ROYAL STATISTICAL SOCIETY

DATA | EVIDENCE | DECISIONS

COMMUNICATION DURING THE PANDEMIC: DATA, STATISTICAL ANALYSES AND MODELLING

Work-in-progress report based on RSS event taking place on 5 April 2022

Summary – lessons learned about communication of data, statistical analyses and modelling

This report documents the discussion at the first of the RSS's Covid-19 evidence sessions.[1] Before the session, we identified six questions designed to discover what lessons could be learnt around communication of data, statistical analyses and modelling – both so that we can learn from what went well, as well as reflecting on where there are areas for improvement. We sought to bring in a wide range of views during the discussion at the event, and these are reported below. However, even over a two-hour meeting, only so many people were able to speak, so this is intended as a reflection of views expressed during the meeting and should not be read as representing the views of the RSS.

The six questions we explored were:

1. How effectively have data providers, government representatives, the media, and commentators communicated with the public on issues involving data, statistical analyses and modelling?
2. How have different statistical definitions impacted how information is communicated to the public?
3. How well has uncertainty been communicated, including in modelling?
4. How good were the visualisations used to communicate with the public?
5. How has data being accessible to the public improved communication – and was data accessible enough for this purpose?
6. Was enough done to try and prevent information becoming misinformation?

List of speakers

This is a summary of the discussion at the first RSS Covid evidence session – the views expressed are those of the speakers. The list of contributors is below:

Main Speakers

- **John Burn-Murdoch** – Chief Data Reporter, Financial Times (FT)
- **Hannah Ritchie** – Head of Research, Our World in Data
- **Oliver Johnson** – Professor of Information Theory, University of Bristol
- **Will Moy** – Chief Executive Officer, Full Fact

Contributing speakers

- **Anthony Masters** – RSS Statistical Ambassador and co-author of Covid by Numbers
- **Simon Briscoe** – Member of RSS Covid-19 Task Force, director at the Data Analysis Bureau
- **Kevin McConway** – Emeritus Professor of Applied Statistics, Open University
- **Liza Hadley** – Researcher, Cambridge University
- **Tom Chivers** – Science writer, i
- **Jen Rogers** -- Vice President for Statistical Research and Consultancy, PHASTAR

[1] This document is a work in progress – a final version will be published in 2023. If you notice errors or omissions please email policy@rss.org.uk.

- **Christina Pagel** – Member of Independent Sage and director, UCL's Clinical Operational Research Unit
- **Paula McCloud** – Head of profession for statistics, FCDO
- **Ben Swallow** – Statistics lecturer, University of Glasgow
- **Jonathan Smyth-Renshaw** – Owner, Jonathan Smyth-Renshaw & Associates Limited
- **Matus Medo** – Data scientist, University of Bern's Department of Clinical Research
- **Ting-Li Su** – Lecturer in Oral Health Statistics, University of Manchester

1. **How effectively have data providers, government representatives, the media and commentators communicated with the public on issues involving data, statistical analyses and modelling?**

Data, statistical analyses and modelling have an important role to play in public health policy during a pandemic – part of their use lies in helping decision-makers understand what is happening and what the impact of their decisions may be. But, especially when the response to the pandemic involves asking the public to make sacrifices, these tools also have an important role to play in communication with the public. Presenting the public with the same information that decision-makers have used to inform their decisions increases the likelihood that people will support their government's decisions. So, we are interested in how effective this aspect of communication was throughout the pandemic.

This question is the broadest of those that we asked at our event and aspects of it – particularly how statistical definitions and uncertainty were communicated – are picked up on in detail in questions two and three.

Hannah Ritchie looked at how the UK had performed in a global context. Since February/March 2020, Our World in Data have published global Covid data (updated every day), including: confirmed cases; confirmed deaths; testing; vaccinations; hospitalisations; excess mortality; and, government responses. This means that their team have seen a range of good and bad practices from across the world. As an example of good practice, Chile provided clean daily updates on GitHub, and other countries have done this via an API. Some countries, on the other hand, still provide irregular updates via YouTube videos (that need to be paused so that the data can be manually copied). There is a wide spectrum globally regarding how data communication has been managed.

Based on this experience, **Ritchie** proposed a few core principles for communicating data:

- Make the **definitions** clear
- Pick the **appropriate metrics**
- Make the data **open and accessible**
- Keep it **simple**
- Provide **context / perspective**
- Be **honest** and transparent

Will Moy argued that the statistical response had been overwhelmingly successful – it was more improvised than was probably ideal, but many people put in a great deal of effort to make the UK's response good. There are lessons to learn – but they should be viewed as such, rather than as failings. **Moy** spoke about the Swiss cheese model of information – you need to collect the right information, analyse it well and communicate it effectively. Those communications need to reach people who can make decisions. If any of those steps go wrong, then decision-makers will not be able to use the best evidence. This is a huge challenge and it is important to recognise just how much has been successful during the pandemic. When the vaccine came out in November 2020, only half of the UK population were planning to get a vaccine, by April 2021 that was up to two-thirds. In April 2022 nine out of ten people had at least one dose. This demonstrates life-saving decision-making based on evidence, and is something to celebrate.

However, when we think about the journey from good evidence to understanding, it is possible to fail at any stage. During the pandemic we saw examples of failure at each stage. The care home catastrophe was an example of

failing at the first stage – there was no process to collect daily data from care providers and we did not know how many people were receiving care in the social care sector, and that cost lives. It was also a predictable problem – the Office for Statistics Regulation (OSR) had raised this in 2019 saying that statistics can only paint a partial picture of what happens to people in the care sector.

At the other end of the process, getting people to act on good information sometimes means using simple slogans. A survey showed that between late March and mid-May 2021 the proportion of adults who found the government's communication about what to do 'very clear' or 'fairly clear' dropped from 90% to 56% -- that's when the messaging urging people to stay home changed to messaging urging people to stay alert. Some government messaging may be crude, but sometimes that is what is needed to turn data into actionable understanding.

Moy offered five reflections on what we've learned from the pandemic.

1. Communicating reliable information is a team effort. These teams need to be funded.
2. Statisticians are powerful and what they choose to measure matters. For all that health statistics have been well communicated during the pandemic, we have failed to communicate related areas, such as migration statistics. The interconnectedness of the UK with the rest of the world and what that means for pandemic management has not been communicated. There has also been nothing like enough clarity on what has happened with the economy – 400,000 people have left the labour market, and this has been underdiscussed. The mismanagement of the exams mess also requires lessons to be learned.
3. We need to listen to our audiences. Good communication starts with listening. The Ask Full Fact project demonstrates how necessary it is to understand our audience in order to communicate with them. A number of groups have been poorly served during the pandemic. Pregnant women have been given confusing information about vaccines. Black people had lower take-up rates of vaccines and were poorly listened to. NHS staff did not have a high enough rate of take-up of vaccines. That came as a huge shock, but indicates there was a need to better engage with that audience. Clinically extremely vulnerable people also needed better quality information. What all these groups have in common is that they are not standard statistical groupings – so we start to see blindspots and poor real world outcomes.
4. The way we listen is by bringing statistics and social research together, which we have not done adequately.
5. A lot of the pandemic communication response has been underpinned by censorship – preventing people sharing certain types of information without a clear legal basis or democratic oversight. We should not underestimate how much the communication landscape has been influenced by control rather than persuasion and we should be very cautious about this. We should be calling in the online safety bill for clearer oversight of this.

Tom Chivers looked at why the media sometimes got communication of data wrong. There was a tendency throughout to focus on worst case scenarios. But there was also the opposite mistake – trying to give wise-sounding "don't panic" messaging. It's not because journalists are stupid or evil (though, they may be innumerate) – there is a fundamental incentive problem for journalists. The incentive is at least partly to make sure that people buy your newspapers or watch your programmes – the exception would be a publication like the FT where the readers are likely to be making financial decisions based on what they read. Questions about the quality of the data being reported can be seen as getting in the way of a good story. This is a tricky problem to get around – there is no point writing a perfectly accurate story that nobody reads. The incentives are genuinely competing and it is not clear how you resolve it.

Christina Pagel emphasised the importance of communicating context to the public. The public have undoubtedly got better at understanding graphs over the pandemic, but data still do not speak for themselves even with the best charts, and there is always going to be vital context that needs to be understood. This has been an issue from the very beginning of the pandemic – eg, understanding and communicating the lags between infections, reported cases, admissions and deaths. The most egregious error is using deaths as the leading indicator, which **Pagel** argues is entirely wrong. There's also the softer context of how all of those indicators change with availability of testing, differential vaccination rates, different cases and different age groups – all of those things have to be understood when you're looking at the headline figures. A really important example is the Omicron variant - you can

have a situation where the overall cases are declining or flat even after the new variant has started exponential growth. To focus on overall cases and say things are going in the right direction could give a totally false perspective on the state of the pandemic and what might happen within a few weeks. So, in this example, it is important to separate out data by variant when presenting Covid statistics. This is especially important as we're doing a lot less real-time identification and sequencing of cases as we're moving away from PCR tests. Finally it is important to try and give people the context of what's *not* on the daily dashboard – eg, things like long Covid and the social impact of Covid and the measures to control it. There is a risk that these do not factor into policy making because they're not a key outcome that we're used to. So**, Pagel** urges people to always provide context and explain what is missing.

Kevin McConway drew on his experience of helping journalists understand the torrent of Covid statistics and on his experience as one of two RSS nominees on an advisory group on data presentation for the Cabinet Office. He makes two points. First, he has been surprised by how little data analysts tend to know about communication. It is not that everyone needs to be an expert in it, but there should be expertise within organisations. ONS is usually excellent in this respect, but others are not. Second, on transparency: it's admirable that the government decided to publish nearly all the papers from SAGE, including the minutes and so on. That is a lot better than keeping it all hidden. But there is a difference between transparency and intelligent transparency, which says that information should accessible, comprehensible, usable, and assessable. The SAGE documents are accessible enough – at least you can find them – but for the public and for many journalists they are not very usable. They are barely comprehensible and so mostly not assessable. In an emergency the temptation is to get on with data collection, data analysis and modelling, but if people can't truly make sense of what you're doing, you won't be trustworthy so you won't be trusted. You have to start early with earning trust as well – otherwise the mistrust and misinformation will already have begun. Intelligent transparency of SAGE documents should have been part of the system – without being subordinate to the government choosing the messaging. Informing not persuading is a key principle of communication during a public health crisis and the work of SAGE needed to be explained.

2. How have different statistical definitions impacted how information is communicated to the public?

Hannah Ritchie, in her presentation, proposed that a core principle for communicating data should be to ensure that the definitions used are clear. As an example, in February 2020 media outlets (eg, BBC, the Guardian) were reporting that 'cases' had doubled from four to eight. We have seen throughout the pandemic that, even before government policy changes, people adapt their behaviour based on how they perceive the severity of the situation. The downside of this is that when the media is reporting low numbers of *confirmed* cases as a low number of 'cases', people did not take it seriously or change their behaviour. So, early on in the pandemic there was a failure to communicate how serious the situation was because it was not made clear that there was a distinction between confirmed cases and the actual number of cases. This is especially important given that people generally struggle to intuitively understand non-linear growth. Our World in Data reacted to this by ensuring that their charts always showed the confirmed numbers of cases and, to give context, built a global testing dataset.

This was bad at the start of the pandemic, but it is a problem that is continuing even in 2022. News outlets still routinely report total number of confirmed cases and deaths as the total number of deaths. The below image, taken from the Guardian on 26 April, shows this.

Latest global Covid-19 data

Total cases	New daily cases
510.24m	**576,087**
Total deaths	New daily deaths
6.22m	**2,412**
% fully vaccinated	New daily vaccinations
58.9%	**11.4m**

% fully vaccinated figure refers to people who have received all doses prescribed by the initial vaccination protocol, as a proportion of the total world population. Data from JHU CSSE Covid-19 Data and Our World in Data at 10.47 on 26 April 2022

In this case, this is despite using Our World in Data's dataset and despite Our World in Data contacting the Guardian to report that this is not the best way to present the data.[2]

Ritchie proposed that a possible solution to this problem would have been for the World Health Organisation (WHO) to collect and publish excess death figures from an early stage in the pandemic – as soon as it was clear that this was the most reliable way to measure the impact in terms of deaths. Our World in Data had a call with them in Spring 2020 to make this point – especially as it looked like some poorer countries weren't being hit at all. At the time the WHO assured them that it was looking at this, but two years on they still present the figure of six million deaths.

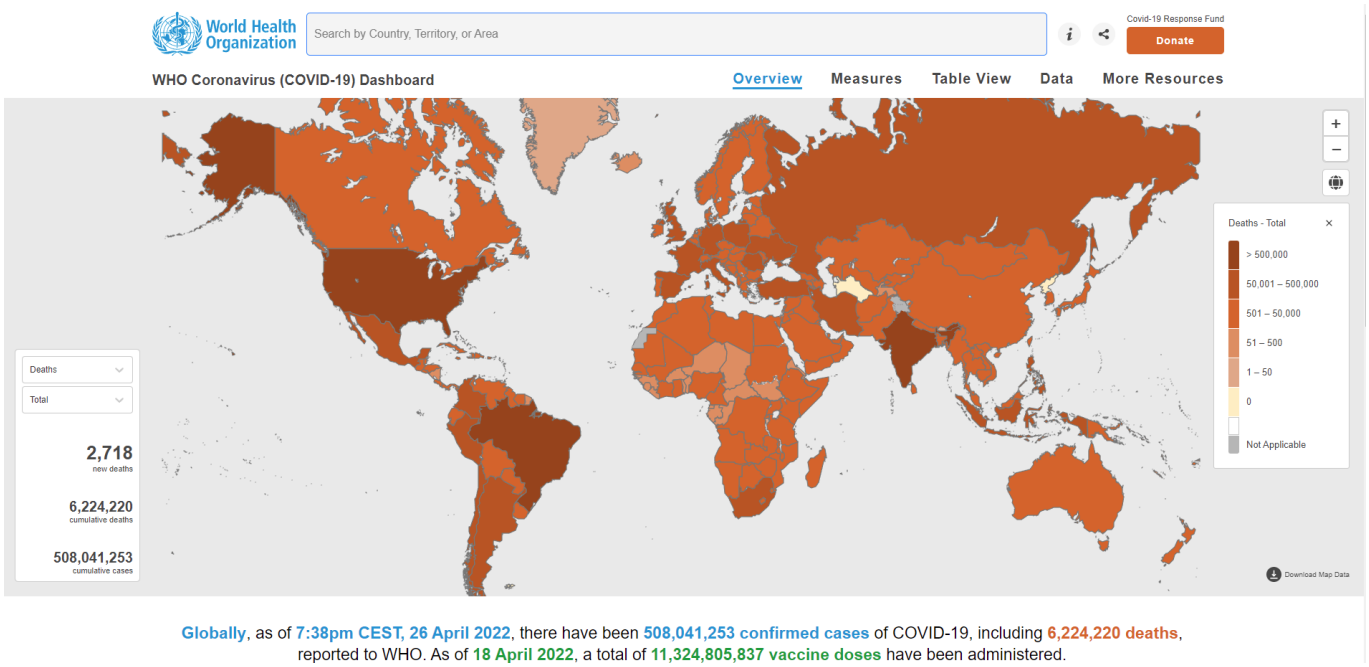

Globally, as of 7:38pm CEST, 26 April 2022, there have been 508,041,253 confirmed cases of COVID-19, including 6,224,220 deaths, reported to WHO. As of 18 April 2022, a total of 11,324,805,837 vaccine doses have been administered.

[2] The Guardian is not alone in this – the practice is widespread.

Taken from https://covid19.who.int/ on 27 April 2022.

We have had to rely on the Economist for some sense of perspective on the situation. Looking at the UK, we have seen that the definition of a death from Covid-19 in England, Scotland and Wales has shifted from registration on the death certificate, to anyone with a positive test, to anyone with a positive test in the past 28 days. It's fine to shift definitions over time, especially if you think you are improving the measure. The current definition does allow for some common objections – eg, people who have had a positive test and then get hit by a bus. But this is uncommon and, by Autumn 2020, 93% of those with Covid-19 mentioned on their death certificate had it as the key underlying cause of death. The UK has not made this definition clear – the number of people with Covid-19 as the key underlying cause of death has not been clearly promoted, leaving a gap in understanding that some people may fill with misinformation.

Another common objection, even with the 'key underlying cause of death' definition, is that people could have died from another cause anyway. This is where excess mortality is important. Excess deaths is obviously not a perfect measure, but it makes it very clear that we have experienced extra deaths due to Covid – as shown in the chart below.

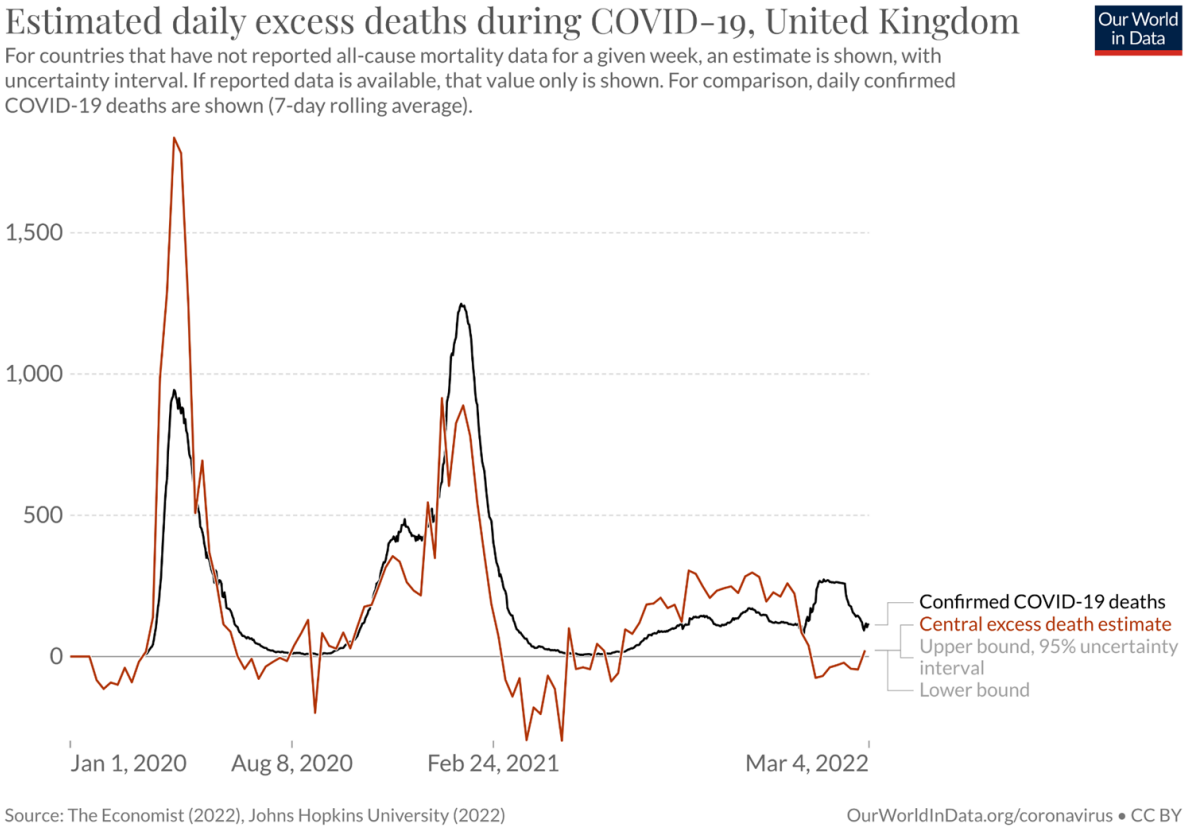

Chart showing excess mortality during the pandemic.

There is no single 'perfect' definition of a Covid death, but there are a range of measures that the government could have used to make it much clearer that: Covid was not 'just' the flu; it was not simply that people dying from Covid would have died from another cause that year anyway; and, most people with Covid on their death certificate died from it as the underlying cause.

As well as confirmed deaths, **Ritchie** also pointed to confirmed cases being a poor choice of metric. There were many examples of world leaders using this statistic in a misleading way. For example, Tanzania's president, John Magufuli, declared Tanzania free of Covid-19 due to not testing for the virus for over a year, before himself dying in unexplained circumstances. At the other end of the spectrum, US President Donald Trump argued that the US was

not actually experiencing a high number of cases – rather they were just performing an uncommonly high number of tests. The more appropriate metric to use is test positivity rate, which provides a more accurate measure of the scale of an outbreak.

Ritchie also suggested that there is an important lesson to be learnt regarding the limitations of static snapshots. This is because with infectious diseases it's the change over time that is most important rather than the situation on any given day (especially given that reported deaths drop at the weekends and peak in the days soon after). Rolling averages (along with being conscious of methodological choices) provide a much more helpful way of understanding data. This is something that the government dashboard is now very good at – and that other countries have been slower at.

Simon Briscoe argued that we can't expect all producers to deliver comprehensive and balanced coverage – for example if they are private sector or a charitable organisation or similar. It is not what is expected of them and it is not their role. It is the duty of public sector bodies to deliver this type of data. This means a single go-to place, to find links to the key data and reports. The ONS and academics received millions of extra pounds but there was no central hub where people could go to find out more information about the data that was in the news that day. This failure says something about the siloes that still pervade the wider public sector. Second, looking at the dissemination of data in the media, some organisations published data – the FT and the Spectator chief among them – but these were shoestring operations, which the public should not have had to rely on. **Briscoe** argued that the BBC with its budget, staff and public duty remit could have done better than it did. A key lesson is that when these crises strike we need public sector organisations to deliver for the wider public and, in his opinion, these organisations did not successfully do this.

Jen Rogers has spoken a lot to numerate journalists and the science media to provide insight into statistics. Her main concerns have been the lack of change in the way that data has been communicated and the lack of change in commentary as the pandemic has evolved. For example, there has been a continuing focus on the estimated daily number of cases – despite uncertainty over whether that statistic is meaningful when you think about the different testing regimens. She has been repeatedly asked about why cases are so much higher than a year before – but the testing situation is so different that it is hard to meaningfully compare different timepoints. There are also numerous commentators who have provided a narrative to improve public understanding and they should be congratulated on that, but not all commentators have acted objectively, and some seem to have an agenda. A lot of commentators have been right by accident at some point during the pandemic and are reluctant to change. But it is important to think about how communicators can be open to changing approach over the course of a pandemic.

Rogers also highlighted two disappointing features of the government's presentation of data. First, the daily update on vehicle use during the first lockdown seemed unnecessary – it felt like data for data's sake and if a statistician is feeling like that, you might expect a similar feeling among the wider public. Secondly, there were a couple of big areas where the government should be held accountable – where they clearly had an agenda and sought to use data to persuade rather than inform. For example, when they were looking at Covid cases in schools and chose to present it as positivity rates rather than cases – they said both ways showed the same thing, but this wasn't true. There has also been a lack of transparency on modelling assumptions – these assumptions could have been peer reviewed.

3. How well has uncertainty been communicated, including in modelling?

John Burn-Murdoch suggested that the modellers themselves had done a reasonable job of conveying uncertainty. For example, consider the below series of charts from modellers at Warwick University.

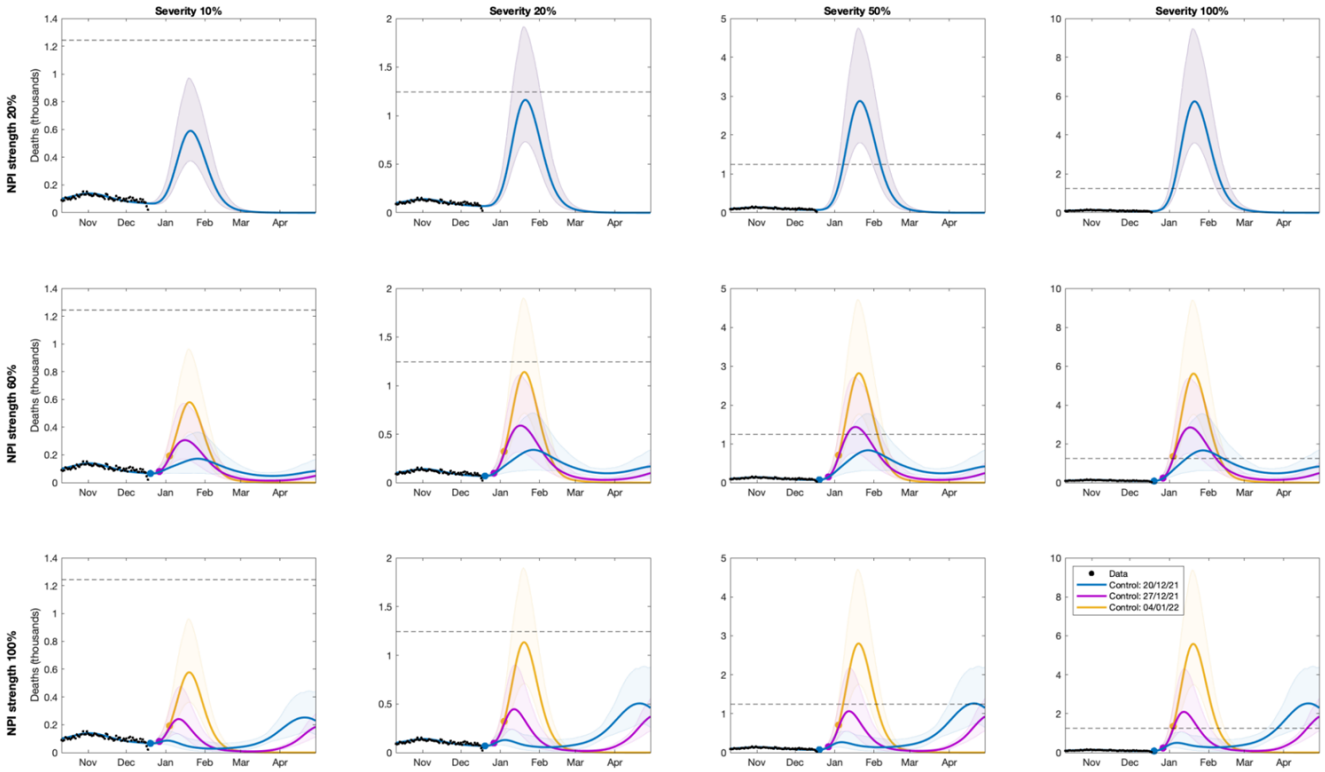

Taken from Keeling et al., *Short-term Projections based on Early Omicron Variant Dynamics in England*, p.12. Showing hospital admissions for the circuit-breaker model with a gradual return to pre-Omicron mixing from 15 January to 15 April 2022.

The figure conveys four different estimates of severity and three different levels of behavioural changes, all plotted in a way that gives an enormous amount of information. It not only shows how much uncertainty there is across different situations, but it also shows how much uncertainty there is within a given situation and how that compares to the information that was available to date.

This allowed media outlets like the Spectator to plot reality against the model, in an effort to push for accountability. One of the examples is below:

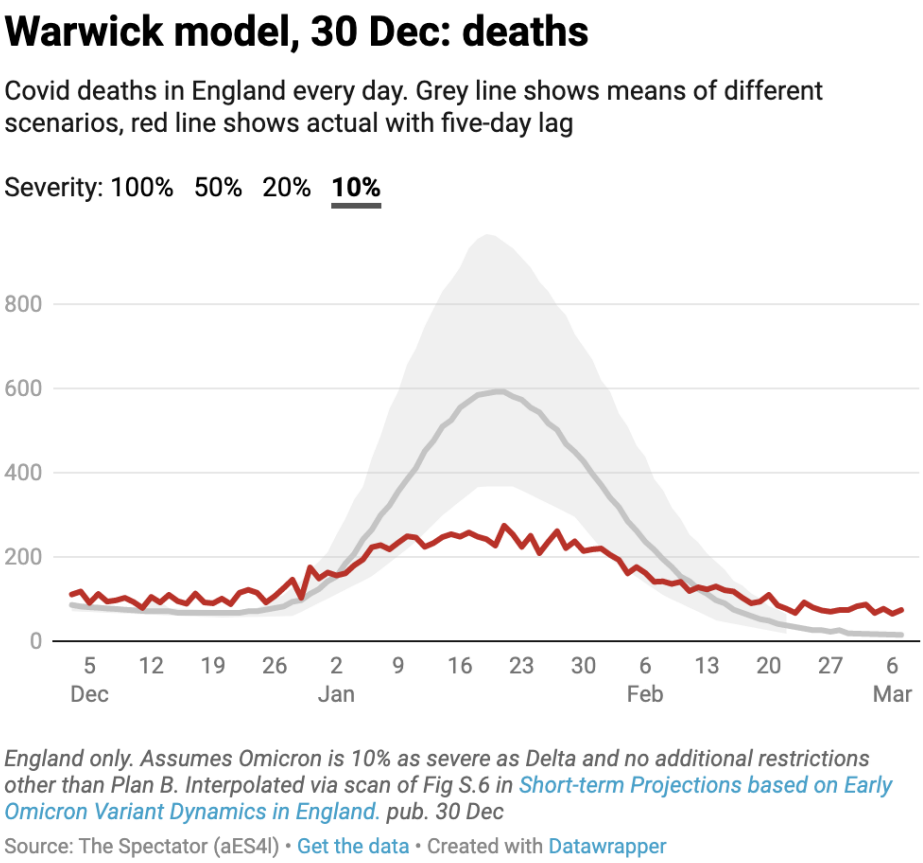

Spectator's reproduction of part of models shown above.

As can be seen here, in one of Warwick's scenarios they slightly overestimated the number of hospital admissions from Omicron, but it was not by a huge amount. The key point is that by providing such a range of scenarios in a transparent manner, they enabled this form of analysis and that was positive.

While modellers generally did a good job of displaying uncertainty, there are other innovative approaches that could be used. Claus Wilkes has argued that hypothetical outcome plots provide an effective alternative to convey uncertainty by showing a variety of different possible modelling outcomes at once. The example below uses random data to show how the same data could produce a variety of different trends. It would be interesting to see if modellers might use more experimental approaches like this to more viscerally show the uncertainty involved in their models.

Burn-Murdoch's example of how a hypothetical outcome plot can communicate uncertainty in an alternative manner. (Press play to run the visualisation).

There is a huge amount of uncertainty around numbers and what they mean. The majority of the public did not get their information directly from the producers of statistics – they relied on government spokespeople and the media to communicate the key messages to them. The uncertainty inherent in the statistics, data analysis and modelling presented a significant challenge for government and the media in how they communicated the information that was informing government decisions. Communicating uncertainty is hard and there is likely to be a limit to the extent to which a typical member of the public can absorb and understand information about uncertainty – but attempting to communicate uncertainty is important to help build trust, and we should attempt to learn lessons from what worked well and what worked less well.

Burn-Murdoch reflected on how effectively the media had communicated uncertainty. He gave examples from December 2021, when the Omicron variant was just emerging. SPI-M (the subgroup of SAGE that gave expert advice on infectious disease modelling and epidemiology) had released a consensus statement on December 15 saying that at the peak of the Omicron wave there were likely to be between 600 to 6,000 deaths per day. This was presented in several newspapers as projections of more than 5,000 deaths per day. Eg:

> *Worst-case scenario modelling by the team…suggests the UK could experience more than 5,000 deaths a day at the peak of the Omicron wave. (*The Guardian*)*
>
> *There could be 5,000 Omicron deaths a day this winter without more restrictions*. (Daily Mail)
>
> *Horrifying prediction of 5,000 Omicron deaths a day unless restrictions tightened.* (Mirror)

These are just some examples of this type of statement. Speaking not just as an observer of this trend, but as someone who has been involved discussions with editors, there is a tendency to say "what is *the* number" and to hone in on it. This is to the detriment of everyone's understanding of how much uncertainty is involved in modelling.

Anthony Masters highlighted the four ways that the Government Statistical Service recommends dealing with uncertainty. First, you can describe the uncertainty with a statement – eg say that it is a survey sample so won't

give a precise estimate. Second, you can describe it by giving a range. Third, you can illustrate the uncertainty using visualisations. Or, fourth, you can tell people what they can and cannot conclude from the data. There were examples throughout the pandemic of good practice in data visualisation from official sources– eg the use of a waffle chart to show vaccine effectiveness. There were some visualisations with multiple overlapping uncertainty ribbons, showing clearly that some observables are consistent with a variety of models. **Masters** argued that there needs to be a better understanding of how the public interpret uncertainty in visualisations.

Ben Swallow comments on the fact that in public discourse, uncertainty is often seen as a negative quantity: either as a lack of knowledge about a process of interest, or the inability to estimate with a degree of accuracy. However it also represents a very important quality, which is the natural variability in the process that statisticians are trying to model. Portraying uncertainty as more than just an error, and as an important quantity for inference and decision-making in its own right is really important and something that statisticians as a community should be encouraging. While the use of intervals in addition to point estimates is definitely a step forward, full uncertainty quantification should mean more than just the population mean plus or minus standard error. Studying variability, particularly at increasingly high resolutions (including allowing for the variability across individuals, across spatial regions and across temporal locations, as well as across variants), and not just aggregating to very low resolutions in both space and time, allows us to gain much better understanding of what's really going on in the wider community. In signal processing it is often the variation that actually determines the information rather than the general trend itself. **Swallow** has been working in collaboration with individuals such as Jasmina Panovska-Griffiths at the University of Oxford to look at the variation in some of the metrics of the state of the pandemic across different spatial resolutions. Their work has shown the important information that we can gain when we don't just aggregate to a single value, but use the interval to try and represent the variability of what is being measured.

4. How good were the visualisations used to communicate with the public?

Burn-Murdoch looked at the Covid-19 dashboard – arguing that, especially compared to other countries, the outputs were exceptional. He gave the example of the below heatmap of cases by age over time.

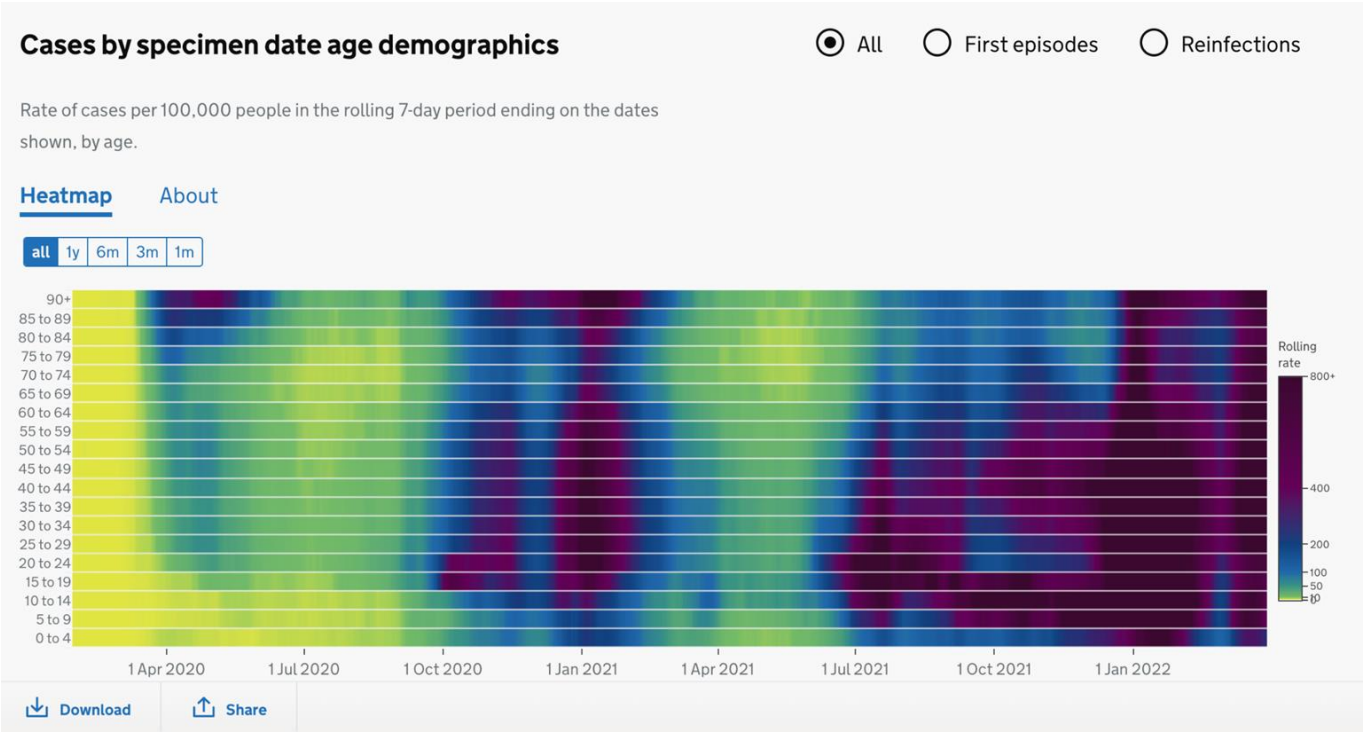

Covid-19 dashboard showing number of cases by age over time.

The granularity of this data and how clearly it is communicated is absolutely world class – there are not many countries who even have this data available, let alone many who communicate it in such a visually clear and striking way.

Another good example is the visualisation below from the time that the Delta wave landed. How age influenced case rates was a critical question at this time – cases were rising, but after the vaccination, were they rising as quickly among older people who had been vaccinated as among younger people who hadn't? The dashboard provided an age breakdown, which allowed us to answer this question in situ on the dashboard.

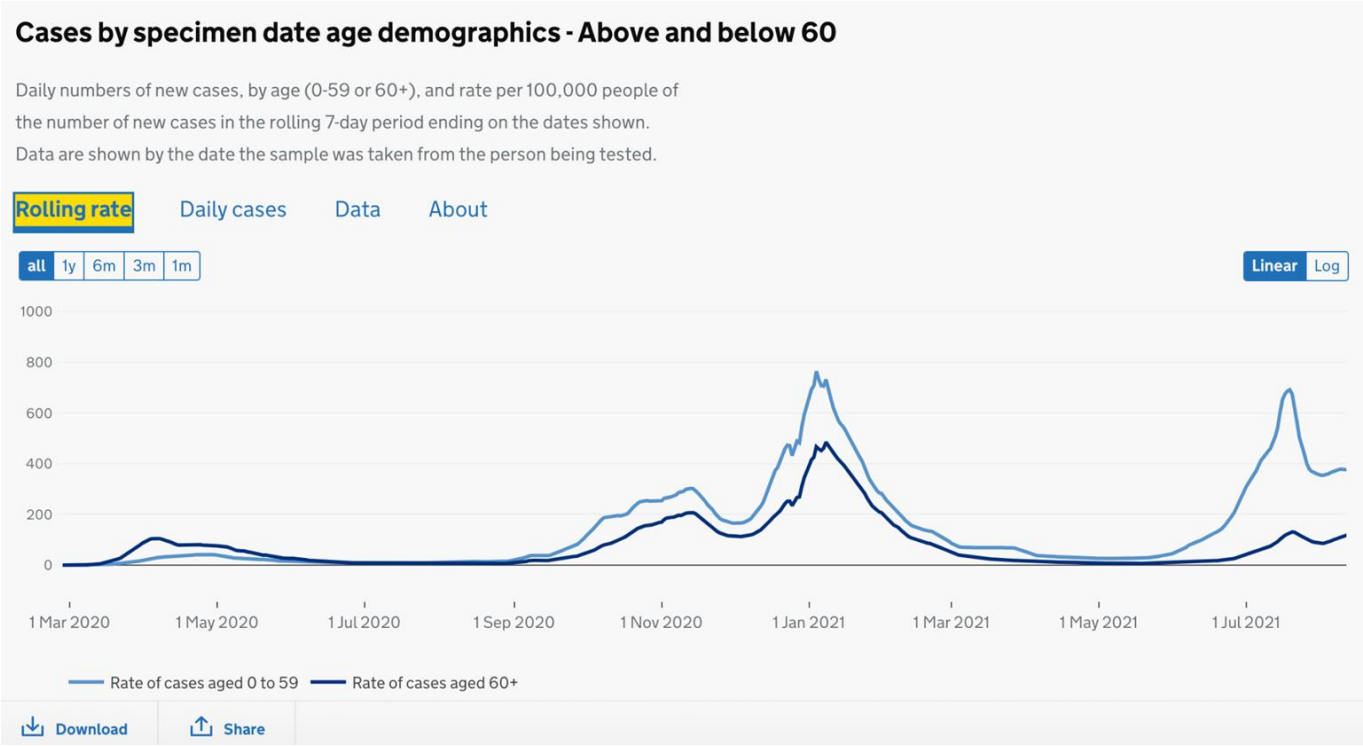

Covid-19 dashboard showing daily number of new cases by age.

Burn-Murdoch suggested that it is not solely about the high quality of the visualisations – but about how quickly they were produced and tailored to enable people to answer pressing questions. This was done well time and time again.

Though the daily press conferences received some criticism, **Burn-Murdoch** argued that presenting data to the nation on a regular basis was a positive. Not many other countries were trying to do that. And there are some examples of especially good practice from those briefings.

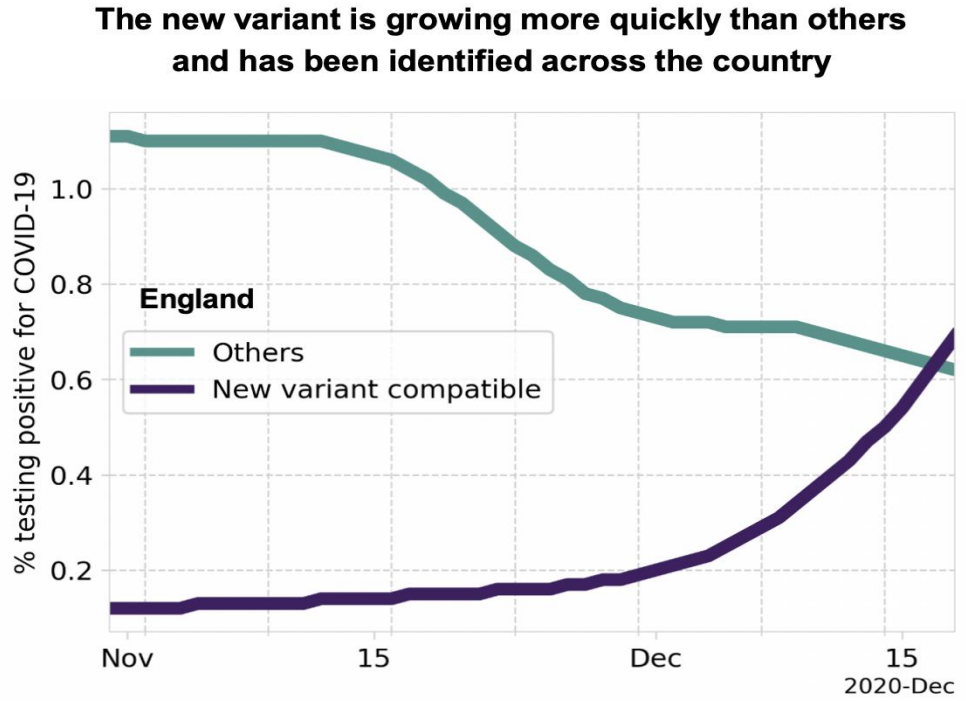

Chart from Government press conference on December 30 2020, showing growth of new variant.

The above example is taken from the period when, for the first time, there were two variants at play – the original strain and the Alpha variant. The chart clearly shows people that the Alpha variant is growing, and has a narrative title, so that even people who do not instinctively understand line charts can appreciate what is going on. It is a good example of how to communicate clearly with a chart that is on screen for only twenty seconds or so.

There were also less good examples. The below example shows heat maps for positive cases by age in different regions of England.

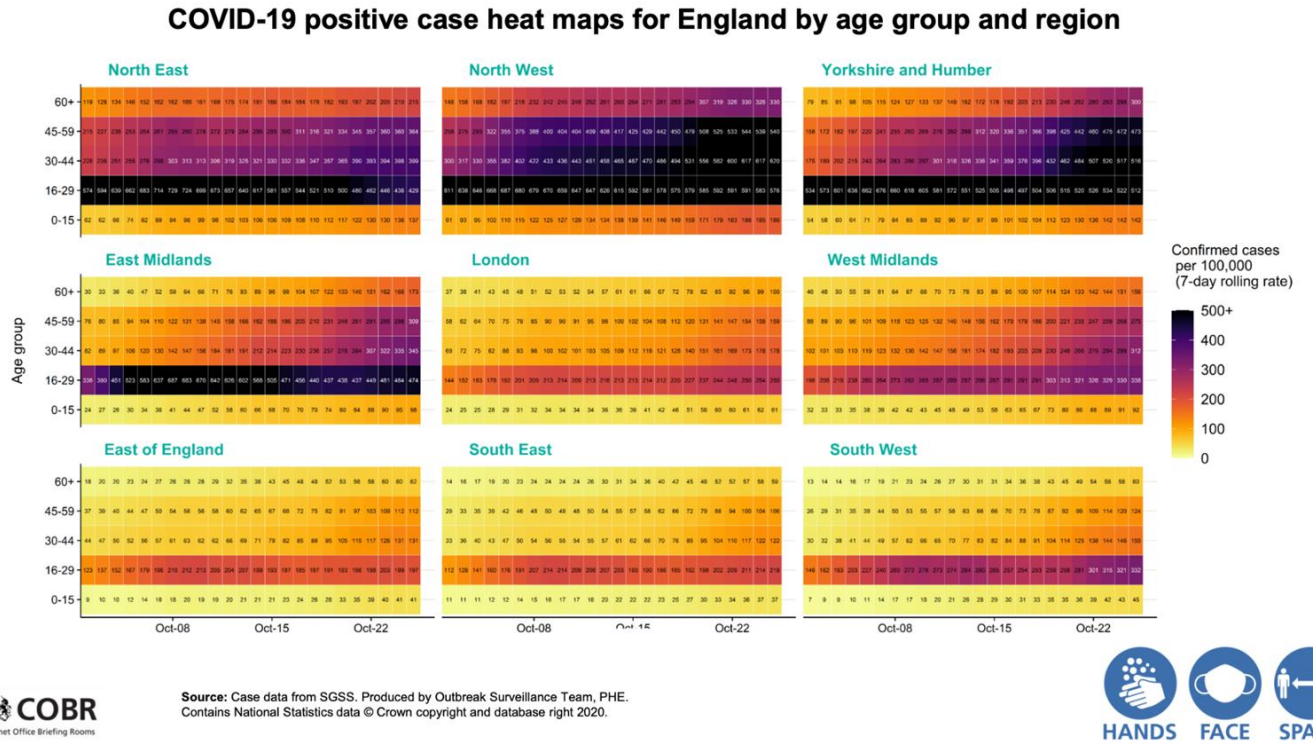

Heat maps of Covid cases by English region shown at Government press conference on October 31 2020.

In the context of the dashboard, heatmaps are great because people have plenty of time to study them. Whereas when nine heatmaps flash up on screen for fifteen seconds or so, there is not enough time for people to interpret them. Here, no narrative title is provided that explains what the key message we should take from the charts is. Even people who understand these types of chart well, will struggle to draw meaning from briefly looking at this type of visualisation. The context in which visualisations are presented is important. In the context of televised press conferences, it is important that the visualisations used have a clear, easily grasped, purpose.

Burn-Murdoch also looked at the standard of visualisations used by the media (and people who used social media platforms to try and shed light on the pandemic). He argued that there are a number of cases of very good practice here – of journalists and others with a public profile helping to explain what was happening to a wider audience.

As examples of good practice, **Burn-Murdoch** pointed to the below charts comparing international case numbers, using a logarithmic scale. There was a general realisation that people were taking the task of visualisation seriously and thinking about what they were trying to communicate and how that could be done creatively.

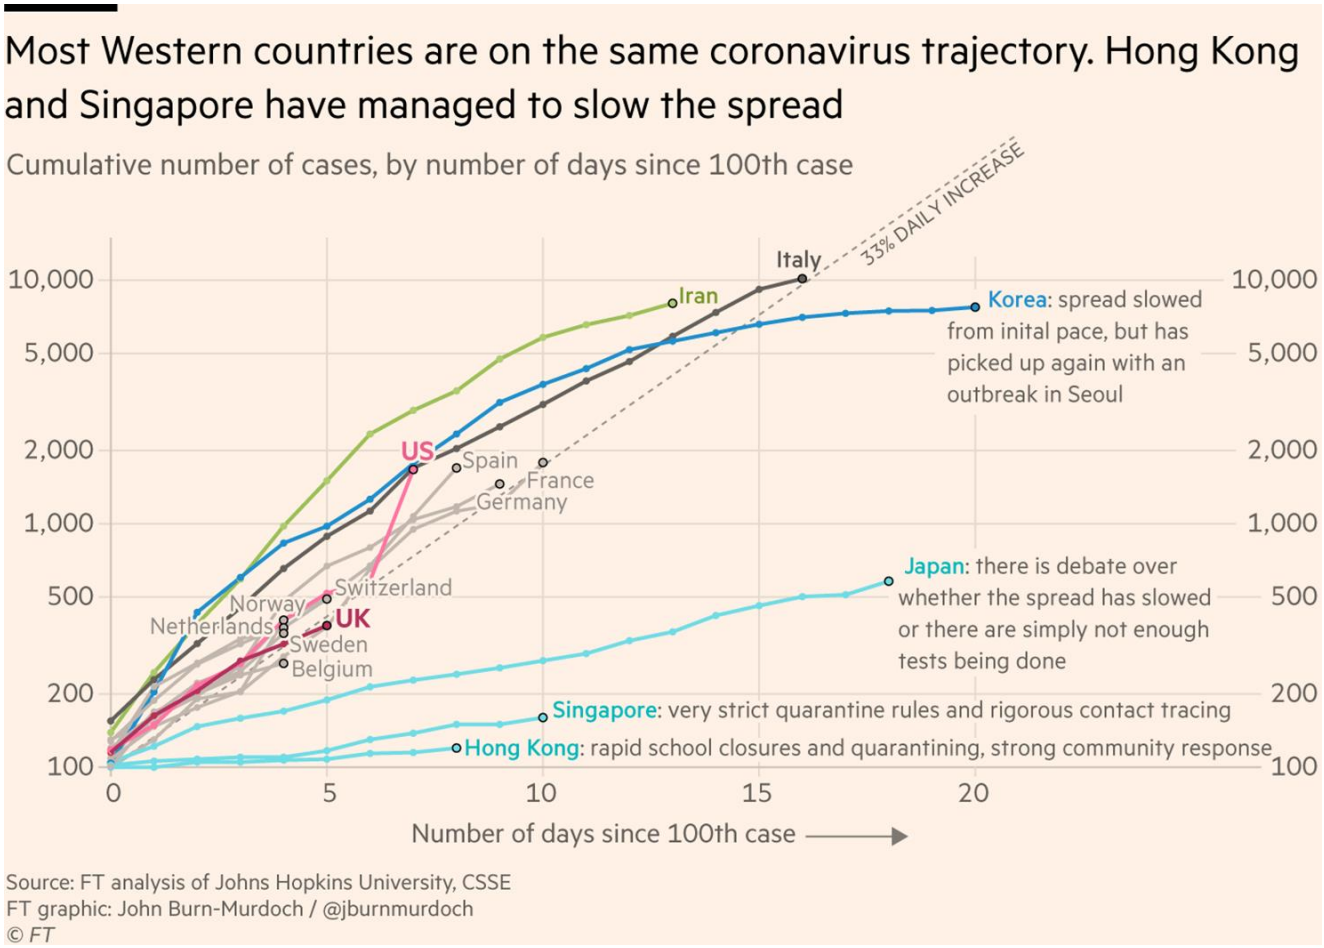

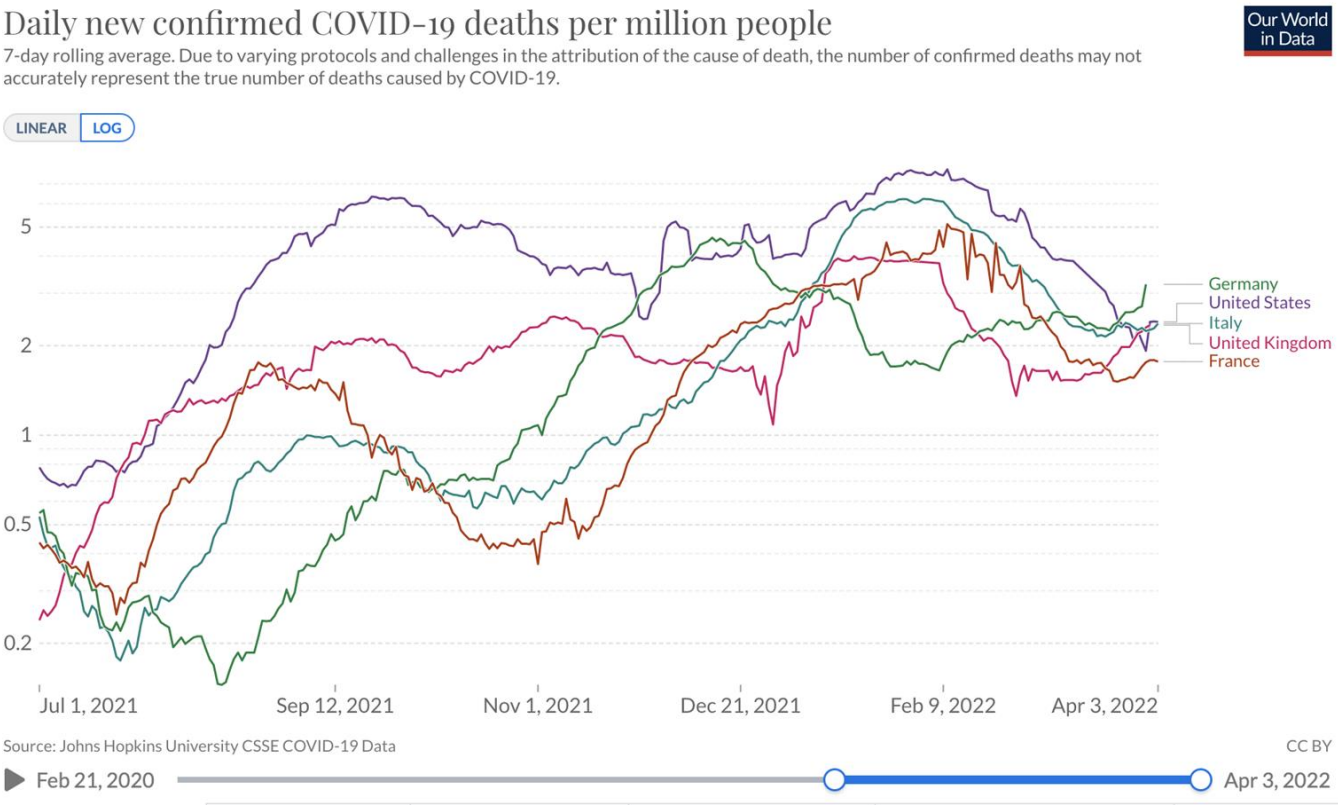

Examples of how a logarithmic scale can be used to helpfully compare international trends.

However, in some parts of the mainstream media, conservative tendencies could limit how useful visualisations were. **Burn-Murdoch** also provided examples of where a conservative tendency – ie, to say that there is a fixed way to presents charts, and to stick with that regardless of whether it is the most effective way to communicate information – meant that visualisations were not especially effective. The below example is taken from the BBC and shows how a linear chart may not give the same level of nuanced understanding as a logarithmic chart.

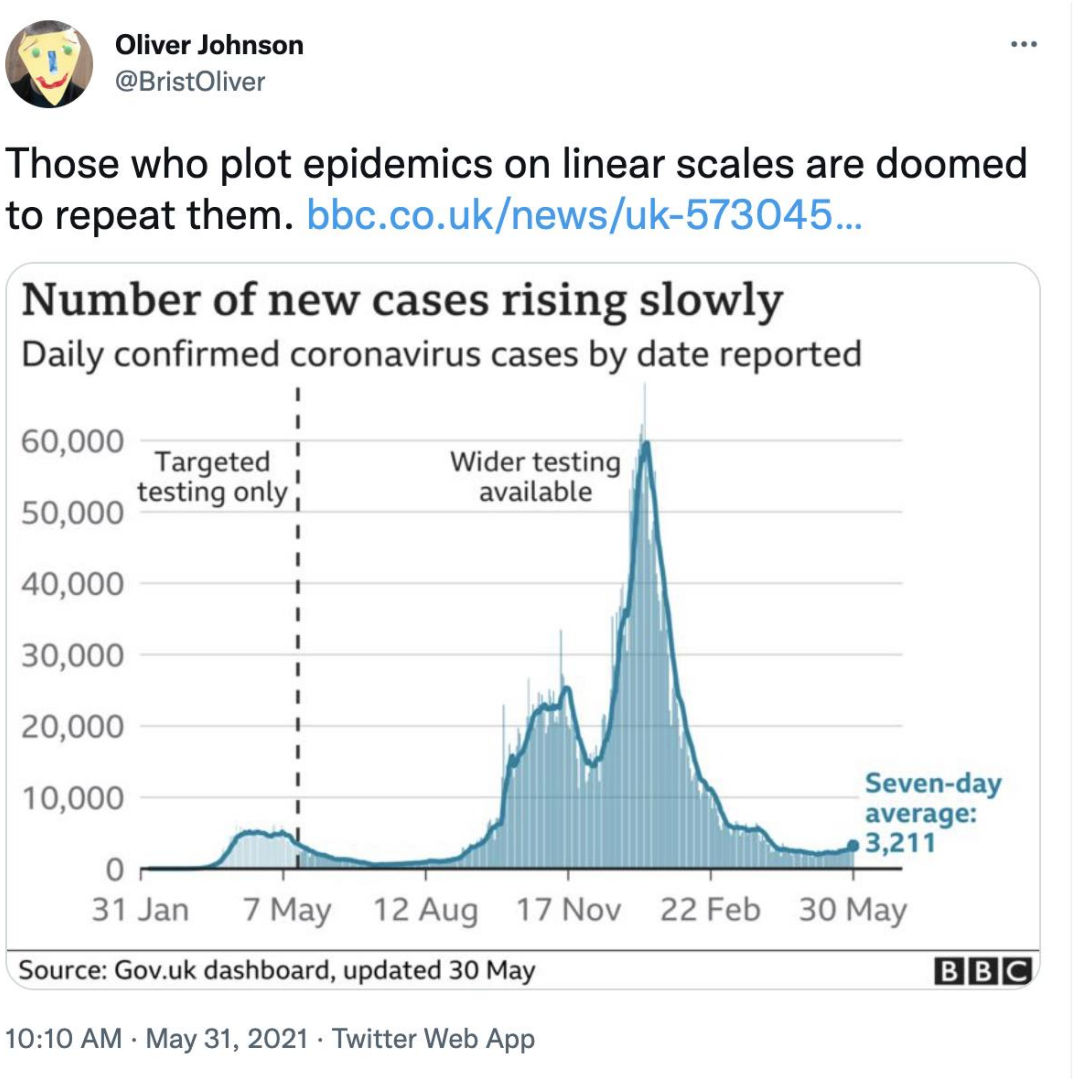

Example of a chart plotted on a linear scale.

A lot of people were confused by the use of logarithmic scales (shown by the increase in google searches of the term over the pandemic). But you don't help to educate people by avoiding presenting data in that way. Rather, the approach should be to acknowledge that it is not intuitive for everyone, and to spend some time explaining the visualisation.

Liza Hadley spoke about the government's daily press conferences. She highlighted the September 2020 press conference, in which a chart showing what would happen if cases doubled every seven days was presented. If a single scenario is given – however it is caveated – it will be misunderstood as a prediction. To avoid falling into this trap, one option is to give multiple scenarios. More recently, the government charts used in press conferences have sought to convey uncertainty using grey bands. But unfortunately, the use of descriptors is inconsistent - many plots do label 95% confidence intervals, but potentially confusing labelling such as "margin of error" or "interval" were sometimes used without being clear what was plotted. A consistent and clear approach to presenting uncertainty would have helped. It would also be helpful, when presenting forward projections, to remove the central

estimates and present only the range of possibilities – this forces people to engage with the fact that uncertainty is present rather than placing too much emphasis on a central projection. Plots displaying Covid-19 data were much stronger – incomplete data was shown in grey and clearly labelled. This was a strength. It helped communicate the uncertainty inherent in all rapid data collection, and this is now well understood. **Hadley** also recommended Nick Holliman's work on visual entropy as a way of communicating uncertainty.

5. How has data being accessible to the public improved communication – and was data accessible enough for this purpose?

Hannah Ritchie pointed to an article by Edouard Mathieu, which argues that data accessibility is more important than fancy dashboards. Dashboards, he argues, can be nice – and the UK dashboard is actually one of the best in the world, in that it combines making data available alongside a helpful dashboard. But many countries have been bad at striking this balance, and have prioritised dashboards over accessible data. If data is available, there is a large community who will dig into it and develop charts that can be used to help explain what is happening – that cannot be done without data.

Ritchie pointed to a number of metrics that could have been useful in the UK, that we either didn't have or that we did have but which weren't readily accessible:

- hospitalisations *for* or *with* Covid
- Covid deaths *because of* Covid or *with* Covid
- earlier breakdowns by age: hospitalisations, vaccinations, deaths
- deaths by vaccination status.

There are some countries that made this data very open, enabling helpful visualisations – examples are below. Switzerland provided this data, and this was helpful to understand how the pandemic was playing out there. Our World in Data attempted to create similar charts for England, but the file was insufficiently stable to easily produce this type of regular visualisation. It is important for the government to make it easy for people to access a reliable source of data.

Switzerland: COVID-19 weekly death rate by vaccination status, All ages

Death rates are calculated as the number of deaths in each group, divided by the total number of people in this group. This is given per 100,000 people.

Our World in Data

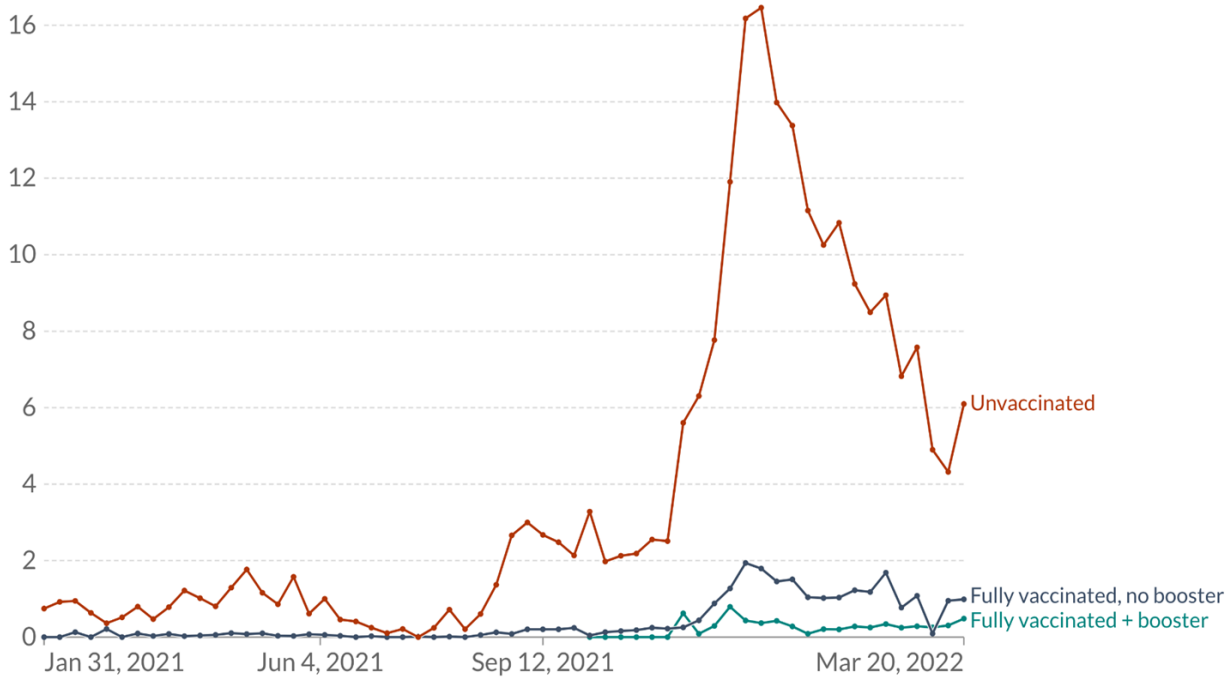

Source: Federal Office of Public Health
OurWorldInData.org/coronavirus • CC BY
Note: Data coverage includes both Switzerland and Liechtenstein. Unvaccinated people have not received any dose. Partially-vaccinated people are excluded. Fully-vaccinated people have received all doses prescribed by the initial vaccination protocol. The mortality rate for the 'All ages' group is age-standardized to account for the different vaccination rates of older and younger people.

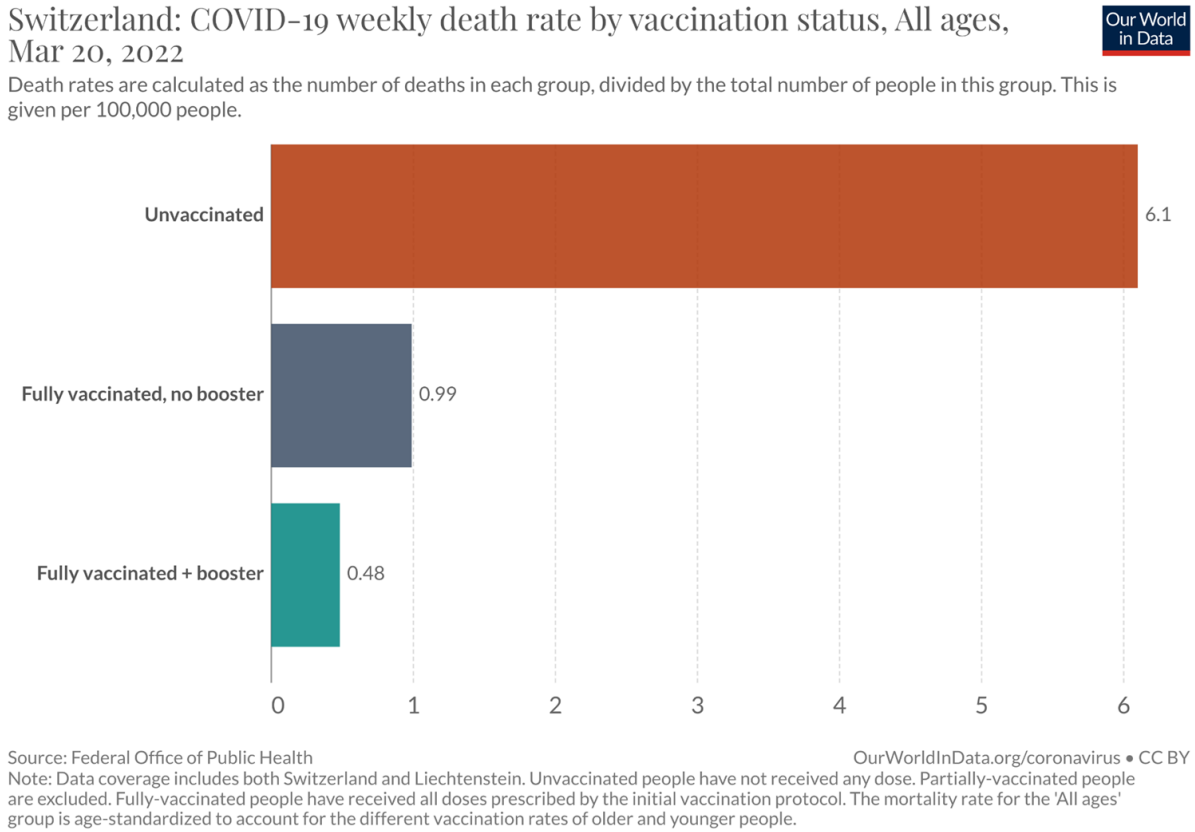

Charts from Our World in Data showing the situation in Switzerland, as examples of the type of chart that it was not possible to produce in the UK due to lack of relevant data.

While the UK was poor in making data accessible at the start of the pandemic, over time the UK has become one of the best in the world. No other country has produced a really good dashboard with accessible data. **Burn-Murdoch** identified Ecuador as another country that started poorly in terms of providing accessible data, but quickly started producing spreadsheets of information and publishing emergency bulletins of excess deaths, which helped understand the situation there.

Ritchie argued that it is important to provide the data that people need to answer key questions, and to communicate it effectively. If this is not done, people will fill the vacuum with misinformation.

In **Oliver Johnson's** view, statistics and data were generally handled well by "the grown ups in the room" – senior government officials, the media, scientists, etc. Scientists went over and above to communicate; press conference slides were clear; most of the media acted responsibly; the public had unprecedented near-live data access; and the dashboard (especially API), Our World In Data, ONS, and UKHSA variant reports etc have been great.

Johnson pointed to semi-official communications from individuals at certain agencies, which have been excellent – numerate individuals using social media (usually Twitter) in order to try and highlight aspects of the pandemic, while not using official organisational handles. On Twitter there is a community of – for want of a better term – 'nerds', says **Johnson**, who have worked in an extraordinary way using open tools and sometimes obscure online sources to give a speedy and visual commentary on the pandemic. This has been a great success and has helped lots of people to better understand the pandemic.

Johnson pointed to what he thought was the most notorious example of data visualisation from the pandemic – a slide presented at the daily government press conference on 21 September 2020 (below), which **Hadley** also referenced (§4).

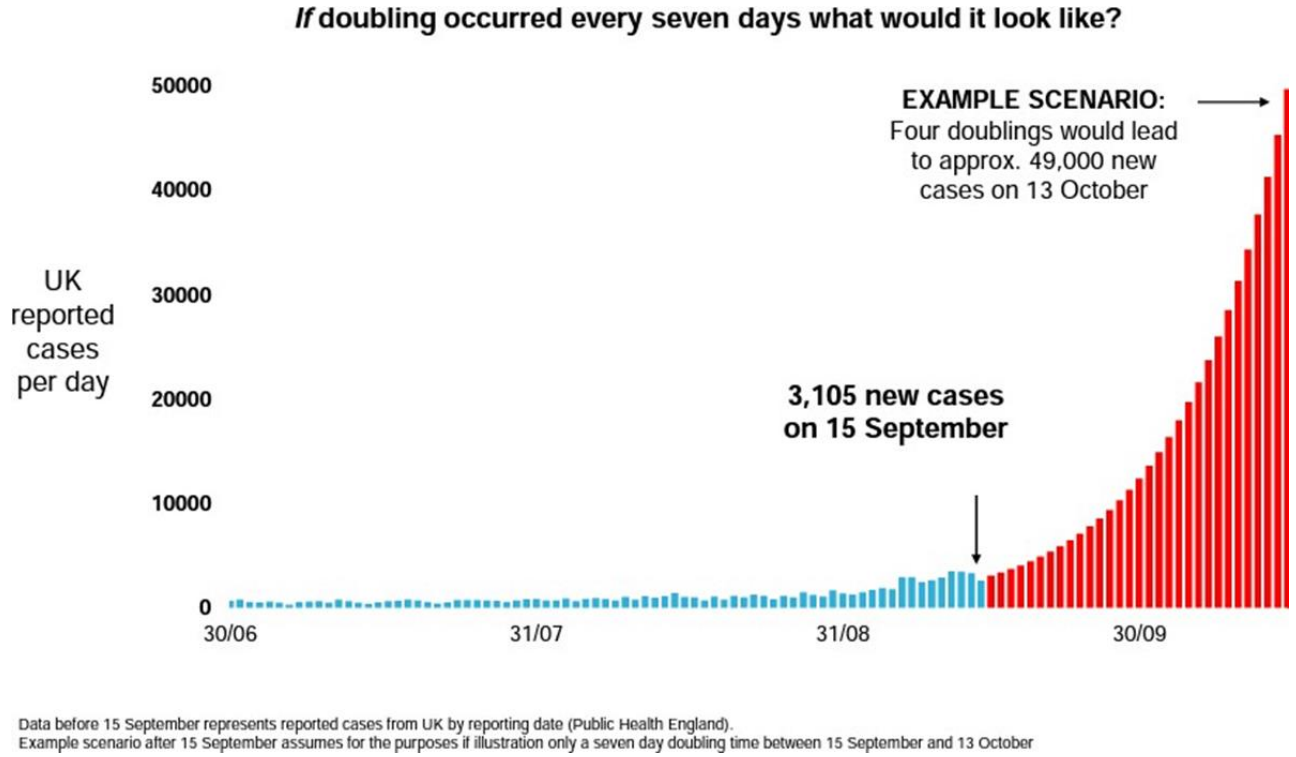

Chart presented at government press conference, 21 September 20202, showing UK reported cases per day and a projected scenario.

The Chief Scientific Officer presented this chart – including the projection based on the scenario that cases double every seven days, leading to 49,000 new cases by 13 October. While this chart says it is not really a projection – it is just showing what would happen if cases doubled every week – it remains a quite unconvincing visualisation. The part of the chart in red does not look likely to follow from the actual data that is available. It is easy to see why people might be sceptical of this type of claim.

Indeed, people have since compared this projection to what actually happened – as shown below.

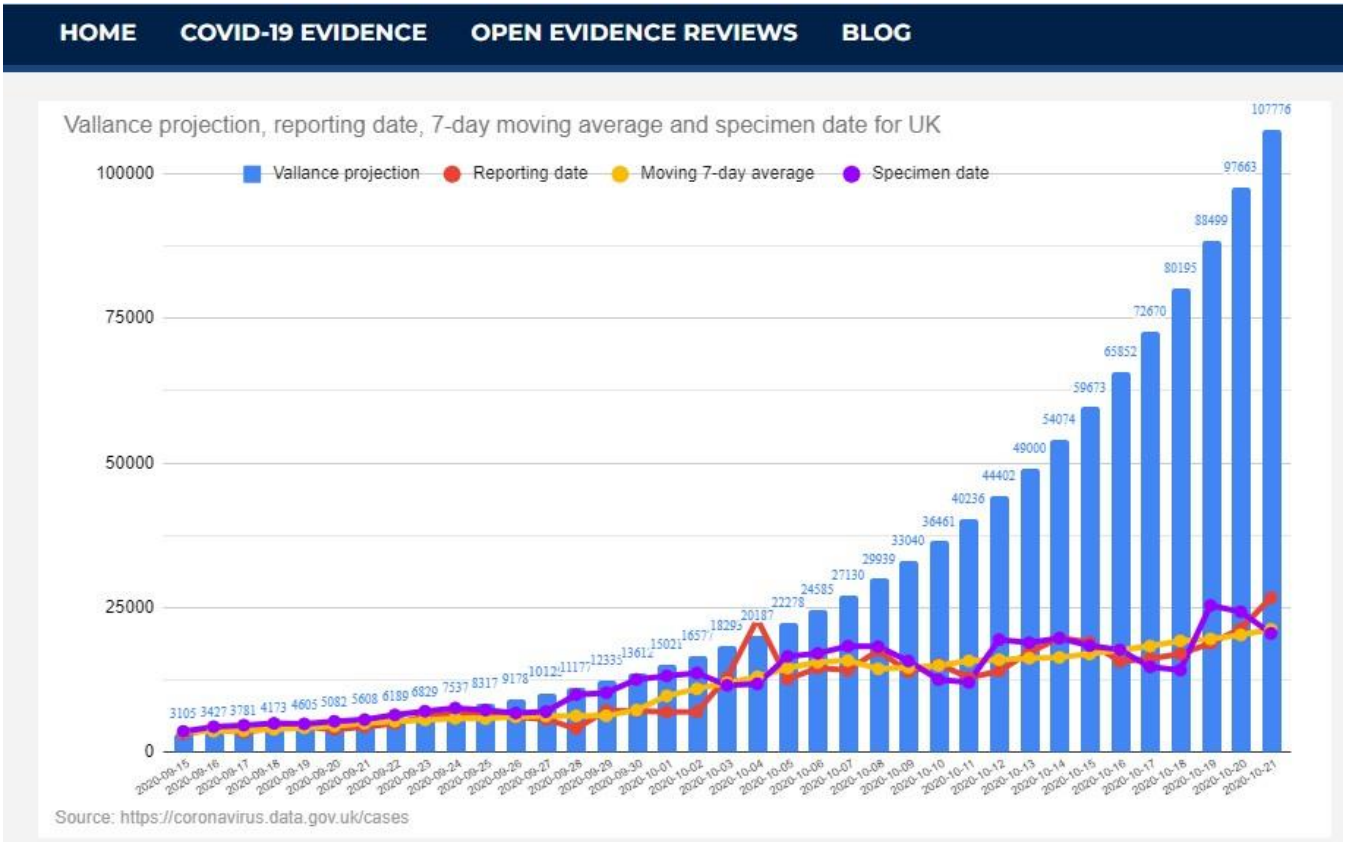

From www.cebm.net/covid-19/tracking-uk-covid-19-cases

This chart shows that while the virus did not spread at the projected rate, cases did continue to rise fairly quickly. This example highlights one of the problems of plotting exponential growth on a linear scale – when you get the growth rate wrong, you get it exponentially wrong. If, instead, you take the same data and use a logarithmic scale, you get a quite different picture – **Johnson** produced the below visualisation to show this. The red circles in the chart are the actual data that were presented at the press conference. **Johnson** fitted a regression line that came out with a doubling rate of 12.4 days, which is slower than the projection presented – but it is more convincing and based on available data. Then, when you add the real data (black circles below) it broadly follows the trend suggested by the log scale. This approach would have been much more convincing.

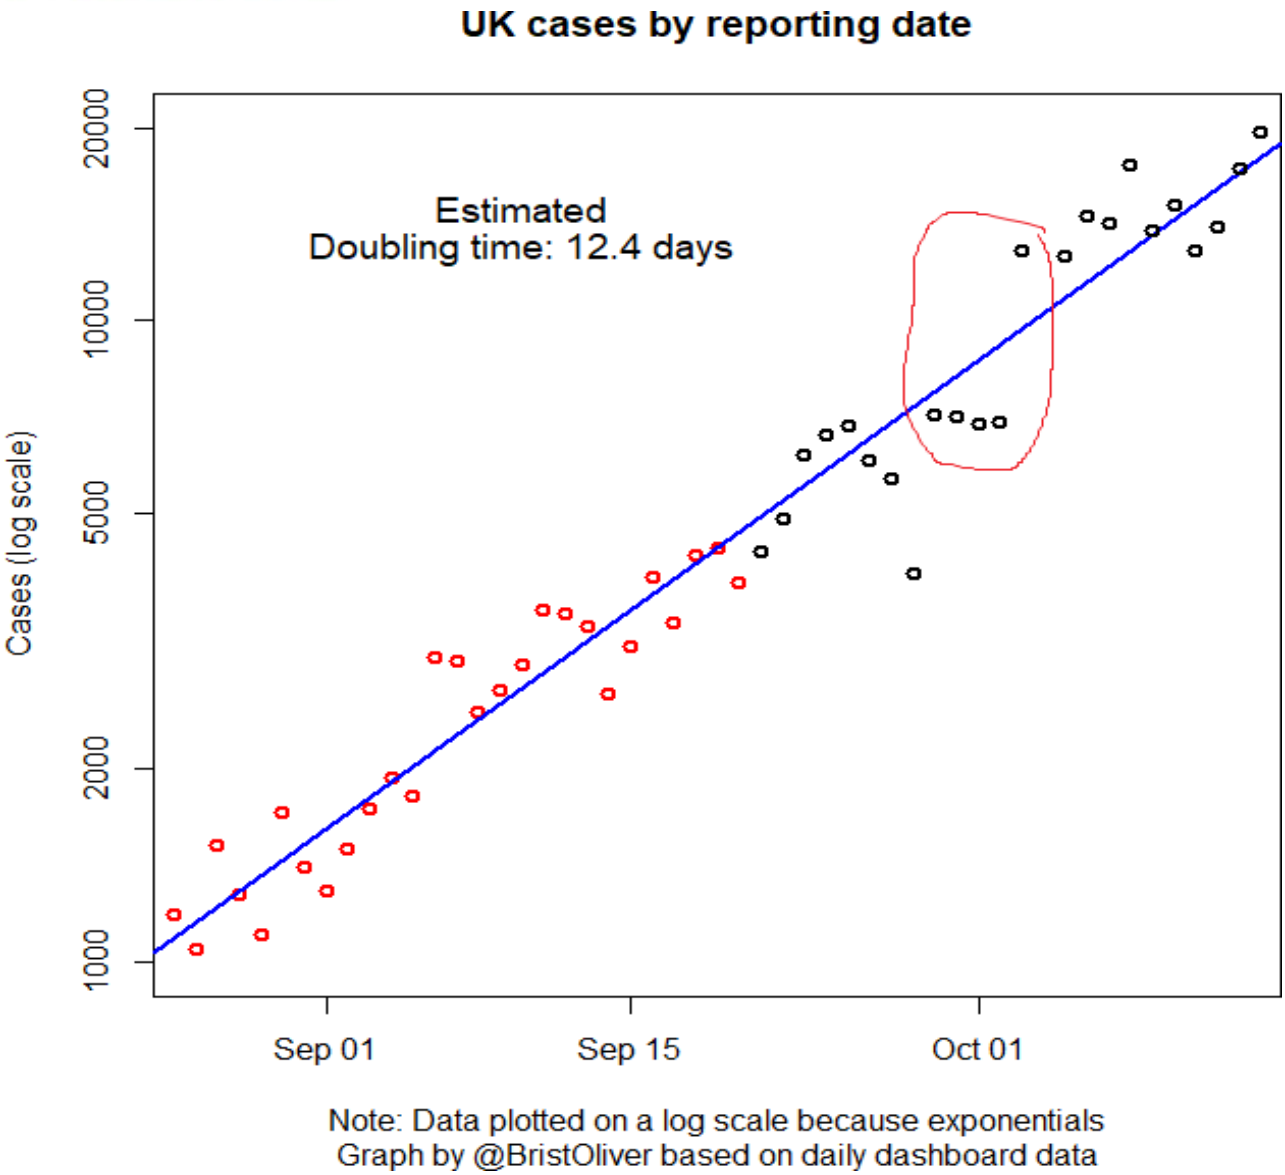

Visualisation produced by Oliver Johnson. Red circles represent the actual data presented by the government on September 21; the blue line extrapolates from that data based on a doubling rate of 12.4 days; black circles represent the actual data after September 21.

By June 2021 organisations seemed to have learnt this lesson – the BBC presented the chart below on 12 June using a more helpful, logarithmic scale.

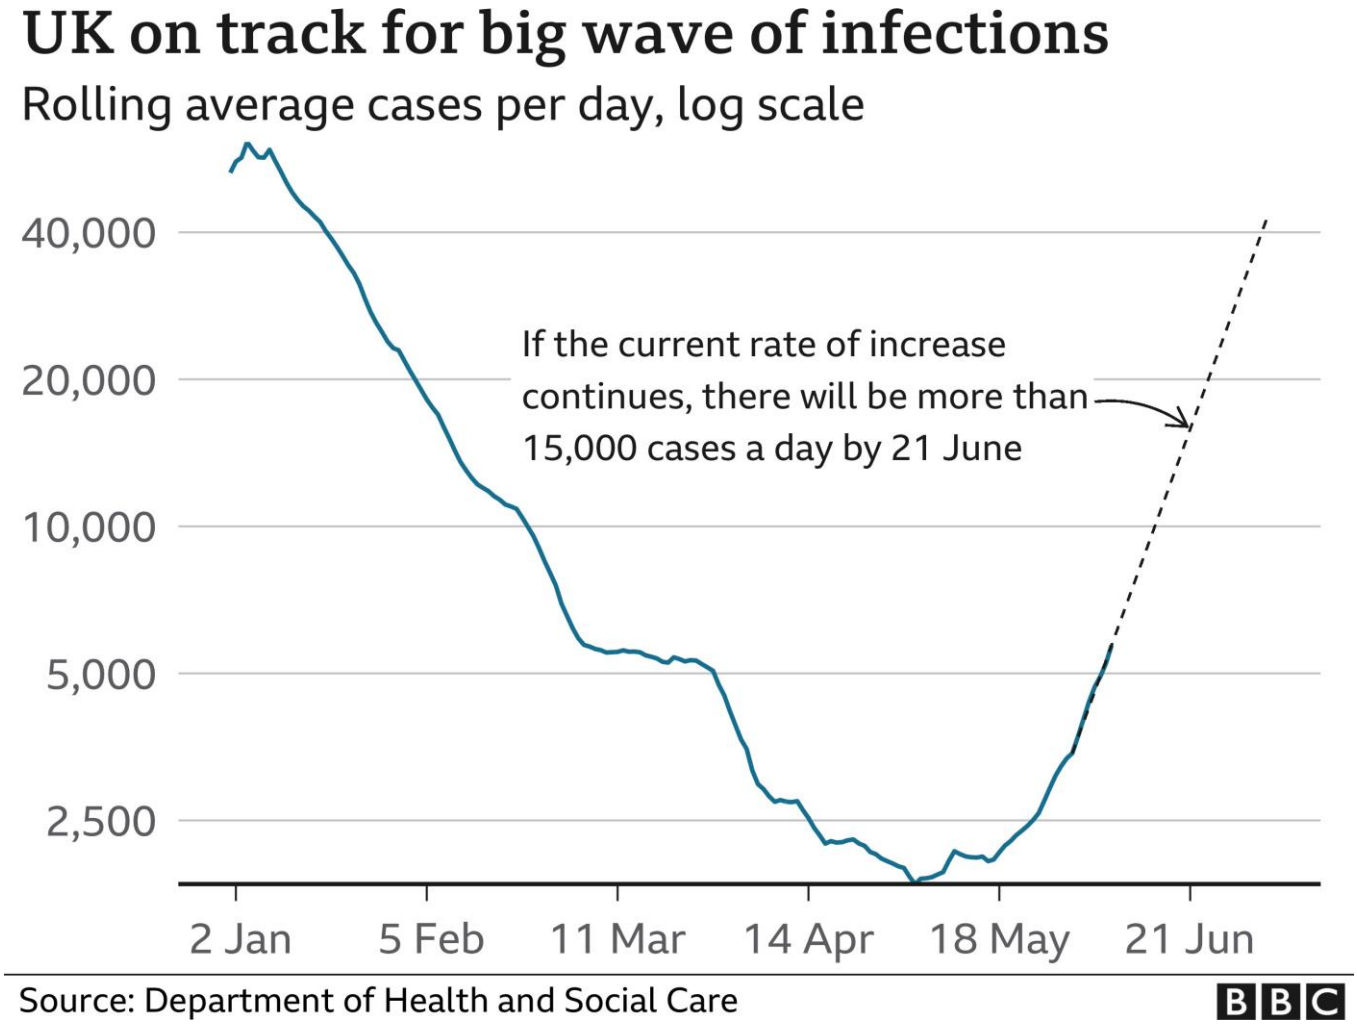

BBC visualisation from June 12 2021 showing projection using logarithmic scale.

Often the choice of scale does not matter that much. But when you are at an early stage and you are dealing with exponential growth, it can be helpful in providing an early warning in a way that looks more plausible. It is true that logarithmic scales may not always be intuitive to the public – but they are not necessarily more difficult to explain than the biological terms that also needed to be explained to the public when communicating about the pandemic.

Johnson also looked at the debate around lateral flow tests, arguing that this had not been the best example of statistical communication. He suggested that the most intuitive metric to use for talking about testing was positive predictive value – ie the likelihood that a positive test result is a true result. Consider the claim from the Guardian on 15 April 2021 that "as few as 2% to 10% of positive results may be accurate". Is this plausible? This can be assessed using Bayes' Theorem. If you translate this to odds – in order to use the odds formulation of Bayes' – you get 0.02 to 0.11. The odds formulation of Bayes' is:

$$\frac{P(\text{infected}\,|+ve)}{P(\text{not infected}\,|+ve)} = \frac{P(\text{infected})}{P(\text{not infected})}\,\frac{P(+ve|\,\text{infected})}{P(+ve|\,\text{not infected})}$$

This says that the odds of being infected given a positive test is the product of three things – it is the prevalence multiplied by the true positive rate divided by false positive rate. So, if you are saying that the odds of being infected given a positive test are as low as 0.02 then you need to have, as an example, something like figures of 0.02% for prevalence, a true positive rate of 50% and a false positive rate of 0.3% – which gives odds of 0.03. This

formulation helps us identify the claims that are inherent in a statement like the Guardian's and think about whether they are plausible. And this is where the unofficial response on social media could help.

Johnson had realised in April 2021 that the Scottish government were publishing a report every week with the number of lateral flow tests and positive lateral flow tests being reported. He presented that information as below.

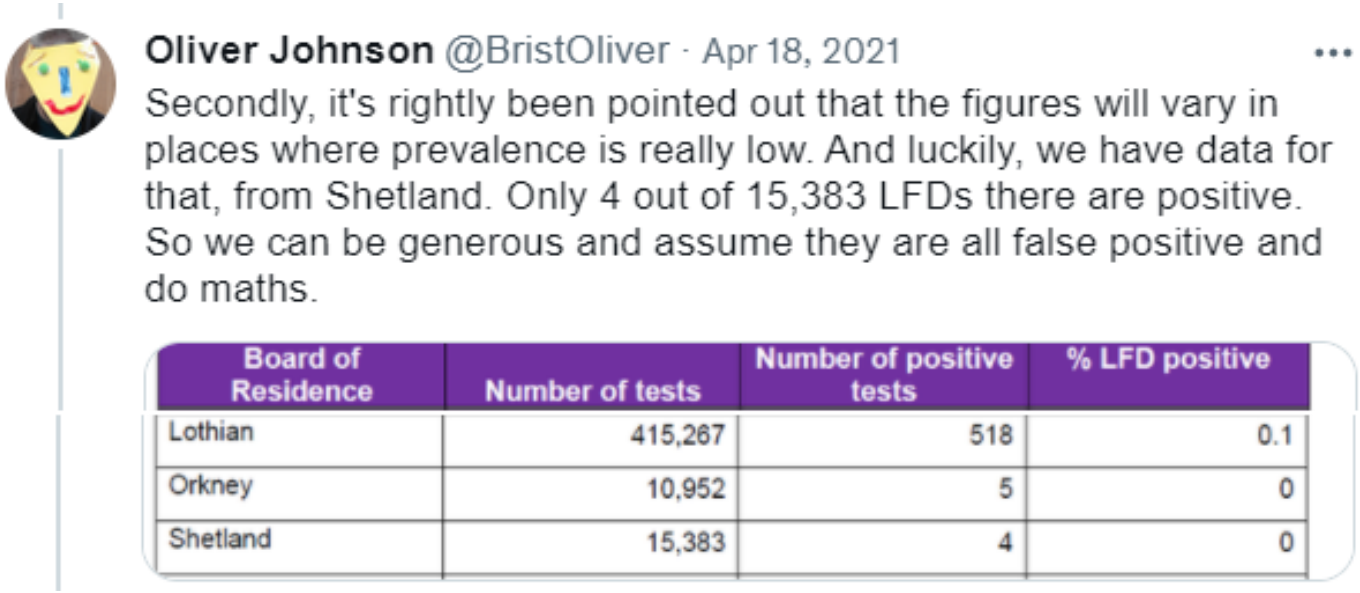

Oliver Johnson @BristOliver · Apr 18, 2021

Secondly, it's rightly been pointed out that the figures will vary in places where prevalence is really low. And luckily, we have data for that, from Shetland. Only 4 out of 15,383 LFDs there are positive. So we can be generous and assume they are all false positive and do maths.

Board of Residence	Number of tests	Number of positive tests	% LFD positive
Lothian	415,267	518	0.1
Orkney	10,952	5	0
Shetland	15,383	4	0

Fig x. Screenshot of Johnson's Twitter thread showing upper range of false positives for lateral flow tests.

If we assume that all tests in Orkney and Shetland were false positives, we can get an upper bound for the false positive rate – ie, (4+5)/(15,383+10,952), which comes to around 0.03%. This is an order of magnitude different from the 0.3% needed for the claim in the Guardian to be right. Around the same time, there was other confirmation of this – on 9 April 2021 there was an update to the dashboard that cancelled 8,010 reported positive lateral flow tests, on the grounds that those cases tested negative with PCR tests within three days of the positive lateral flow result. This allowed people to work out the proportion of positive tests that were correct – which was around 82%. That was consistent with what you could deduce at the time using Bayes' Theorem.

This is an example of how it is possible to use metrics that weren't specifically put out for that purpose, to accurately estimate the positive predictive value of a lateral flow test.

Using similar techniques, it was also possible for people to look at statistics and identify problems before they were officially recognised. From September-October 2021, an Immensa lab in Wolverhampton was sent around 43,000 false negative results. Below is a screenshot from Twitter on 10 September identifying the problem – looking at daily ratios of positive cases.

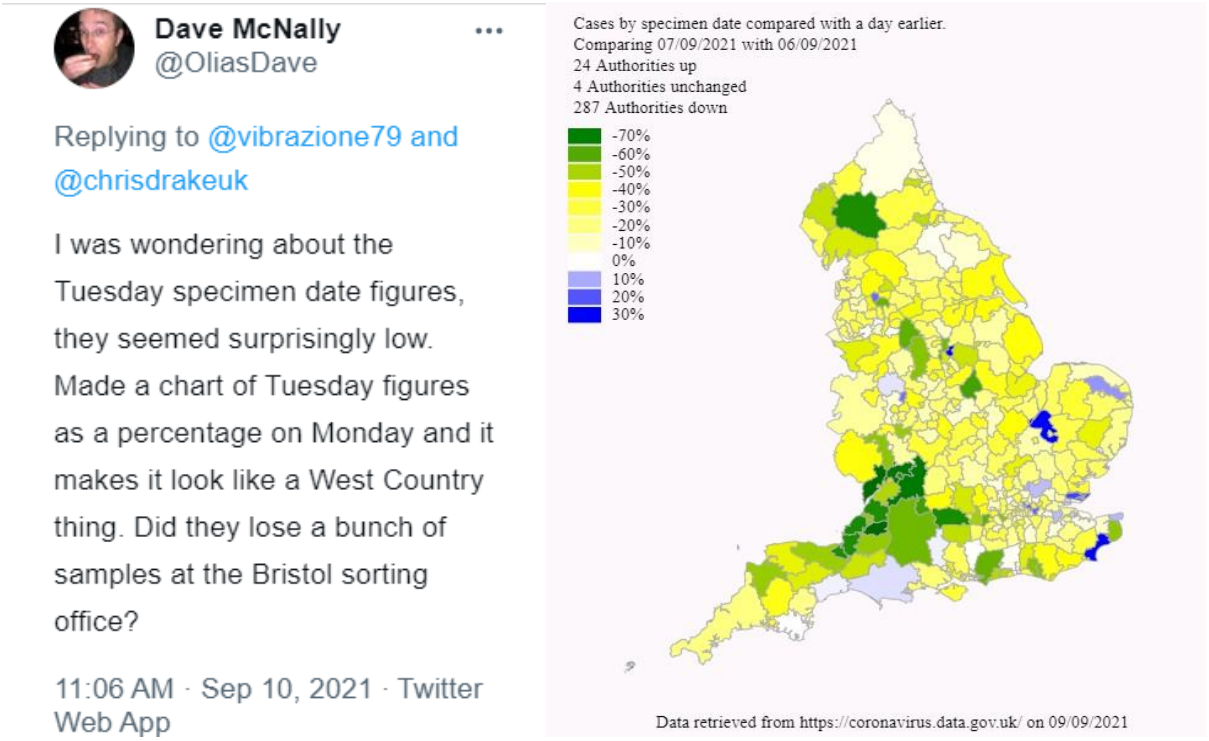

Immensa issue identified on Twitter using publicly available data.

In case citizens spot something like that, it would be helpful if there was a mechanism for them to feed that back to officials so that it can be investigated at an earlier stage.

6. Was enough done to try and prevent information becoming misinformation?

Burn-Murdoch argued that when a notorious misinformer, such as Joe Rogan, shares official government documents as evidence of something that you know is not true, that is a good indicator that not enough was done to combat misinformation. The latest intended solution involved presenting information that is known to be potentially misleading, and then explaining the information in footnotes – an example of this is shown below. This is not an effective approach – the vast majority of people do not read footnotes, and people who would present the data to misinform others are not going to present the government's explanation of the data. It would be more effective, **Burn-Murdoch** suggested, to remove the information that is being misused.

	Cases reported by specimen date between week 9 2022 (w/e 6 March 2022) and week 12 2022 (w/e 27 March 2022)		Cases presenting to emergency care (within 28 days of a positive test) resulting in overnight inpatient admission, by specimen date between week 9 2022 (w/e 6 March 2022) and week 12 2022 (w/e 27 March 2022)		Death within 28 days of positive COVID-19 test by date of death between week 9 2022 (w/e 6 March 2022) and week 12 2022 (w/e 27 March 2022)		Death within 60 days of positive COVID-19 test by date of death between week 9 2022 (w/e 6 March 2022) and week 12 2022 (w/e 27 March 2022)	
	[see information on population bases and unadjusted rates in footnotes 1 and 2 below this table]							
	Unadjusted rates among persons vaccinated with at least 3 doses (per 100,000)	Unadjusted rates among persons not vaccinated (per 100,000)[1,2]	Unadjusted rates among persons vaccinated with at least 3 doses (per 100,000)	Unadjusted rates among persons not vaccinated (per 100,000)[2]	Unadjusted rates among persons vaccinated with at least 3 doses (per 100,000)	Unadjusted rates among persons not vaccinated (per 100,000)[2]	Unadjusted rates among persons vaccinated with at least 3 doses (per 100,000)	Unadjusted rates among persons not vaccinated (per 100,000)[2]
Under 18	1,454.0	1,711.7	3.1	9.6	0.0	0.0	0.0	0.0
18 to 29	3,118.8	941.6	5.4	8.2	0.1	0.0	0.1	0.2
30 to 39	4,324.7	1,085.6	6.8	7.4	0.2	0.3	0.4	0.5
40 to 49	3,957.8	955.3	6.0	7.7	0.2	0.3	0.6	0.7
50 to 59	3,303.4	779.8	9.0	12.9	0.5	1.6	1.2	2.4
60 to 69	2,814.9	572.8	14.3	22.1	1.5	5.9	3.8	9.1
70 to 79	2,161.5	532.1	36.6	58.8	6.8	20.2	13.4	30.3
80 or over	2,023.7	775.6	117.9	123.5	44.6	87.4	84.4	121.8

[1] Comparing case rates among vaccinated and unvaccinated populations should not be used to estimate vaccine effectiveness against COVID-19 infection. Vaccine effectiveness has been formally estimated from a number of different sources and is summarised on pages 4 to 14 in this report.

The rates are calculated per 100,000 in people who have received either 3 doses of a COVID-19 vaccine or in people who have not received a COVID-19 vaccine. These figures are updated each week as the number of unvaccinated individuals and individuals vaccinated with 3 doses in the population changes.

The case rates in the vaccinated and unvaccinated populations are unadjusted crude rates that do not take into account underlying statistical biases in the data and there are likely to be systematic differences between these 2 population groups. For example:

- testing behaviour is likely to be different between people with different vaccination status, resulting in differences in the chances of being identified as a case
- many of those who were at the head of the queue for vaccination are those at higher risk from COVID-19 due to their age, their occupation, their family circumstances or because of underlying health issues
- people who are fully vaccinated and people who are unvaccinated may behave differently, particularly with regard to social interactions and therefore may have differing levels of exposure to COVID-19
- people who have never been vaccinated are more likely to have caught COVID-19 in the weeks or months before the period of the cases covered in the report. This gives them some natural immunity to the virus which may have contributed to a lower case rate in the past few weeks

Taken from UKHSA variant surveillance report. The vaccinated and unvaccinated rates in the first two columns have been used to question the effectiveness of the vaccine. The footnote provides an explanation intended to counter this.

Burn-Murdoch also presented an example of his work at the FT, which raised some interesting ethical questions. The below chart shows how as vaccinations have been rolled out and immunity has increased, Covid-19 became gradually less lethal. This was presented with the headline “Vaccines and Omicron mean Covid now less deadly than flu in England”. There is a question as to whether charts and headlines like this give ammunition to people who want to say that Covid is no longer an issue, everything is now fine, and we can forget about it. This is a grey area, and there is a question as to whether it is right to publish things if there is a reason to believe they could be misused by misinformation campaigns.

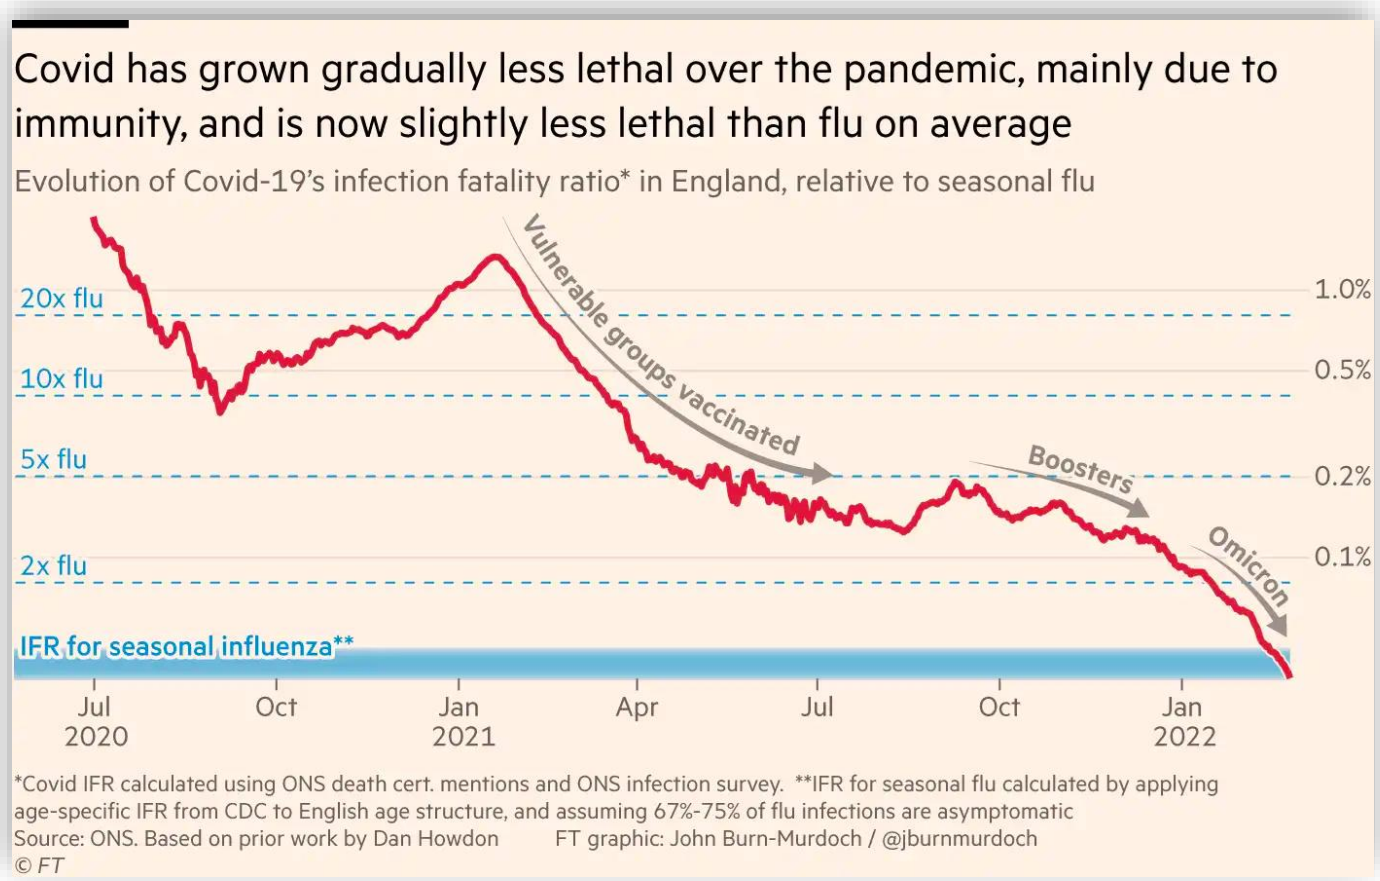

Will Moy detailed the steps that Full Fact took to counter misinformation. They wrote over 800 fact checks related to the pandemic – vaccines, deaths and testing were the most common topics. They are first responders to information quality problems and around twenty million people visited their website – giving some indication of the appetite for this type of information. This was helped by Google embedding their fact checks into its search results.

Full Fact asked for hundreds of corrections over the pandemic. The media were generally pretty good at following these requests, whereas politicians were more reluctant to issue corrections.

Moy also suggested some ways that we could more systematically intervene to prevent the spread of misinformation. Full Fact have been working with numerous public bodies to notice early examples of their information being misused, to prevent it as quickly as possible, and to pre-empt it. He gave the example of the yellow card system, which allows people to report adverse experiences after receiving a vaccine (a very important safety net). That database is available online. People misunderstood something happening after a vaccine as something happening *because* of the vaccine. We have seen this misunderstanding turned into adverts in local newspapers – misrepresenting the evidence as showing that vaccines could kill you. The regulator has forced these adverts to be removed. We have seen it quoted by MPs in parliament. This is a powerful example of how a simple collection exercise that was never intended to be controversial suddenly turned into ammunition for people on different sides of an argument. The MHRA have gradually updated the messaging around the yellow card to stress what it does and doesn't do. But in general, our public bodies have not been paying very much attention to how what they are communicating can be misused and the urgency of responding to that quickly. It is Full Fact's experience that when those misuses become established, it is much harder to stamp out.

The second type of systematic intervention that Full Fact carried out involved flagging to counter-disinformation teams in government that, when the vaccine was released, there would be a surge of misinformation. When the

vaccine launched, these government teams were therefore well-prepared to counter misinformation, and that has contributed to the high vaccine uptake.

There were examples of good practice around pre-empting misinformation. For example, it was known that reported deaths would drop at the weekends so these figures were always reported with very clear caveats. It is important, for **Moy**, to publish data even when there is a risk of it being used disingenuously – but producers need to be alert to the possibility that their data could be misunderstood, and provide very clear explanations to minimise the risk of this.

Matus Medo discussed the role that scientific papers and slow editorial processes can play in contributing to the spread of misinformation and how this might be improved. His direct experience with this issue was after *Science* published a paper on the impact of population-wide antigen testing in Slovakia. **Medo** and colleagues identified important methodological problems and modelling issues in this paper and produced a technical comment on this paper, aiming to have it published in *Science.* It was rejected because the journal had decided not to publish any technical comments on articles related to Covid to speed up publishing. You could provide technical comments on any subject apart from Covid. Another example he gives is from the *European Journal of Epidemiology*, which published a paper claiming that Covid incidence is unrelated to vaccination levels in the US. Several researchers became aware of issues, and eventually managed to publish some comments alongside the original article on the journal's website. However, it took almost three months for this to happen. By that time, the original article had been accessed two million times, indicating the huge impact the paper managed to make on public debate in the meantime. These two examples show that poor peer review in combination with traditionally slow editorial mechanisms have directly led to scientific papers published by prestigious journals becoming direct sources of misinformation for the general public. To address the problem, **Medo** suggests introducing a 'quarantine' for scientific papers.

Ting-Li Su emphasised the importance of conveying statistics around long Covid. She argues that the impact of long Covid should be regularly highlighted and reported while communicating the risk of Covid to the public. Messaging focuses on fatality, which does not give an accurate reflection of the impact of Covid on individuals or society. This is especially true for children, for whom the disease can be long-lasting and debilitating – the current impression from the government and media is that Covid is mild among children, and this should be counterbalanced. Statistics on long Covid are lacking and not very accurately measured – we don't know whether the severity of the disease impacts the chance of developing long Covid and we don't know whether vaccination status impacts the likelihood of being affected by long Covid. Especially as more people are getting reinfected with Covid, it would also be helpful to have a sense of whether long Covid is as likely on reinfection. **Su** suggests that long Covid should be added to the reporting dashboard.

Jonathan Smyth-Renshaw suggests two areas where more could have been done. He argues that reporting cases without reference to the number of tests undertaken is misleading, as sample size is a key piece of information in any analysis. There was also a problem with the use of models that were wildly wrong – models using misinformation were used to create more models of misinformation. In particular the models used to empty hospitals into care homes resulted in thousands of early deaths. **Smyth-Renshaw** also voiced thoughts on an area of concern for the future – how misinformation could have serious consequences. He suggests that the impact of prophylactic vaccines is still to be understood with respect to Covid. He expresses concern that misinformation informing past decisions will undermine information needed to manage strains of Covid in the future.