the public health officials in Taiwan for this timely response; I doubt anyone in the UK realised what we were facing that early on. The UK ended up well behind not only Taiwan but also Japan, South Korea and most other Asian countries.

Second, Taiwan responded vigorously, recognising the importance of rigorous case finding, contact tracing, isolation and quarantine immediately. Those are the measures you'd expect to implement if you were prepared for SARS. If – like the UK – you were prepared for influenza your focus would be more on social distancing.

There are three main reasons for taking a different approach to pandemic influenza.

First, most influenza cases are mild with non-specific symptoms, so hard to detect at all.

Second, influenza has a shorter generation time, roughly two to four days. It's four to six days for novel coronavirus. That difference matters because it makes influenza harder to stop using measures such as self-isolation of symptomatic cases and contact tracing – everything has to happen even faster, almost impossibly so.

Third, the maximum R number for influenza was expected to be considerably lower than it was for novel coronavirus, meaning that less stringent – and therefore less harmful – social distancing measures would be enough to control an influenza pandemic.

There are lessons to be learned for the future from the SARS-influenza dichotomy. One that I'd emphasise is not to plan for just one kind of pandemic threat, a point I've been making to governments and international agencies for years. Another is to act fast in the early stages of a pandemic, something I tried hard but failed to make happen in the UK in January 2020.

I am less convinced by the argument of Jeremy Hunt MP and others that the best way for the UK to tackle novel coronavirus was to be more like South Korea. Putting aside the fact that the UK is a quite different place than South Korea, that proposition wasn't made until March 2020 when it was already far too late. If we had been more like South Korea when we were doing our pandemic preparedness planning in the years leading up to 2020 then – I agree – that would have made a difference. If we had ended up being more like Taiwan that would have been better still.

New Zealand

I have mentioned New Zealand several times already as another example of a country – like Taiwan – that fared much better than the UK. Briefly, New Zealand ramped up quarantine of international arrivals during February 2020, closed its borders on March 19th, went into full lockdown on March 25th, waited until levels of infection were close to zero, and came out of lockdown. Surveillance for new cases continued and if there was evidence of community transmission strict local lockdowns were immediately re-imposed. Australia took a similar approach, though they were more frequently forced into lockdowns.

New Zealand fared much better than the UK during the first year of the pandemic. The public health burden was far smaller, just twenty-five deaths. Though there was significant economic damage – particularly to the tourism and higher education sectors – this too was less than in the UK. At the same time, it is not true to say – though some have done so – that New Zealand is a model that the UK could have followed.

The difference is in timing. We're not talking about calendar date – as it happens, New Zealand's first lockdown came two days after the UK's, not before – but timing in relation to the state of the national epidemic. New Zealand's lockdown came before there was widespread community transmission of novel coronavirus in that country. The UK's did not.

Studies of virus genomes show that thousands of infections were imported into the UK in January and February 2020 and community transmission was well-established by March. The majority of those early introductions into the UK came from Spain and France. New Zealand is far less well-connected with those countries.

To achieve the same benefit, the UK would have had to close its borders completely several weeks before New Zealand did so on March 19th. Once the epidemic was firmly established in the UK there was no way back, no route to the near-elimination state achieved by New Zealand. This is apparent from the second waves that countries around the world experienced towards the end of 2020. Every country with a first wave of similar scale – adjusted for population size – to the UK experienced a major second wave. The same is true for every country whose first wave was half, a quarter or even a tenth as bad as the UK's.

Whatever the deficiencies of the UK's pandemic response, our failure to emulate New Zealand wasn't one of them. New Zealand never had to solve the problems the UK and every other badly affected country faced, so it's a meaningless comparison.

Sweden

Now let's switch to a comparison much closer to home: Sweden. The Swedish public health authorities – led by Anders Tegnell, State Epidemiologist – said from the outset that they believed they would be living with novel coronavirus for a long time and their pandemic response had to be sustainable.

Sweden was criticised for adopting a so-called 'herd immunity' strategy. This isn't an accurate description of their policy given that they implemented a range of measures to limit the spread of the virus. As a consequence, they remained well below the herd immunity threshold throughout 2020. Whatever label is attached to their strategy, the Swedish experience in 2020 compared unfavourably with its Scandinavian neighbours. Naturally enough, this led to further criticism of their approach.

The relevant comparison, however, is between Sweden and the UK. In 2020, the two countries had similar epidemics but Sweden imposed markedly less stringent restrictions – well short of the full lockdowns in the UK – and incurred substantially less economic harm as measured by the reduction in GDP. For all but the oldest students, most schools stayed open in Sweden. The Swedish experience is profoundly uncomfortable for those who insisted that a strict lockdown was necessary to suppress the virus.

Whenever this point was raised there was a chorus of objections that Sweden was a completely different country to the UK and it was misleading to compare the two. As a debating point, this would have carried more weight if the objectors hadn't been people who – often at the same time – were campaigning vociferously that the UK should be more like South Korea. Nevertheless, let's look at the arguments that Sweden is different in more detail.

One of the alleged differences between the UK and Sweden is that their demographics differ in ways which might affect the spread of the virus. Though there is some truth in this for the UK as a whole, it is not true for the whole of the UK. So, let's compare Sweden with Scotland.

Sweden has a slightly more – I stress *more* – urban population than Scotland: 87% versus 83% live in urban areas. This matters because novel coronavirus spreads more rapidly in cities. On the other hand, Scotland has a slightly lower proportion of people in one-person households and a slightly larger average household size. This matters because the virus spreads efficiently within households.

It wasn't likely that these modest differences in opposing directions would have had much influence on the course of the two epidemics. To find out if that was correct, my team compared the trajectories of cases and deaths per capita in Scotland and Sweden throughout March and April 2020, including the all-important exponential growth phase before either country had implemented countermeasures. They were almost identical. The most important arbiter – the novel coronavirus itself – did not behave as if there was much difference between the two countries.

This leaves the question as to how Sweden managed to control novel coronavirus without lockdowns unanswered. The rate of decline of the Swedish epidemic in spring and early summer was somewhat slower than in the UK and other European countries but it was achieved without closing most schools (except to the oldest children) or most restaurants and – most important of all – no instruction to stay at home. To me, that is yet more evidence that getting the R number below one can be achieved without a full lockdown.

Detractors of the Swedish approach then resorted to the claim that Swedes were more likely to follow voluntary guidelines and take their own precautions against infection, alleging that the same approach wouldn't work in the UK. If Sweden's avoidance of lockdown was really due to vaguely defined cultural differences then we could try to identify suitable cultural levers to achieve the same result in the UK. Better that than unthinkingly assuming that nothing other than lockdown would work here.

The critics were quick to claim that Sweden's approach had 'failed' when – like every other country in Europe – Sweden experienced a second wave in late 2020. Those claims were misleading: Sweden's second wave was discernibly less severe than the UK's and it was brought under control at much the same time, again without imposing a full lockdown. The next wave – as the alpha variant took off – led to more restrictions, as it did everywhere in Europe, but still no full lockdown.

I take two lessons from Sweden's experience. First, it is possible to control novel coronavirus without imposing a full lockdown. Second, a government is more likely to take that course if they recognise from the outset that the measures they take have to be sustainable. As to whether Sweden did better than the UK, although many Swedes have been critical of their own country's response I have yet to hear of any who would swap their 2020 for ours.

Africa

Africa offers a different perspective on the balance of harms from novel coronavirus. I have worked with fellow scientists in Africa for over thirty years and in 2020 I was the director of a multi-national health research partnership called Tackling Infections to Benefit Africa (our acronym – TIBA – is a Swahili word meaning 'to cure an infection'). Our nine African partners are all leading biomedical research centres. TIBA was heavily involved in the pandemic response both in individual countries and – working closely with the World Health Organization Regional Office in Africa – internationally too.

My last overseas trip before travel restrictions came into force was a visit to TIBA's partners in Sudan in February 2020. My colleague Maowi Mukhtar introduced me to the team leading Sudan's pandemic response. He was keen to emphasise that, although budgets were small and resources were limited, his colleagues had a lot of recent experience with infectious disease outbreaks ranging from cholera to Rift Valley fever.

As we sat in one of Khartoum's endless traffic jams on the way back from that meeting, Maowi and I discussed the likely impact of novel coronavirus. At the time, public health experts and many commentators were predicting a catastrophe as the virus would quickly overwhelm weak health systems. Maowi and I demurred. We both remembered similar concerns about the swine flu pandemic in Africa ten years before that had come to nothing. We didn't think novel coronavirus would be nearly such a big problem as we were being told. We weren't the only ones; in September that year I was on the panel for a World Health Organization African Region press conference alongside a Ugandan colleague who was adamant that the threat of novel coronavirus to Africa had been exaggerated from the outset.

For the most part, our expectations look to have been correct. By the end of the 2020 there had been fewer novel coronavirus-related deaths – less than thirty thousand – reported across the entire sub-Saharan Africa region than there had been in most larger European countries. Around half of the deaths in Africa were in South Africa, the most developed country in the region, but even there the per capita mortality rate due to novel coronavirus was well below the figures for Brazil or Mexico or, for that matter, the UK.

I recall arguing with a reporter from the *New Scientist* who insisted that this must be down to poor surveillance and massive under-reporting. I didn't agree with the implication that huge numbers of unexplained deaths due to a viral pneumonia were being missed in most African cities. Under-reporting was a problem; the Institute for Health Metrics and Evaluation later estimated that an estimated 70% of novel coronavirus-related deaths in Africa went unreported, a little higher than the global estimate of 60%. Even so, correcting for this did not put the public health burden due to novel coronavirus in Africa anywhere near that of other infectious diseases such as malaria, tuberculosis or AIDS.

There were several reasons – even back in February 2020 – to expect that the novel coronavirus pandemic would be less severe in much of Africa.

One was that the environment – particularly the climate, perhaps linked to a more outdoor lifestyle – was less conducive to the spread of the virus, implying a lower R number and smaller epidemics. South Africa, of course, has a more temperate climate and is more urbanised than most other African countries, which is consistent with this explanation.

Another difference is demography. Africa's population is much younger and so a much smaller fraction is at risk of a severe outcome of infection. Simply based on age alone, we'd expect the infection fatality rate in Africa to be around a third of that in western Europe. There was a concern about multi-generational households but, on the other hand, there are few care homes in Africa.

My colleague Gordon Awandare at the University of Ghana had a different idea. He wondered if the generally higher level of infections circulating in African populations meant that people's immune systems were trained not to over-respond to novel coronavirus. That was plausible: a lot of the damage caused to patients is actually inflicted by their own body's attempts to keep the virus in check.

This proposition has echoes of the so-called 'hygiene hypothesis', the idea that if our immune systems have too little to do then they can malfunction, turning against harmless antigens – leading to allergies – or even self-antigens – leading to auto-immune diseases. The hygiene hypothesis is controversial because it implies that some level of infection is 'good' for us and most public health practitioners would strongly disagree with that.

Gordon's idea was subtler. There's plenty of evidence that different infections can interact with one another in ways that affect the outcome for the patient. These simultaneous infections are called co-infections. Sometimes co-infections make matters worse but sometimes – perhaps by regulating our immune systems – they result in less severe disease. Co-infections might well reduce the pathogenicity of novel coronavirus, in Africa and elsewhere.

Environment, demography and co-infections might work in Africa's favour but in the early stages of the pandemic evidence started to emerge in the UK that people of black African descent were more at risk of severe Covid-19. This was a direct concern to me as we are a mixed-race family – my wife, Francisca Mutapi, is black Zimbabwean (and, incidentally, the first black female professor in the University of Edinburgh's four hundred year history). Having looked at the data, we felt that not all this increased risk was directly attributable to ethnicity. It was at least partly explained by environmental factors correlated with ethnicity such as occupation and housing conditions, and by underlying health conditions that are more prevalent in black people in the UK.

Francisca was working tirelessly on the pandemic response in Zimbabwe. A key issue was the provision of care by an already hard-pressed health service. Intensive care units were few and far between and there was little prospect of increasing capacity quickly. So she took another tack. We did some work with Meghan Perry – an infectious disease consultant at the Western General Hospital in Edinburgh – demonstrating that the biggest life saver was the provision of oxygen to patients who needed it. Francisca set about persuading the Zimbabwean government to invest in a new oxygen production facility – a far more realistic goal than installing large numbers of intensive care units. Globally, this issue didn't receive the attention it deserved until the spring of 2021 when a severe shortage of oxygen in India cost many lives.

Another concern was access to PPE. Francisca's colleagues in Zimbabwe did a rapid study of novel coronavirus exposure in hospitals and found that the non-clinical staff were just as much at risk as doctors and nurses. They needed face masks urgently and Francisca procured a substantial donation from a large pharmaceutical company to fill the gap in supply.

Zimbabwe and most (but not all) other African countries took the threat from novel coronavirus very seriously. The majority imposed lockdowns at one time or another. Some were highly restrictive, some less so; some were strictly enforced, some barely enforced at all. From the outset there was concern that these lockdowns would cause harms out of all proportion to the threat posed by the virus.

Organisations such as the World Bank, UNICEF and Save the Children published a steady stream of reports forecasting huge increases in poverty, malnutrition and starvation, together with health crises caused by interruptions to the delivery of childhood immunisations and malaria control programmes. Educational harms would be huge and closing schools would have wider consequences as well, given that many health and welfare programmes in Africa are centred on schools. They anticipated that altogether these would result in far more deaths than would be caused by novel coronavirus.

We can and we will debate whether lockdown caused more harm than novel coronavirus in the UK and other western European countries. I don't think there'll be much debate about that issue in Africa though. For every country in sub-Saharan Africa – with the possible exception of South Africa – there can be little question that the cure will turn out to be far worse than disease. That's a lesson for everyone to learn, not least the international agencies directing the global pandemic response.

World Health Organization

The World Health Organization has been a positive force in global health since it was founded in 1948. During the novel coronavirus pandemic, the World Health Organization was – as you'd expect – widely seen as the most authoritative source of information on novel coronavirus and of guidance on measures that countries should take to control its spread. For the most part, it fulfilled that role well, accepting that its advice changed as we learned more about the new virus – that was inevitable.

Sadly, though, the World Health Organization got the biggest calls completely wrong in 2020. I'll talk about three issues in particular: the slow acknowledgement of the severity of the threat; the failure to recommend an international travel ban when it could still have made a difference; and the promotion of lockdowns long after it became clear that in many places they were causing more harm than the novel coronavirus itself.

Let's start with the delays to acknowledging the threat. The World Health Organization has two main levers at its disposal.

The first is to declare a public health emergency of international concern, which it only did on January 30th, though novel coronavirus had been detected outside China seventeen days earlier.

The second lever is to declare a pandemic, which it did on March 11th, weeks after it was apparent to most epidemiologists that a pandemic was well under way. It matters that these declarations were made far too late because they are necessary precursors for what the World Health Organization calls 'co-ordinated international action'.

One co-ordinated action that could have made a difference right at the beginning of the pandemic was a comprehensive ban on international travel. This is not hindsight. Ever since the 2003 SARS epidemic, epidemiologists have been well aware of how fast a new virus can spread through international travel, especially air travel. We knew that to have any chance of averting a pandemic action must be taken very quickly indeed. Ideally (from an epidemiological perspective), China would have closed its borders completely in early January 2020.

Next best would be for the World Health Organization to have acted on January 13th when the first novel coronavirus case was confirmed outside China, in Thailand. That was never likely though. As late as February 29th they declared that *the World Health Organization continues to advise against the application of travel or trade restrictions to countries experiencing Covid-19 outbreaks.* The only redeeming feature of this declaration is that by the end of February it was already too late for a travel ban to make a decisive difference in many countries.

It wasn't only the World Organisation who was fault. I should point out that few, if any, UK scientists were calling for the border to be closed at the outset of the pandemic. I didn't mention border closures in my January 2020 e-mails to the CMO Scotland. Why not? My less than satisfactory explanation is that I didn't think it could or would be done – it was too big a step, too quickly, and the suggestion would have been seen as massively disproportionate. I was almost certainly right about that but it doesn't get me – or anyone else – off the hook. It still would have been correct advice and I didn't give it. Looking back, I made a judgement about the credibility of the advice I was providing that wasn't mine to make.

We will need to work out how to persuade the world to suspend international travel immediately should we find ourselves in a similar position in the future. Such drastic action wouldn't be warranted every time we find a new human virus. As we shall see later, that happens almost every year, and most of them do not cause pandemics. Yet it would have helped enormously had it been done early in 2020. Even if a global travel ban wasn't implemented quickly enough to prevent a pandemic entirely, it would have slowed it down, stopped the virus being seeded so widely, and bought the world some time.

I explained the World Health Organization's rationale for recommending lockdowns in Chapter 4. To recap, lockdown was conceived as a time-limited, localised intervention with the aim of eradicating novel coronavirus entirely. It was never intended to be a sustainable intervention implemented on a global scale. It should have been abandoned once it became apparent that eradication was not feasible (with the exception of countries that were in a position to try for elimination – and to completely close their borders).

Yet the World Health Organization continued to promote lockdown even as cases and deaths soared all over the world. The lockdown strategy clearly wasn't working and was causing immense harm at the same time – the worst possible outcome. The key was to recognise much earlier on that novel coronavirus was not going to be eradicated and to promote sustainable responses to the pandemic. The World Health Organization did finally come round to this view, but not until October 2020. Even then, it never spelled out what an alternative to lockdown might look like for countries unwilling or unable to inflict such a damaging policy on their people.

It remains to be seen where all this leaves the World Health Organization in the long run. No doubt there will be calls for it to be given more powers and to be more agile, and also calls for it to be reformed or even replaced. A report by one high-level body – the Independent Panel for Pandemic Preparedness and Response – published in May 2021 mentioned all those options without making a recommendation, but did acknowledge that the early global response to the pandemic was woefully inadequate.

Whatever the World Health Organization's fate, we need firmer international commitments to underpin a co-ordinated international response to the next pandemic, as we do to other global health challenges such as antimicrobial resistance. Experience of climate change – another problem that cannot be solved unless countries act together – shows this will not be easy. It has to happen though.

CHAPTER 16

SAGE SCIENCE

The fact that the UK had one of the worst novel coronavirus epidemics in the world sits uncomfortably with the widely held view that the UK is good at science. Using any measure of scientific achievement – numbers of research papers published, numbers of times those papers are cited by other scientists, number of Nobel Laureates – the UK comes out as punching well above its weight. That strong science base should have given us an advantage during the pandemic, yet we fared so badly.

Role of advisors

The UK also had the advantage of a well-developed system for feeding scientific advice into government policy. This process is the responsibility of scientific advisors appointed by all key government departments and by the devolved administrations. The most senior of the network of advisors are the Government Chief Scientific Advisor (CSA) and the Chief Medical Officer (CMO) for England.

In 2020 the CSA was Patrick Vallance and the CMO was Chris Whitty. The CSA chairs the Scientific Advisory Group for Emergencies, usually known as SAGE. SAGE is convened to deal with crises such as flooding, nuclear accidents or pandemics. It first met to discuss novel coronavirus on January 22nd 2020 and continued to meet at least weekly throughout the year.

SAGE handpicks many of its members according to the nature of the crisis before it. The 2020 membership was extremely well qualified and brought a wealth of expertise and experience of infectious diseases to the table. The same was true of the SAGE sub-committees: SPI-M, which advises on epidemiological modelling; SPI-B, which advises on behavioural science; and NERVTAG, which advises on viral threats.

The majority of scientific advisors are unpaid. The CMOs and CSAs are on the government payroll but the rank-and-file are seconded, pro bono, from their regular jobs, be these in universities or government agencies. Even so, as you'd expect, advisory committee work became the main activity for many of us in 2020–21, often eating up every working hour and more.

I know many of the scientists on these committees well and have the utmost confidence in them. I know the quality of the work they do. The advisory system gave the UK government direct access to some of the best expertise in the world. Yet, somehow, the UK still ended up with one of the world's worst novel coronavirus epidemics. So what happened? Did the scientists give good advice, which the government ignored? Or did the scientists give bad advice, which the government acted upon? There are no clear-cut answers to those questions but it seems to me that there were faults on both sides.

One of maxims within government is: 'advisors advise, ministers decide'. In some ways it's fair comment; most of the time scientists play a limited role in policy development. One SAGE member expressed those limitations in saying: 'the role of the scientific advisor is to help the government do things ever so slightly less wrong'. The problem with this is that it lets the advisors off the hook by minimising their responsibility for the actions that are taken on the basis of their advice.

On the other hand, during the pandemic we often heard ministers say that they were following the science. Ministers have used that line for as long as I can remember, and I've always been wary of it. It lets ministers off the hook by hiding behind the scientists. The truth is somewhere in between. Greg Clark MP – Chair of the House of Commons Science & Technology Select Committee – captured this when he said to Chris Whitty that ministers are not being given much leeway when the government's own Chief Medical Officer tells them that if they didn't impose a lockdown immediately the NHS would be overwhelmed. It was a fair point.

I have advised the UK government on infectious diseases for more than twenty years, covering a variety of difficult topics including pandemic influenza, foot-and-mouth disease and bovine tuberculosis. My aim is always to explain as best I can the pros and cons of the options on the table, not to advocate for one or the other. As I said on Radio 4's *Today* programme one morning in June 2020, I felt that I had done my job as an advisor if I was heard, and if I was understood. That's all.

Not everyone saw it that way. It became routine during the pandemic for scientific advisors and commentators to push for one policy or another. As citizens they have every right to do that, but when there is a constant stream of government scientific advisors in the media calling for a lockdown – as there was during the second wave in October 2020 – then we have crossed a line. This is not advice, this is advocacy.

Such behaviour would be an issue even if we were dealing with a purely public health problem, but we were not. The ramifications of lockdown extend far beyond any short-term public health benefit. Yet, as Patrick Vallance has openly admitted, SAGE was never tasked with the bigger picture, its primary role was to advise on the controlling the epidemic. It was exceptionally well qualified to do that but – given that so many of its members were clinicians and public health experts – its advice was always going to err on the side of averting an immediate health threat with too little emphasis on long-term harms.

Several explanations have been put forward as to why the UK's science advisory system wasn't more successful in 2020: group-think, unconscious bias, tunnel vision, hubris, discouragement of dissent and lack of diversity, to name a few. I accept each of these could have played a part, though I'd add that most of these charges could be levelled at almost every committee I have ever known. There's no question that everyone involved was doing their best.

The question that bothered me more was whether the government was using its own advisory system to best effect, whether the advisors were being asked the right questions. Government seemed to spend much more time thinking about tactics – such as how many people could meet or what times pubs had to close – than about strategy. Even senior ministers opined on minutiae such as the maximum length of a queue for takeaways, a question surely best left to the public health agencies. Strategic questions were left unasked and unanswered.

For example, SPI-M was never asked to model alternatives to lockdown. Most of the work on identifying alternatives was done outside not inside the science advisory system. That's hard to justify – it's not as if lockdown was uncontroversial, so why not use the huge scientific resource available to government to try to find another way?

A sub-optimal relationship between science, scientific advice and policy was an unexpected weakness given that the UK had invested in both pandemic preparedness planning and a sophisticated system for communicating scientific advice to ministers. This wasn't the only line of communication, however. The novel coronavirus epidemic was the biggest news story of the year by far and there were plenty of voices eager to participate in the discussion.

Any number of scientists – not least many of us on the advisory committees – were vocal in the media (usually 'speaking in a personal capacity' – a phrase that some found irritating but signified only that they were not hearing an official pronouncement by any government department or agency). Throughout the year there was an almost daily public airing of topics such as the necessity of lockdown, the pros and cons of keeping schools open, what relaxations could be permitted at Christmas and who should be the first to be vaccinated.

I think it was right that this happened. Exactly the same debates were happening in research seminars and advisory committee meetings, and the issues being debated there had too big an impact on too many lives to not be discussed freely and openly. The public discourse kept pace with the scientific discourse pretty well.

Sometimes the scientists spoke in more co-ordinated fashion. These included well-respected bodies such as the Royal Society of London and the Academy of Medical Sciences. There were also pop-up groupings such as Independent SAGE – led by ex-CSA David King, who I'd worked with on several occasions since the foot-and-mouth disease epidemic of 2001 – and lobby groups such as the Tony Blair Institute for Global Change. All had useful contributions to make, though they didn't always agree with government policy.

I suspect that the net result of having so many competing voices was that it became easier for politicians and policymakers to ignore all of them. That said, I am strongly of the view that every voice has the right to be heard, taking the view attributed to Voltaire by his biographer Evelyn Hall: 'I may not agree with what you say but I'll defend to the death your right to say it'. I must add the standard rider that with freedom comes responsibility. That's true of scientists as much as anyone.

Doing research at pace

There was certainly a lot of science to explain, comment and advise on. The pace of scientific research on novel coronavirus during 2020 was extraordinary and unprecedented. Thousands of research groups around the world – including my own – were re-purposed to work on the pandemic threat. Tens of thousands of scientific papers were written on every conceivable aspect of the biology of the virus and all sorts of ideas for tackling the pandemic.

The volume of research output was far too great for any one individual to keep track. The Usher Institute – where I'm based at the University of Edinburgh – quickly set up teams to carry out rapid reviews and syntheses of the current evidence on pretty much any topic relating to novel coronavirus. This was an unheralded contribution but a valuable one and I hope it will become standard practice for major scientific issues.

In normal circumstances, the process of translating scientific knowledge into policy is painfully slow. The link between smoking and lung cancer was first identified in 1950 but it took decades before serious efforts were made to reduce smoking.

There is a well-defined sequence of events: the research is done; it undergoes peer review; it is published in a scientific journal; more research is done that replicates or refutes the findings; a consensus is reached among scientists; this is conveyed to policy-makers and fed into the policy-making process. At every stage the science is subject to scrutiny and challenge, and new evidence may emerge that shifts the consensus. Entire books have been devoted to unpicking how this process works.

During the pandemic the time between doing the research and informing policy – normally years or decades – was reduced to weeks or even days. Scientific advice had to be based on whatever relevant information was available at the time and sometimes policy decisions had to be made with a minimum of evidence, but there was no way around this. I recall explaining to a Scottish Government minister that it was entirely possible the advice that my colleagues and I were giving could change as new evidence emerged. The minister completely accepted this but pointed out that it was difficult for politicians if they then had to change policy as a result – they would be accused of U-turns and flip-flopping. Following the science wasn't going to be as easy as it sounded.

It is right to ask whether the pace of work meant that there was too little quality control on the research being done and that resulted in poor scientific advice to government. I don't think it did. I am not aware of any key decision being made on the basis of scientific evidence that hadn't been scrutinised by one committee or another. In any case, the question the policy-makers should be asking is not what does any individual scientist or scientific paper say but what is the consensus view. A strength of the UK's advisory system is that it aims to look at all the available evidence and reach consensus, while acknowledging disagreement where disagreement exists.

In my view, the bigger problem was that shifting consensus can be a slow process. I felt this most acutely when it came to schools. My reading of the evidence from as early as March 2020 was that it was safe to keep schools open, but it took much longer for a majority of my colleagues to be convinced. I have mixed feelings on this. Of course it was a disaster that so much schooling for millions of children was lost. However, it would have been wrong for policy to change on the basis of a minority scientific opinion. When the evidence became overwhelming, the consensus changed.

I think the lesson here is more about prejudgement. The assumption that children were virus factories and schools would drive the pandemic – taken straight from the influenza chapter of the epidemiology text books – was so ingrained that it took a greater weight of evidence than it should have to overturn it.

As I've already mentioned, a crucial step in the scientific process is the publication of research findings following peer review. Like most working scientists, I have often taken issue with peer review – especially when a cherished paper is rejected by a leading journal – but it can be a positive and constructive process and, overall, my own publications have been the better for it. In any case, despite numerous attempts, no-one has yet come up with a better system for the quality control of research outputs.

Peer review isn't perfect though. Not only do good papers get rejected but, just as worrying, bad papers sometimes get published, such as the retracted paper on the safety of the MMR vaccine in *The Lancet* that I mentioned in Chapter 3. For that reason, I have always argued that peer-reviewed publication in a scientific journal is not enough by itself. Further corroboration is needed before scientific evidence is considered robust enough to inform policy.

In 2020, peer review was largely bypassed anyway. The process of publication following peer review was simply too slow. In normal circumstances it takes several months and – even if it were accelerated for particularly important papers – there would still be a delay of weeks before publication.

The solution was the publication of pre-prints, advance copies released prior to peer review. This quickly became the norm and novel coronavirus research was sometimes described as science by pre-print. Many scientists complained that if the paper hasn't been through the peer review process we had no way of judging its quality. Actually, we did: we could make that judgement for ourselves, something we should be doing anyway if drawing on the work to inform policy.

In my view, science by pre-print worked surprisingly well in 2020. I'd go so far as to say that in many respects it worked better than the conventional route of publication in scientific journals. I have already mentioned the pros and cons of the peer review process but the biggest problem with scientific journals is the editors.

Journal editors have immense power over what research gets published yet they are not accountable to the scientists whose work they judge, nor to the users of that research, such as government. The vast majority of journal editors are well meaning and good at their jobs, but they are not all knowing and they are no more free from bias – conscious or unconscious – than the rest of us.

Some journals dealt with this better than others, openly acknowledging the problem. I was particularly struck by a blog written to accompany a paper on modelling herd immunity published by one of the world's leading journals called, simply, *Science*. In the blog, the journal's editors said that they were concerned that the paper gave a different result to other modelling studies and might be interpreted as downplaying the severity of the pandemic. Nonetheless, the paper had passed peer review and they felt that it was in the public interest to publish.

Science was more enlightened than most; some other leading journals seemed to favour certain perspectives on the pandemic over others. Science by pre-print lets us see all the novel coronavirus research there is, not a subset of it filtered by editorial whim.

Most people, of course, don't read scientific journals – their understanding of the science is filtered through the media. I had countless conversations with journalists in 2020 and – putting aside the inevitable desire for a punchy headline – my experience was that the great majority wanted to get the facts right, even if they had their own slant on the story.

I recall talking to a *Mail on Sunday* journalist one Saturday morning about a specific entry in a table in a scientific paper, purely to confirm whether or not the data were consistent with an article the paper was preparing on novel coronavirus in schools. That's impressively diligent journalism.

In this book so far, I have been critical of BBC television's coverage of the pandemic – their main news programmes repeatedly misrepresented the risks from novel coronavirus – but it wasn't all bad. If you dig into the health and science sections of the BBC News website James Gallagher, Nick Triggle and colleagues were providing informed and more balanced coverage.

The media's efforts to explain the science of novel coronavirus and its implications were greatly helped by the work of the Science Media Centre. The Centre was set up in 2002 with the remit of improving communications between scientists and journalists. This was the era that Ben Goldacre described in his 2008 book *Bad Science*, characterised by alarmist reporting of threats to our well-being and over-hyped claims for possible cures and remedies. The Centre's aim was to deliver more accurate and better balanced science reporting and it came into its own during the novel coronavirus pandemic.

The Science Media Centre's way of working is to highlight the latest scientific reports and publications and post comments from those on its extensive list of scientific experts. These posts are often widely picked up by journalists and news agencies in the UK and beyond. The Centre often links individual scientists with individual reporters, and regularly holds briefings for the media on hot topics – I spoke at several of these myself.

The Science Media Centre provided journalists with an accessible source of authoritative scientific comment on the epidemic and its ramifications which was used to full effect, and we all got better quality media coverage as a result.

Science that made a difference

The UK was a major contributor to the global output of research on novel coronavirus – as I said, we have a strong science base. The UK spends tens of billions of pounds on scientific research every year, through government, industry and charities such as the Wellcome Trust. This huge amount of money is distributed to many different pots. I have wondered for years whether the right amount of money was going into the right pots. The novel coronavirus pandemic convinced me that it was not.

The Rolls-Royce disciplines of biomedical science are molecular biology, cell biology, immunology, physiology and the 'omics' (a collection of high-tech approaches including genomics, transcriptomics, proteomics and metabolomics). These are the disciplines that have the highest status within the scientific hierarchy, get the most investment and the newest buildings, publish in the top journals and receive the most prestigious prizes.

Much of the work done in these Rolls-Royce disciplines is not immediately practical and is described as 'fundamental' science. As such, it is aggressively defended whenever there is any suggestion that science budgets should reflect utility as well as the more nebulous concept of 'excellence'. A lot of fundamental biomedical science is done using animal models and I once had a head of department who regularly called by my office waving the latest cover of a leading science journal exclaiming, 'Look Mark, another great advance for mousekind!'

Towards the end of 2020 I was reflecting on the contribution of the Rolls-Royce disciplines to the pandemic response so far. The answer I came up with was: practically nothing. It is true that a huge amount of research was being done in these subjects, and it is also true that fundamental science of this kind provided the underpinning for the vaccines developed in 2020. Also, as we've seen, genomics did start to make an important contribution in December. So I am not arguing they are worthless, far from it. I am simply pointing out – politely but firmly – that science's most prestigious disciplines had minimal impact on the ground during the first year of the pandemic.

The disciplines that saved lives during 2020 were ones much further down the scientific pecking order: clinical medicine (mostly patient care), behavioural science, epidemiology and public health. These are less Rolls-Royces and more battered old Land Rovers, lacking both funding and prestige. The interventions that really made a difference – improved patient survival, contract tracing and outbreak investigation, identification of risk factors and advice on hygiene, face coverings and ventilation – were all delivered by these disciplines, many of them at a remarkably early stage. That's how science saved lives in 2020.

I can illustrate this by referencing the research work done on the most important risk factor for severe illness due to novel coronavirus, age. Clinicians and epidemiologists recognised age as a risk factor before the UK epidemic had even begun. This was vital information and allowed patient management protocols and public health interventions to be tailored accordingly. The knowledge that age is a risk factor undoubtedly saved lives. It also showed that it was safe to keep schools open. Within six months, epidemiologists had measured other risk factors so precisely that we could identify 5% of the population who were over fifty times more likely to die from novel coronavirus infection than the other 95%.

It would have helped to know precisely why age and other risk factors were important because then we might be able to identify those most at risk more reliably. There were several contenders: senescence of the immune system, expression of specific molecules on cell surfaces, or more recent exposure to other infections. Yet, despite decades of enormous investment in fundamental research on ageing, cell physiology and immunology, we still don't know the answer. The Land Rovers delivered; the Rolls-Royces never got out of the garage.

I can make a similar case for intervention technologies, meaning vaccines, drugs and diagnostics. New vaccines and therapies were always going to be our main hope for finding a way to live with novel coronavirus. The efforts during 2020 that went into delivering them as quickly as possible were truly remarkable. Vaccine and drug development, and all the fundamental science that underpins those endeavours, have been worth every penny invested in them.

All of that is true but it still misses one crucial point: for the first year of the pandemic, vaccines (outside clinical trials) didn't save a single life in the UK – we didn't have any (though Russia and China were a little quicker). Drugs performed better. Dexamethasone, for example, significantly improved survival of severe cases in hospitals. Its effectiveness in treating Covid-19 cases was demonstrated through a fast-track clinical trial called Recovery – another terrific example of good, practical, high-quality science done at extraordinary pace. Even so, Dexamethasone is not a new drug. Completely novel therapies, such as plasma containing antibodies against novel coronavirus, were less successful.

Contrast this with diagnostics. The RT-PCR test was available within weeks and used around the world. Without it, we would have been working in the dark, case finding and outbreak investigation would have been severely compromised. A few months later, cheap, rapid diagnostic tests had been developed and mass testing on a huge scale became feasible, opening up new possibilities for reducing individual risk and suppressing the virus.

As an epidemiologist, I have always seen diagnostics as a crucial public health tool yet diagnostic development receives nothing like the prioritisation and investment given to vaccines and drugs. Back in 2006, I co-authored a major UK government report – called a Foresight study – on the importance of diagnostics for infectious diseases. The report had a modest impact on investment in the field, but it didn't stop one of the most senior medical scientists in the country telling me that diagnostics weren't cures and so weren't that important. Novel coronavirus proved such attitudes wrong. Of course a RT-PCR test isn't a cure, but when it is used to trigger self-isolation and contact tracing it saves lives.

When we are faced with this kind of crisis again – and I'll say more about that later – I hope that we will have learnt the lessons that 2020 so painfully taught us. A lot can happen while we are waiting for a new drug or a new vaccine to solve the problem, many people can die and a horrendous amount of damage can be done. During that period we will be relying once more on medics learning how to treat patients better, epidemiologists learning how the disease spreads, public health practitioners working out what interventions are most effective and developers of diagnostics giving them all the tools they need to do their work. I hope that in whatever time we have before the next pandemic we will have nurtured those disciplines so that they are there, ready to help when needed.

CHAPTER 17

WHAT SHOULD HAVE HAPPENED

The July 2020 article I wrote for the *New Statesman* concluded with the words *I fear that history will judge lockdown as a monumental mistake on a truly global scale.* So I suppose I shouldn't complain that I was introduced on the BBC's *Andrew Marr Show* two months later as a lockdown sceptic. I don't like the term though; it has a ring of 'climate change denier' or 'flat earther' about it.

The case against lockdown

My lockdown scepticism is rooted in the long list of harms caused by the lockdowns implemented in the UK, the unintended but predictable consequences of trying to control novel coronavirus by shutting down society. My hope – and my main motivation for writing this book – is that lockdown scepticism will become the mainstream view.

I must repeat that I am not playing down the threat of novel coronavirus. For many people Covid-19 is a dangerous disease, though the majority of the population is at relatively low risk. The pandemic always warranted a robust public health response, nationally and internationally. For me, that's not the issue; the issue is whether that response had to be lockdown. I have three reasons to claim that lockdown was an over-reaction to novel coronavirus.

First, we didn't take adequate account of the damage that lockdown would do to lives and livelihoods. We will be paying the price for a generation or more.

Second, the public health benefits of lockdown were over-estimated, skewing the argument in favour of that course of action, however damaging it was.

Third, there were much less damaging ways to save lives and protect the NHS. Those alternatives are the subject of this chapter.

It goes without saying the best way to avoid lockdown is not to have an epidemic in the first place. I have already talked about two courses of action that might have made that ideal outcome possible. We could have implemented border controls and international travel bans right across the world in January 2020. That didn't happen. Then the UK could have taken those same actions unilaterally in February. That didn't happen either. I don't think either option was realistic at the time, too many decision-makers had yet to accept the seriousness of the situation. Taiwan and several other countries were much quicker off the mark. That's a crucial lesson to learn: we will need to make better decisions faster should we find ourselves in similar circumstances in the future.

Having said those actions should have been taken, I must add a crucial rider: it's far from certain how well they would have worked. Remember that border controls delay but do not prevent the importation of cases, so we'd face a continual threat. Even if the UK had successfully contained every novel coronavirus outbreak we experienced, I don't for a moment imagine that epidemics could have been prevented in every single country in Europe, the Americas, the Middle East, South

Asia and Africa. Perhaps more countries would have been able to sit out the pandemic behind closed borders in the manner of Australia and New Zealand – with a lower death toll as a result – but the virus would still be there, life would be far from normal, and we'd still have to find long-term, sustainable solutions. We'd all have to learn to live with novel coronavirus eventually.

For now though, let's focus on the situation that the UK faced from early March 2020 onwards. Given what we knew at the time, could we have done better? Finding the best path through the pandemic was never going to be easy, but the path we took led to the UK suffering both one of the highest mortality rates in the world and a hugely damaging series of lockdowns. Is it possible that we could have saved lives and kept out of lockdown as well? My answer to that question is an emphatic yes.

Alternative views

There are different perspectives on the way the UK should have responded to the pandemic. We have already encountered the view that no death due to novel coronavirus is acceptable. Taken to its limits this argument says we should lock down immediately, as tightly as possible, and stay there until we have rolled out an effective vaccination programme. No country in the world was prepared to take such an extreme course of action. There has to be a balance between the public health harms due to novel coronavirus and the economic and other harms caused by lockdown.

A high-profile proposal for an alternative to lockdown was published in October 2020 in the form of the Great Barrington Declaration. The declaration was named for the home of a think-tank called the American Institute for Economic Research and was led by a group of US and UK scientists.

The Great Barrington Declaration advocated an approach where the vulnerable are protected but the virus is left to circulate in the remainder of the population until enough people have been infected to reach the herd immunity threshold. According to its website, around fifty thousand academics and clinicians signed up to the declaration. I was not one of them. I thought there were serious problems with the Great Barrington approach.

One problem was that the epidemic envisioned in the Great Barrington Declaration would be far larger than the one the UK experienced in 2020. Even if cases were entirely confined to the less vulnerable majority, tens of millions of people would be affected, many times more than were actually infected that year. By definition, this group is less much likely to suffer severe disease and die, but a small fraction of a very large number is still a large number. This was not a viable strategy: we would still end up with many more people needing to be hospitalised than the NHS could cope with.

The second problem was the absence of a convincing plan for confining the epidemic to the less vulnerable majority in the first place. This would have to happen for however long it took for the herd immunity threshold to be reached, most likely several months. I have argued throughout this book that we could and should have protected those who most needed protecting far better than we did, but there are limits to how well that could be done. If the virus were circulating freely it was inevitable that some infection would spill over into the vulnerable minority and cause a substantial burden of death and disease in that group too.

There was a third problem with the Great Barrington Declaration. Back in October 2020 we did not know whether acquired immunity to novel coronavirus would be effective enough to get us over the herd immunity threshold. It was risky to assume in advance that immunity would be as good as it is for infections such as measles – that could turn out to be over-optimistic. If we failed to reach the herd immunity threshold in a matter of months but rather – as some epidemiological models predicted – suffered repeated waves of infection then the strategy wouldn't work at all.

For a time, the Great Barrington Declaration received a lot of attention from the media and politicians in both the UK and the US. This provoked an unrestrained backlash from medics and scientists who lined up to denounce the declaration as dangerous, unethical and poor science. The criticism was so vehement that several journalists asked me why the scientific community was so unwilling to listen to new ideas. When an alternative approach is rejected that forcibly it creates a climate of 'it's our way or nothing' and that's not a good impression to give. Dogma has no place in science. If discussion of alternatives were to be stifled then we were condemned to lockdown.

Back to basics

From my point of view, the most unfortunate consequence of the Great Barrington Declaration debate was that it hardened the view that 'protecting the vulnerable' was somehow inextricably linked to a herd immunity strategy of one kind or another. It wasn't. It is true – as Jeremy Farrar wrote in *The Times* in September 2020 – that suppressing the virus is incompatible with the build-up of herd immunity. In other words, you can have lockdown or the Great Barrington approach, but you can't have both. That's not the whole picture though. There is a middle ground and I believe that is where the answer lies.

Let's start with the evidence. Even before the UK first went into lockdown in March 2020 we knew that novel coronavirus is far more dangerous to the elderly, the frail and the infirm than it is to healthy young adults and children. I do agree with the

Great Barrington Declaration that this simple, straightforward epidemiological fact should have shaped our response. Instead, the UK administrations continued to act as though everyone was equally at risk. They even actively promoted that misperception to try to bolster acceptance of the lockdown strategy.

Having established who is at risk the next step is to devise ways of protecting them. As we saw in Chapter 4, suppressing the spread of the virus wasn't enough by itself; huge numbers (possibly a majority) of people who died of Covid-19 got infected during periods of lockdown. Lockdown didn't save these people; we needed to do more. During the first wave in the UK all we actually did was tell the vulnerable to shield themselves through self-isolation – which was neither desirable nor sustainable – and, belatedly, increase biosecurity around care homes. In or out of lockdown, these interventions weren't enough. Some have claimed that this proves that protecting the vulnerable doesn't work. That is like saying that fire-fighting doesn't work because we failed to put out a raging inferno by throwing a bucket of water at it. The truth is that we didn't really try to protect the vulnerable.

I will turn to economics to demonstrate how little we tried. It will be a while before we know the full economic costs of the pandemic but the UK reportedly had to borrow an extra three hundred billion pounds in the 2020–21 financial year. By definition, we incurred that debt because of the pandemic so it wasn't an either/or choice between controlling novel coronavirus and sparing the economy; other things being equal, a smaller epidemic is better for the economy. On the other hand, we should not have been willing to pay any price at all to control the virus. It was always possible for the cure to be worse than the disease, so we needed to find the right balance.

Working out how much higher or lower the costs would have been had we taken a different approach is a counterfactual problem. We came across counterfactuals in Chapter 4 when trying to estimate how many people would have died if we hadn't gone into lockdown. That was a job for epidemiologists. Now we have to work out what the economic costs would have been in this alternative scenario. The answer will again depend on how we imagine people – and businesses, banks and governments – would have behaved in a different set of circumstances. That will always be disputable and no doubt economists will be debating the question for years.

I will look at the imbalance in our response in a much more straightforward way. The UK spent an enormous amount of money on its pandemic response. Most of that money was linked to the measures taken – including lockdowns – to suppress the virus. By comparison, how much did we spend on protecting the vulnerable, both in care homes and the wider community? The answer is not much at all.

There are about twenty thousand care homes in the UK and well over thirty thousand people in those care homes died with novel coronavirus in 2020. (For context, it's worth noting that over one hundred and fifty thousand care home residents die in a normal year). According to the DHSC, the UK spent about five billion pounds on care home support. That's two hundred and fifty thousand pounds per care home. Several times as much could have been spent on protecting care homes – enough to greatly improve biosecurity measures – and that would still amount to a tiny fraction of three hundred billion pounds.

That's one part of the challenge, but the majority of vulnerable people don't live in care homes. Across the UK, over two million vulnerable people in the community were advised to shield in 2020 (increased to four million in 2021). There are about seven million more households that have people over seventy years old or who are otherwise at elevated risk. More than forty thousand people in these groups died of novel coronavirus infection in

2020. I haven't been able to find estimates of government spend on shielding. It can't have been much. Those asked to shield were given advice and offered help with their grocery shopping. Other vulnerable people got nothing more than a letter advising them to take extra precautions.

Each of those households would have benefitted from personalised advice on how to implement Covid-safe measures and being provided with the means to do so, together with routine testing of contacts and support for shielders if they needed to self-isolate, and for those they were shielding if that happened. We could have spent several thousand pounds per household – a substantial sum by any standards – on these measures and that would still amount to a small fraction of three hundred billion.

I am particularly struck by the neglect of shielders, especially care-in-the-home workers and informal carers. These people had the responsibility of standing between the people they are caring for and novel coronavirus. For most of 2020, they got minimal recognition and received no help.

All in all, the under-investment in protecting the vulnerable is one of the most baffling failures of the UK's pandemic response. To say we spent a fortune on suppressing novel coronavirus doesn't even begin to cover it. The UK took on more debt than at any time since the Second World War and economists anticipate we will be servicing that debt for generations to come. We spent, in comparison, a trivial amount on care homes and almost nothing on protecting the vulnerable in the community. We should and could have invested in both suppression and protection. We effectively chose just one.

U-turn on protecting the vulnerable

Ideas such as protecting the vulnerable, cocooning (protecting the vulnerable by protecting their shielders) and segmenting (targeting interventions at the appropriate subsets of the population) ought not to have been controversial. They were sensible and proportionate approaches to tackling novel coronavirus. Nonetheless, they were fiercely resisted by scientific advisors who had mistakenly convinced themselves that they were not compatible with suppressing the virus. The irony is that they were all adopted in one form or another at various phases of the epidemic, just not quickly or widely enough.

Once vaccines became available, the attitude towards protecting the vulnerable completely changed – you could call it a U-turn. When the Joint Committee on Vaccination and Immunisation published its priorities in December 2020, care home residents, the elderly and the most vulnerable were right at the top of the list. This was broadly welcomed with no mention of the implication that if protecting these vulnerable groups was the right strategy once we had a vaccine then it was the right strategy all along. Vaccination was the most – probably by far the most – effective way we could protect the vulnerable, but it wasn't the only one.

Cocooning was adopted even earlier by offering health care workers priority access to testing to reduce the risk to their charges. Care home workers were also among the top priority groups for vaccination for the same reason. This too was broadly welcomed yet, at the same time, millions of shielders in the wider community continued to be ignored. If cocooning was the right strategy for care home residents then it was the right strategy for the most vulnerable in the community too.

Again, this changed when vaccines became available. The vaccination priority list included social care workers from the outset and was amended in February 2021 to include anyone in the community who received the carers allowance or was designated as a main carer.

Segmenting was an important part of the UK government's novel coronavirus strategy at the other end of the risk spectrum. Allowing children to go to school while most adults were in lockdown was segmenting. Prioritising the re-opening of universities in the autumn of 2020 was also segmenting. If segmenting was the right strategy for schoolchildren and students then it was the right strategy for other segments of the population too.

The realisation has half-dawned. It is widely accepted that the UK's response to novel coronavirus failed care home residents. We now need to accept that it failed the vulnerable in the community as well. We also need to recognise that we failed those at low risk from novel coronavirus too, in their case by imposing not one but three national lockdowns (two in Scotland) that harmed everybody.

I have stressed throughout this book that lockdown was not and was never intended to be a sustainable intervention. Once the virus was firmly established – as it was in the UK from late February 2020 onwards – lockdown could only defer the problem; it was not a solution. Lockdown is what you do when you have failed to keep the virus under control in other, more sustainable ways.

How to avoid lockdown

Let's review the tools we had available to tackle the novel coronavirus epidemic in 2020 and early 2021. I've already mentioned interventions that protect the vulnerable; these include infection control in hospitals, biosecurity for care homes and support for vulnerable people and their shielders in the community. In addition, there were several ways we could suppress novel coronavirus transmission.

First, there were Covid-safe measures: hygiene, PPE, physical distancing, ventilation and so on.

Next, there was self-isolation of cases and their contacts, which depended on effective case finding.

Finally, as what should have been a last resort, there was a series of increasingly restrictive social distancing measures, culminating in lockdown.

We needed to make the best possible use of every other option available to give ourselves the best chance of avoiding lockdowns in 2020 and 2021. It was a trade-off: the better the alternatives were working, the less need there'd be for any social distancing measures, lockdowns included.

It's hard to argue that the UK did its best to avoid lockdown. Despite a huge investment in testing capacity we never fully got to grips with case finding, so large numbers of infections went undetected. We never managed to ensure full compliance with the need to self-isolate either. We were also slow to recognise that we could negate the need for social distancing by making contacts Covid-safe, and we never fully exploited the potential of testing for ensuring Covid-safety.

One reason why we kept turning to social distancing and lockdowns rather than relying on Covid-safe measures was our attitude to risk. The majority of public health officials I encountered in 2020 were extremely risk averse. That attitude was understandable in the early stages of the pandemic and it was defensible if we believed the crisis would be short-lived, but it was no basis for a sustainable novel coronavirus response. We needed interventions that were proportionate, not designed to mitigate the tiniest health risk imaginable without any regard to the impact on our daily lives.

Even where the risk was real, we learned that it was possible to carry out many activities safely. We saw this demonstrated every time we watched sport on our TV screens. The key was Covid-safe protocols integrated with regular testing. If a rugby club can manage that then I don't see why most workplaces

– including mine, a university – cannot. The approach was eventually extended more widely once point-of-care diagnostics became available, but that could have been done much sooner and it could have gone much further.

Social activities too could have been made safer if everyone had access to test-on-request. Instead of embracing this solution we wasted months – and endured two lockdowns – arguing about how to deal with false positives.

There is precedent for slowing the spread of an infectious disease by making contacts safe. For many years, this has been the advice given by public health agencies to prevent HIV/AIDS. HIV is spread mainly by sexual contacts (though also by sharing drug injection equipment). The major focus of anti-AIDS campaigns is on safe sex and knowing your HIV status – which means getting tested – not on banning sex altogether.

Though we could have done more to make our contacts Covid-safe, facilitate self-isolation and protect the vulnerable, I am not claiming that we could have avoided all need for social distancing in 2020 and 2021. I do think, however, that if social distancing measures were required they could have been targeted to deliver maximum public health benefit for minimum indirect harm. That would mean targeting interventions by infection and immunity status, by age and risk group, by specific location (such as a hospital or care home) and by geography. There's nothing new in that – we did many of these things in 2020 – but, yet again, we failed to turn this line of thinking into a coherent strategy and pursue it vigorously enough.

There is already a name for a strategy of this kind: precision public health. It means delivering the right intervention, at the right time, to the right people. It's the exact opposite of a crude, catch-all intervention such as lockdown. If, as is often said, precision is the future of public health then we took a big step backwards in 2020. There were huge differences across the population both in the harms caused by novel coronavirus and the harms caused by lockdown. This was an epidemic crying out for a precision public health approach and it got the opposite.

As well as thinking about where to target our interventions we also needed to think about when to introduce them. The answer is that when the epidemiological data indicated that additional measures are needed – any measures, not just social distancing – these should have been implemented immediately. The guiding principle for avoiding lockdown is that early action can be less drastic action.

This is not quite the same as the frequently heard argument that the UK should have locked down sooner than we did in March 2020 and again in November. I am arguing for a quicker but also a less drastic response, well short of lockdown but still preventing the NHS from being overwhelmed. The rationale is simple: we don't have to impose harsh restrictions in an attempt to drive down levels of infection if we don't let them rise up in the first place.

The triggers for relaxing measures are the mirror image of the ones for imposing them. We knew that increasing numbers of hospitalisations and deaths were 'baked in' – as it was often expressed – once case numbers start to rise. The same is true – but in reverse – when case numbers started to fall. A cautious policy-maker would want to wait until the trend is clear, but I'd argue that a halving of case numbers would have been evidence enough.

Waiting too long compounds the harms done by the lockdown. You will recall that, thanks to the way exponential decay works, the public health benefit of keeping restrictions in place falls off steeply over time, making long lockdowns much harder to justify than short ones.

My emphasis on removing lockdown restrictions as quickly as possible may sound surprising – after all, it's frequently argued that the UK made a mistake by exiting from the March 2020 lockdown too soon. This became the rationale for the slow and overcautious exit from the January 2021 lockdown. It's not correct though. The problem isn't when you start relaxing restrictions, it's where you end up. If you relax too far then the epidemic will get out of control again.

I carefully didn't say 'too far, *too soon*' in the previous paragraph. Novel coronavirus has no memory and it doesn't reward us for being patient. If the R number is above one then the epidemic will take off, regardless of how long it has taken to reach that point. We saw this play out in the gradual lifting of restrictions after the first wave – being patient then didn't spare us from the second wave. There is some merit in unlocking gradually: it makes it easier to go back if the situation changes and it allows more time for alternative measures to be put in place that reduce or remove the need for social distancing in the future. There is no merit in waiting too long to begin the process.

The question of which social distancing measures can be safely relaxed after a lockdown is perhaps the toughest to answer of all. There was good evidence even at the time of the first lockdown that we could have removed some restrictions – especially those on outdoor activities and on schools – much more quickly than we did.

Nor do I accept that the instruction to stay at home was justified beyond the first six weeks. The great majority of individuals, employers and organisations could have been advised, supported and trusted to take precautions to reduce the risk. After all, many day-to-day activities could be made safe and were: several million people did not stay at home because their work was considered 'essential'. This wouldn't be normal life – particularly with regard to social mixing and contacts between households – but it would have been considerably closer to it, perhaps with voluntary social distancing guidelines replacing legal enforcement.

There's an ideological dimension to any preference for voluntary over enforced social distancing, and this coloured discussions of the topic, even among scientists. Yet there's a practical scientific question to be asked as well. We need to know whether the collective impact of voluntary social distancing would be enough to meet the policy objective of not overwhelming the NHS.

You will recall that this approach worked reasonably well in Sweden, but it was claimed that it wouldn't work in the UK and we didn't try. This was despite polls consistently showing that people were greatly concerned about Covid-19 and willing to change their behaviour to reduce the risk to themselves and others. Indeed, it's hard to imagine that everyone would ignore the state of the pandemic around them. So, although no formal analysis has been done, I suspect that voluntary behavioural change could have had a substantial impact if we'd adopted that approach in the UK. It may not have been enough by itself, but it didn't need to be, there were other measures we could use as well.

By the time of the second and third lockdowns we had all we needed to put in place a more comprehensive and sustainable approach to tackling the novel coronavirus epidemic. What was lacking was proper planning for and urgent investment in alternatives to social distancing (regardless of this being enforced or voluntary). I think the UK would have done that planning and made that investment had we accepted from the outset that the virus was here to stay. We would have realised that sustainability was more important than using lockdown to temporarily drive down the incidence of infection to low levels only to have it rebound again.

Let me sum up what should have happened. Above all, we should have done far more to protect those most vulnerable to novel coronavirus – that would have reduced the mortality rate regardless of whatever else we did. Social distancing measures should have been introduced more quickly when needed and removed more quickly too. There should have been a concerted effort to shift rapidly from mandatory social distancing to Covid-safe measures and self-isolation of cases and contacts as the preferred means of suppressing the virus, coupled with a mass testing programme designed to find cases and make activities safer and augmented by voluntary social distancing based on public health advice.

We could have done all this in 2020 and 2021 – no hindsight required. If we had, the UK would have spent far less time in lockdown and still suffered a substantially lower burden of death and disease. That's why I'm a lockdown sceptic.

Mistakes made and lessons learned

There will be many lessons learned from the novel coronavirus pandemic. Numerous issues will be raised by the inevitable inquiries as we reflect on what went wrong in the UK and internationally. I am most interested in what prevented us from taking those steps that would have both saved lives and prevented lockdowns. As I see it, we made three cardinal errors: we repeatedly succumbed to optimism bias, we fixated on lockdown, and we focused too narrowly on the public health harms caused by the pandemic.

The first error was the failure by government – and some advisors too – to accept the scale, severity and duration of the unfolding crisis. They were warned. If you turn back to Chapter 1 you'll find the advice that I and my colleagues were giving in January 2020. I think you'll agree that our advice was accurate. The one point I overstated was the potential impact on mortality. That's because I was using the World Health Organization's initial estimate of the infection fatality rate which – thankfully for all of us – turned out to be too high. Even so, the mortality rate in the UK had doubled by the end of March 2020. In all other respects, what we said would happen did happen.

The reason we were able to give such accurate advice at such an early stage – less than three weeks after the new virus was first reported – is that this was a textbook pandemic. It progressed over the course of 2020 much as any infectious disease epidemiologist would expect. The science wasn't difficult – my undergraduate classes have worked through these kinds of scenario. The maths was straightforward. The hardest thing to grasp was the sheer scale of the coming crisis.

This reluctance of policy-makers to accept the seriousness of the situation persisted throughout 2020. When I first warned of the pandemic's likely impact I was told that everything was under control. When I said that we needed to do mass testing on a scale of millions I was told I was being unrealistic. While I was talking about a coming second wave the talk in policy circles – at least in Scotland – was all about the pipe dream of elimination. When I said the vaccine roll-out would not 'end this by Easter' – as some were claiming – I was called a pessimist. When I said that the virus was here to stay I was accused of being fatalistic.

There is a phrase for this kind of attitude: optimism bias. Government didn't plan for the pandemic the scientists said was coming because they gambled that the scientists were wrong. I wish we had been, but we weren't. Government didn't plan for a second wave because they imagined the worst was behind us. It wasn't. Government didn't plan for a sustainable response because they wouldn't accept that the virus is here to stay. It is.

The second error was our fixation on suppressing the virus through lockdown as the 'right' way to deal with the pandemic. There are echoes here of Daniel Kahneman's concept of anchoring. Lockdown was the first approach used to tackle novel coronavirus and most of the world latched onto it. Lockdown only made sense in the context of eradication, but when the World Health Organization was forced to abandon its eradication strategy it did not change course and lockdown quickly became the international norm.

Once that happened it was extremely difficult for any country to take a different path. There is comfort in following the crowd even while it is stampeding in the wrong direction. We wouldn't let go of lockdown even after the evidence of the harm it was causing became so compelling that the World Health Organization itself came to reject it. Governments around the world – the UK included – couldn't change course because to do so would be an admission that they'd taken the wrong path to begin with. Instead, they doubled down and imposed lockdowns again.

From a government perspective, lockdown had big advantages: it didn't require any forward planning, there was no need to build capacity in advance and no direct financial cost. All lockdown took was government decree and a modicum of enforcement. It was a lazy solution to a novel coronavirus epidemic as well as a hugely damaging one.

Avoiding lockdown would have required a lot more effort, but the necessary knowledge, technologies and systems were all there well before the end of 2020. The problem was that none of the UK administrations were fully committed to making them work. I suggested on the BBC's *Andrew Marr Show* in March 2021 that we should regard going into lockdown as a failure of policy in the same way that we regarded the NHS being overwhelmed as a failure of policy. If governments and their scientific advisors had taken that stance then we'd surely have spent far less time in lockdown.

The third error was to manage this public health emergency as if it were *only* a public health emergency, as if it were only about the burden of death and disease caused by novel coronavirus. It wasn't. The pandemic had huge ramifications for health care provision beyond Covid-19, for mental health, for education, for the economy and for the well-being of society. We needed to manage all those harms at the same time, yet the dice were always loaded in favour of suppressing novel coronavirus at – almost literally – any cost.

Part of the fault lay with the make-up of the scientific advisory committees. The advisory system was dominated by clinicians and public health specialists who weren't looking at the bigger picture – they weren't asked to do so and weren't competent to do so if they had been. Which is why they kept recommending lockdown.

We didn't hear many other voices. The loudest voices in opposition to lockdown in the UK came from educationalists concerned about the long-term impact of closing schools on our children. We didn't hear enough from economists, especially those who could speak to the long-term impact of the damage caused to the economy, to health care provision and to education.

Nor did we hear enough from ethicists. Why wasn't there a much noisier debate about the ethics of the UK's pandemic response? As a parent, I am deeply uncomfortable with the fact that our strategy to tackle novel coronavirus did such serious harm to children and young adults. We deprived them of their education, jobs and normal existence, as well as damaging their future prospects and leaving them to inherit a record-breaking mountain of public debt. All this to protect the NHS from a disease that is a far, far greater threat to the elderly, frail and infirm than to the young and healthy, an objective that could have been achieved by other means. If there was a justification for burdening the younger generations in this way then I'd have liked to have heard it. To me, it seems morally wrong.

Once novel coronavirus took hold it was inevitable that 2020 would be an exceptionally difficult year, but there were options at every turn. If we'd trusted ourselves, our data, our systems and our science then I've no doubt we'd have made better choices. Lives and livelihoods could have been saved, lockdowns largely avoided and far less damage done. Instead, we were mesmerised by the once-in-a-century scale of the emergency and succeeded only in making a global crisis even worse. In short, we panicked.

At the beginning of this book I listed the things I did not expect to happen during the pandemic. I did not expect that elementary principles of epidemiology would be misunderstood and ignored, that tried-and-trusted approaches to public health would be pushed aside, that so many scientists would abandon their objectivity, or that plain common sense will be a casualty of the crisis. Yet – as I've explained – these things did happen, and we have all seen the result.

I didn't expect the world to go mad. But it did.

CHAPTER 18

DISEASE X

I spent two cold, grey days in January 2017 at a World Health Organization meeting in Geneva. I was one of a number of experts brought in from all over the world to advise on viral pandemic threats that should be high priorities for scientific research. After lengthy discussions, our committee came up with a sensible-sounding list that included Ebola, Lassa fever, MERS and SARS.

Myself and a few others present – including Peter Daszak, Director of Ecohealth Alliance, a New York-based non-profit organisation – wanted to add one more item, and we lobbied hard for it to be included on the list. We thought the next pandemic was just as likely to be caused by a virus that we didn't even know about yet. We succeeded in persuading the committee and the concept of 'Disease X' was born.

The next pandemic threat

The committee's conclusions were published on the World Health Organization website and received plenty of attention. The following year the committee took the Disease X concept one step further by thinking more about the kind of disease it was most likely to be. They came up with a small number of suggestions that included *highly pathogenic coronaviral diseases* *other than MERS and SARS.* It is hard to imagine a more accurate prediction: no-one can say the world wasn't warned. The World Health Organization's recommendations for further research didn't have much impact though. Most of the world remained fixated on flu.

It's right that we should be concerned about flu. There was an influenza pandemic – the swine flu pandemic – as recently as 2009. This involved a relatively mild strain of a type called H1N1 and the public health impact was, by pandemic standards, quite minor. One of the things that kept swine flu in check was that many older people had been exposed to related strains that gave them at least partial immunity. That was fortunate: Spanish flu – which killed around fifty million people worldwide in 1918–20 – was also caused by a H1N1 strain. By grim coincidence, a new H1N1 strain – of unknown pathogenicity – was reported from China in 2020, but has not spread widely so far.

There are even more lethal strains of influenza, collectively known as bird flus. Fortunately, these strains do not spread between humans, though scientists have been concerned for a long time that they might evolve the ability to do so. If that happened we could be facing something at least as bad as Spanish flu. This is why so much effort had gone into preparing for an influenza pandemic and we need to stay prepared.

Coronaviruses might have played second fiddle to influenza in the past, but that's all changed now. We learned very quickly that novel coronavirus is a close relative of the SARS coronavirus, so close it is considered to be a strain of the same virus species. SARS was considerably more lethal than novel coronavirus – it had a case fatality rate of 11%, more than ten times higher. It's fair to say that the world had a narrow escape in 2003 when an incipient SARS pandemic was halted and the virus was eradicated – one of the World Health Organization's proudest moments.

We were not so fortunate in 2020. This means that there has been a near-pandemic and a full-blown pandemic caused by SARS-like viruses in less than two decades. You do not need to be an expert on emerging infectious diseases to work out that we should already be worrying about SARS-3.

We must not, however, fall into the trap of preparing only for another pandemic caused by a SAR-like coronavirus or a new strain of influenza. Public health has a history of preparing for the pandemic we've just had, not the one we have next. The report of that World Health Organization meeting in Geneva in 2017 was intended to make the point that the next pandemic virus could be something quite different and present distinct challenges.

My research team has been working on the 'what next?' problem for over twenty years and we've made some progress. Back in 2001 we published the first systematic study showing that the majority of recently emerged infectious diseases have been viruses acquired from animals.

The technical term for an infection acquired from non-human animals is a zoonosis. SARS, Ebola and novel coronavirus are all viruses of zoonotic origin. Two or three new ones are found in humans in an average year, but most of them are only a minor threat to public health. Scientists are still arguing about how many potential new human viruses are still 'out there', but we do agree that it helps to try to narrow down which ones we should be most concerned about and we are beginning to find ways of doing that.

In 2016 my team published a list of around thirty zoonotic viruses that were rare but we knew could spread in human populations. We deemed these to be significant epidemic risks. While we were doing this work three viruses on the list – Chikungunya and Zika (both mosquito-borne viruses) and Ebola – all erupted into major epidemics, so our approach seems to work. All the viruses on the World Health Organization's research priorities list were already on ours, so we were in good agreement with our colleagues too.

As well as identifying known viruses that are a pandemic threat, we are also thinking hard about Disease X. In 2020 – building on earlier work by Peter Daszak and others – we published a map of the world showing where new human viruses had been discovered and where we estimated the next ones were most likely to appear. All of the eighteen new human viruses reported in the preceding decade were discovered in the locations we designated 'very high risk' or 'high risk'. Wuhan, China, was one of those locations.

In 2019 we published a pre-print reporting an analysis of the phylogenies of almost two thousand different viruses based on their genome sequences. We reported that new human viruses with epidemic potential – candidates for Disease X – are most likely to come from lineages that already contained viruses with epidemic potential, but emerge as separate introductions into humans from an animal source. Again, our approach works: novel coronavirus fits that pattern perfectly, it is closely related to another virus – SARS – that has epidemic potential but it jumped into humans independently, probably from horseshoe bats.

Emergence of new viruses

Much has been written about how new viruses emerge. I first became interested in the subject after reading Laurie Garret's book *The Coming Plague* in 1995. The emergence of any new virus – be it Ebola, MERS or anything else – is a fascinating story in its own right. No-one could predict the precise sequence of events that led to the emergence of a virus like Nipah, a story that involves bats, mangoes and a new way of farming pigs in Malaysia. I am sure that the story of the emergence of novel coronavirus will prove just as complicated, if we ever find out.

We believe we know when novel coronavirus emerged. Analysis of viral genomes supports the idea that the virus first jumped into humans in December 2019, though it is possible that it happened a month or two earlier than that. Some have proposed that it did emerge earlier and this was covered up. There are even suggestions that it had arrived in the UK by that time. I'm doubtful; this isn't a virus anyone could hide for long. December 2019 looks right to me.

That's when, but what about where? We know that many of those early cases of novel coronavirus were linked to the South China Seafood wholesale market in Wuhan, China. This is a so-called 'wet' market, selling a wide variety of live animals. The market is located just ten kilometres from the Wuhan Institute of Virology, China's leading coronavirus research laboratory. Coincidences happen and there may be no link between the virus outbreak and the lab, but this is a big one and I doubt China will ever throw off the suspicion that there was a connection.

Early in 2021 the World Health Organization dispatched a team of scientists to Wuhan to investigate. The report they published in February of that year received – to put it politely – a mixed reception. I accept that no-one could describe the report as definitive, but it does contain a lot of helpful information and some interesting analysis.

The team's preferred explanation for the emergence of novel coronavirus was that it was a spill-over event involving contact with an infected animal outside the laboratory, perhaps in the wet market, which is quite possible. Another suggestion was that it might have been brought into China in frozen food, which seems highly improbable. The report confirmed that there was no evidence that the virus was a product of genetic engineering – we'd see the signature of that in the genome, and we don't (though the idea still re-surfaces from time to time). In its conclusion, the report went further by more or less dismissing the notion that the Institute of Virology was directly involved in any way at all.

That doesn't necessarily get the Institute off the hook though. An intriguing hypothesis was put forward by *The Sunday Times* in July 2020. They reported that researchers from the Wuhan Virology Institute sometimes visited bat colonies in southern China to collect samples. They suggested that, despite whatever precautions were taken, one of the researchers had been exposed to novel coronavirus during one of these visits and brought it back to Wuhan. This seems a plausible scenario but, as I said, we may never know for certain. We know all too well what happened afterwards.

Thinking about next time

What happened afterwards was a disaster, but was far from the worst pandemic that the world could have faced. I explained in Chapter 1 that the scale, speed and severity of a pandemic are largely determined by R0 (the basic reproduction number), the generation time and the infection fatality rate.

Novel coronavirus isn't exceptional by any of those measures. There are viruses with much bigger values of R0: measles, for example. There are viruses with faster generation times, such as influenza. There are viruses with far worse infection fatality rates; Ebola is one. It's easy to imagine a scenario where even the harshest measures we took in 2020–21 will not be enough. What would we do then?

That's a question we need to start thinking about now. As our collective experience of 2020 has illustrated so starkly, every one of us has a profound interest in making sure that we have a good answer. Our governments and international agencies need to have the knowledge, the tools and the systems in place to respond quickly and effectively in the face of future threats from emerging infectious diseases. However, I won't conclude with a grand vision for meeting that challenge – that would be a whole new book in its own right. I shall end with a personal story.

In April 2020 I was asked by Emma Barnett on Radio 5 Live how it felt to see the crisis that I had been concerned about for so many years playing out in real life on a global scale. In effect, Emma was asking how did it feel to have been right all along. I made some vague reply, but I was consciously dodging the question. Here's how it truly felt.

In Chapter 1 I explained how I came to write a series of e-mails to the Chief Medical Officer of Scotland in January 2020. One of these referred to a hot-off-the-press estimate of the infection fatality rate for the new virus, just below 5%. This virus was highly transmissible – much more so than flu – and if one in twenty of those infected were going to die then we were facing a global catastrophe. I suspected that 5% was an overestimate. On the other hand, I knew that the new virus was closely related to SARS, which was even more deadly, so I couldn't be sure. On the evening of January 24th I was talking this through with my wife, broke off in mid-sentence, put my arms around her, and burst into tears. That's how it truly felt.

The first year of the novel coronavirus pandemic was a desperate time for humanity, but thankfully it wasn't as bad as I'd feared that January night in Edinburgh. This won't be the end of the story though. There will be another pandemic sooner or later and that one may well be worse. We need to be prepared.