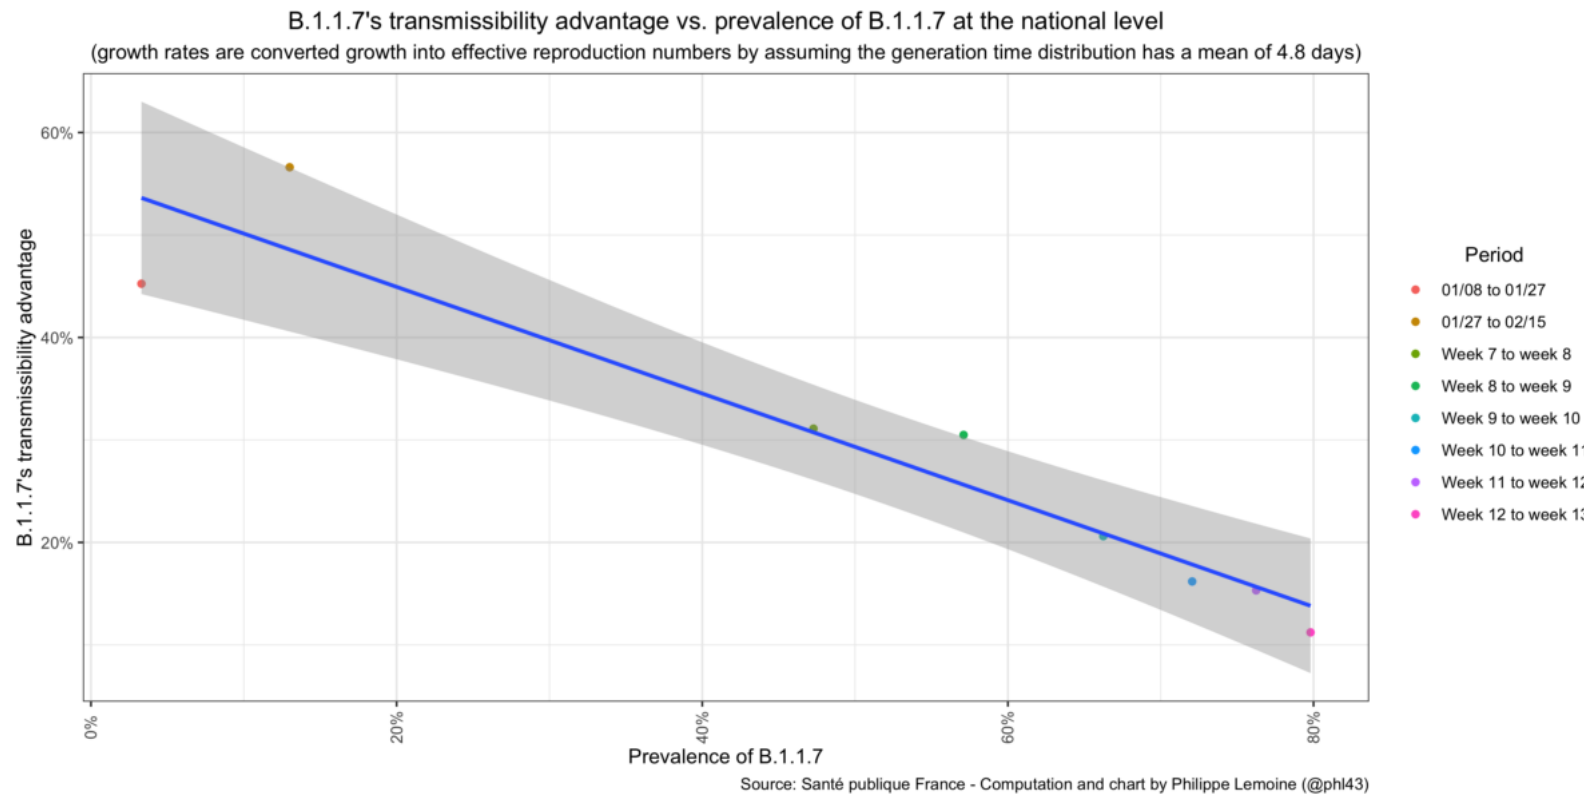


The British variant of SARS-CoV-2 and the poverty of epidemiology

April 9, 2021 by [Philippe Lemoine](#)



Summary

- B.1.1.7, a variant of SARS-CoV-2 that emerged in England last year and has been spreading rapidly everywhere since the beginning of the year, is claimed to be far more transmissible than the historical lineage. For instance, based on the early expansion of B.1.1.7 in France, Gaymard et al. (2021) found that it was between 50% and 70% more transmissible, which was used to recommend more stringent restrictions since if that were accurate a massive surge of incidence would be inevitable unless transmission were significantly reduced.
- However, this estimate is based on fitting a simple exponential growth model to 2 data points in January and is extremely sensitive to the assumptions made about the generation time distribution, about which I argue there is considerable uncertainty. When I replicate Gaymard et al.'s analysis but properly take into account this uncertainty by trying a wider range of possible assumptions about the generation time distribution, which I was able to do no thanks to them since not only did they not publish their code but they also didn't reply when I asked them for it, I find that B.1.1.7 could be anywhere between 21% and 72% more transmissible, a much wider range than what is reported by Gaymard et al.
- If I fit the same type of model but use more recent data instead of relying only on the growth rate of B.1.1.7 in January, I find that B.1.1.7's transmissibility advantage ranges from 16% to 42% depending on what assumptions I make about the generation time distribution, which is much lower than Gaymard et al.'s 50%-70% range. However, this conclusion is based on the assumption that B.1.1.7 has been growing exponentially in France since the beginning of the year, a view that is clearly falsified by the data even though some epidemiologists inexplicably continue to hold it.
- Most epidemiologists probably realize that, but argue that the explosion they predicted in January fail to materialize thanks to the decision to advance the curfew from 8pm to 6pm in January and the school holidays in February. So rather than assume a simple exponential growth model, I try to model the effect of government interventions and school holidays on transmission in a way similar to what the epidemiologists who advise the French government are doing, except that I use that model to estimate B.1.1.7's transmissibility advantage instead of assuming it's between 50% and 70% more transmissible. This approach leads to the conclusion that, depending on what assumptions we make about the generation time distribution, B.1.1.7's transmissibility is somewhere between 22% and 53% more transmissible.
- I explain that, even if B.1.1.7's transmissibility advantage had remained constant (which this approach implicitly assumes), this estimate would not be reliable because, as I argue elsewhere, this kind of model rests on strong and totally unrealistic mechanistic assumptions. So I also try an econometric approach to estimate B.1.1.7's transmissibility advantage, which also assumes that it has remained constant, but is agnostic on the underlying mechanism. However, the estimates are even more all over the place than with the previous approach, with B.1.1.7's transmissibility advantage ranging from -12% to 98% depending on what model I use and what assumptions I make about the generation time distribution.
- Thus, even if we assume, as epidemiologists systematically do (but rarely acknowledge explicitly), that B.1.1.7's transmissibility advantage has remained constant, Gaymard et al.'s claim that it's 50% to 70% more transmissible is completely unwarranted, since there is far more uncertainty than that. However, I show that B.1.1.7's transmissibility advantage has not remained constant in France since the beginning of the year, but has rapidly fallen as this lineage rose in prevalence and I estimate that it's now only 11% more transmissible than the historical lineage.
- Unfortunately, epidemiologists apparently couldn't be bothered to check that B.1.1.7's transmissibility advantage has remained constant, they just assumed it was and plugged Gaymard et al.'s 50% to 70% estimate into the models they use to make their projections, which unsurprisingly predicted that incidence would soon blow up like never before. Of course, this didn't happen, but instead of admitting they were wrong and revising their assumption that B.1.1.7's transmissibility advantage has remained constant, they just offered ad hoc rationalizations for why the projections didn't come true even though they were right and continued to advocate for stringent restrictions.

The number of COVID-19 cases has recently started to increase again in several countries on both sides of the Atlantic. If you listen to epidemiologists in the media, the culprit is B.1.1.7, a variant of SARS-CoV-2 that first appeared in the UK and which they claim is far more transmissible than the historical lineage. Since it has rapidly expanded everywhere it's been introduced, there is no doubt that B.1.1.7 is more transmissible or at least that it initially was. But how much more? Depending on the study, epidemiologists give various ranges of estimates, but B.1.1.7's transmissibility advantage is always estimated to be very high. For instance, according to this [study](#) based on British data, this advantage is estimated to be somewhere between 43% and 90%. In this post, I will focus on [Gaymard et al. \(2021\)](#), another recent study based on French data that puts B.1.1.7's transmissibility advantage between 50% and 70%, with a central estimate at 59%. Moreover, epidemiologists don't merely claim that B.1.1.7 is far more transmissible than the historical lineage, they claim that this transmissibility advantage is constant.

They don't just make those claims in scientific publications, but also in the media, where they often don't show any caution about their estimates. For instance, after that French study was published, one of the co-authors, who also happens to sit on the scientific council that advises the French government on the pandemic, gave an [interview](#) to *Le Monde*, in which one would be hard pressed to find any trace of doubt about the validity of their estimates. It is thus unsurprising that journalists and commentators talk about the hypothesis that B.1.1.7 has a constant transmissibility advantage in that range as if this were established fact. However, not only is this hypothesis not established fact, but I think at this point it's overwhelmingly unlikely to be true. Unfortunately, most epidemiologists don't seem very interested in discussing the data that are hard to reconcile with that hypothesis, so in this post I will take a closer look at Gaymard et al. (2021) to explain how they concluded that B.1.1.7 was 50% to 70% more transmissible than the historical lineage and show that it's not very serious.

Let me start by explaining quickly what Gaymard et al. did. There will be a bit of math, but don't worry it won't be long and it's okay if you don't understand everything, I will state everything you need to know in order to understand the rest of the post in plain English. In order to estimate B.1.1.7's transmissibility advantage, Gaymard et al. assumed a simple exponential growth model:

$$C_h(t) = C(0)(1 - p_0)e^{r_h t} \text{ and } C_v(t) = C(0)p_0e^{r_v t}$$

where $C(t)$ is the total number of cases at time t , $C_v(t)$ is the number of cases due to B.1.1.7 at time t , $C_h(t)$ is the number of cases due to other variants at time t , p_0 is the prevalence of B.1.1.7 on January 1, r_v is B.1.1.7's growth rate and r_h is the rate of growth of the other variants. From this model, after doing a few algebraic manipulations, it's possible to derive the expression for the prevalence of B.1.1.7 at time t

$$p(t) = \frac{1}{1 + \frac{1-p_0}{p_0}e^{(r_h - r_v)t}}$$

which as you can see depend on the initial prevalence of B.1.1.7, the rate of growth of B.1.1.7 and the rate of growth of the other variants.

What Gaymard et al. (2021) are interested in, however, is how B.1.1.7's reproduction number R_v relates to the reproduction number of the historical lineage R_h . In their model, they assume that $R_v = (1 + \alpha)R_h$, i. e. B.1.1.7 has a constant transmissibility advantage of 100α percent. Thus, the hypothesis that B.1.1.7 has a *constant* transmissibility advantage is not something they found in the data, it's something they assumed at the outset. This is a very important point, because that assumption is hardly obvious and, as we shall see, there are very good reasons to doubt that it's true. But if it's false, then even if the other assumptions of their model were true (which as we shall see shortly they definitely are not), their results are worthless. Anyway, since they are interested in R_h and R_v , Gaymard et al. need a way to connect those quantities to the equation for $p(t)$. This can be done if we know the distribution of the generation time, the time between the moment someone is infected and the moment they infect someone else, because it's possible to derive the growth rate of the epidemic if you know the reproduction number and the mean and standard deviation of the generation time.

So Gaymard et al. assumed that B.1.1.7 and the historical lineage were both growing exponentially at a rate determined by their reproduction number and the generation time of SARS-CoV-2. In order to estimate α , B.1.1.7's transmissibility advantage, they assumed that R_h remained fixed at 1 since the beginning of the year and that the generation time had a mean of 6.5 days and a standard deviation of 4 days. (They also tried with a value of 0.9 and 1.1 for R_h and with a mean of 5.5 days and a standard deviation of 3.4 for the generation time.) Once you make those assumptions, to predict how the prevalence of B.1.1.7 is going to change over time, you just need to estimate p_0 and α . The French health authorities conducted a survey to determine the prevalence of B.1.1.7 on January 7-8 and then again on January 27. Gaymard et al. used the estimates of B.1.1.7's prevalence at those dates in order to fit the model with a statistical technique called Markov chain Monte Carlo and estimate p_0 and α . You can think of Markov chain Monte Carlo as trying many possible values of p_0 and α , each time using them to simulate how $p(t)$ should have changed over time if those were the actual values, so as to figure out which values of p_0 and α result in the best fit of the data under the other assumptions of the model.

Here is the figure in the paper that summarizes the results they obtained with that method:

As you can see, in their central scenario (R_h of 1, mean of generation time of 6.5 days), B.1.1.7's transmissibility advantage is estimated to be 59%. (In case you're wondering, 501Y.V1 is just another name for B.1.1.7.) As panel D shows, in early February, the prevalence of B.1.1.7 predicted by the model fitted on data from January was very close to estimates that were not used to train the model. In sensitivity analyses, they tried different assumptions about R_h and the mean of the generation time (including a scenario in which the generation time is longer for B.1.1.7 than for the historical lineage), but the model still found a very large transmissibility advantage in every case. The range of 50% to 70% is apparently obtained by looking at the different point estimates they obtain in each of the 6 scenarios they tried.

Now, if you just look at that figure, this conclusion seems pretty convincing. It means that R is going to increase a lot as the prevalence of B.1.1.7 rises, which is why many epidemiologists have called for a strict lockdown in order to prevent the epidemic from getting out of control. However, upon closer examination, the whole thing unravels pretty quickly. First, even if we accept that a simple exponential growth model is adequate (which as we shall see we definitely should not), their sensitivity analyses clearly are not. In particular, it's worth spending some time on the choice they made for the mean and standard deviation of the generation time, because as we shall see the results are very sensitive to it. (The generation time is the time between the moment someone is infected and the moment they infect someone else.) As we have seen, in their central scenario, they assumed a mean of 6.5 days and a standard deviation of 4 days. In order to justify that choice, they cite this [paper](#) on the transmission of B.1.1.7 in England at the end of 2020, which itself cites Flaxman et al.'s infamous [paper](#) on the effect of non-pharmaceutical interventions in Europe during the first wave. (In case you're interested, I already took apart that paper, which is probably the most cited study on the effect of non-pharmaceutical interventions, in another [post](#) a few months ago.) However,

Flaxman et al. didn't even use an estimate of the generation time distribution, but with an estimate of the distribution of the serial interval they found in this [study](#) based on early Chinese data, which estimated a mean of 6.3 days and a standard deviation of 4.2 days.

The serial interval is the time between symptoms onset in the infector and symptoms onset in the person they infect, so it's not actually the same thing as the generation time. It's true that, as long as the incubation periods of infectors and infectees in a transmission pair are independent and identically distributed (which is reasonable), they have the [same](#) mean, but the generation time usually has a smaller standard deviation. The generation time is very difficult to estimate because it refers to unobservable events, which is why people frequently use estimates of the serial interval to approximate it, but again those are different concepts. Moreover, in part because it's so hard to estimate, very few studies have tried to estimate the generation time of SARS-CoV-2. I have only found 5 of them, for a total of 7 different estimates, which exhibit a large amount of heterogeneity. (I have created a [repository](#) on GitHub with the code of the analyses I did in this post, where you can also find the list of references.) The means ranged from 2.82 days to 5.5 days, while the standard deviations ranged from 1.51 days to 2.96 days. In other words, not only are there very few estimates of the generation time distribution, but they are all over the place.

There are more estimates of the serial interval, which is still informative even if you're interested in the generation time since, as we have seen, they have the same mean and the standard deviation of the serial interval provides an upper bound for that of the generation time. A recent [meta-analysis](#) found 23 estimates of the serial interval. The means range from 4.2 days to 7 days, while the standard deviations range from 0.95 to 5.8, so again they're all over the place. (Another [meta-analysis](#) I will come back to shortly found means ranging from 3.95 days to 7.5 days, but estimated a serial interval as short as 2.09 days using data collected by Public Health England.) So whether we look at published estimates of the generation time or the serial interval, the mean chosen by Gaymard et al. is on the higher end, but as we shall see assuming a longer generation time results in a higher transmissibility advantage for B.1.1.7. Although it's easier to estimate the serial interval than the generation time, because unlike infection symptoms onset is an observable event, it's still extremely difficult and those estimates should be taken with a huge grain of salt. Indeed, not only are they generally based on very small samples (typically a few dozens people), but the data are usually very low quality. For instance, researchers have to rely on people's memory to determine the time of symptoms onset, which it's generally unreliable because of [recall bias](#). The samples also probably aren't representative, if only because asymptomatic cases are underrepresented.

To make things worse, be it for the generation time or the serial interval, estimates are almost exclusively based on Chinese data from early in the pandemic. This is problematic for at least 2 reasons. First, we are interested in the generation time in France, but there is no reason that it's the same as that in China and in fact the literature suggests there are very large between-country differences. Second, even in the same country, there is no reason to think the generation time distribution remained the same over time and in fact there are excellent reasons to think it has shrunk over time. For instance, there was a shortage of tests in the first months of the pandemic, but as they became more widely available it became easier for people to find out they have been infected and they were able to isolate more quickly, which presumably reduced the generation time since once people are isolated they have fewer opportunities to infect other people. This hypothesis has some empirical support from a study which found that, after non-pharmaceutical interventions were introduced in China, the mean of the serial interval decreased from 7.8 days to 2.2 days. (The meta-analysis cited above doesn't include this paper, but it includes another [paper](#) from some of the same co-authors that found a mean of 5 days, presumably because it aggregates people from both periods.) Thus, not only is the assumption made by Gaymard et al. about the mean of the generation time on the higher end of the published estimates, but those estimates probably overestimate the generation time in France in 2021. This is also true of the mean of 5.5 days which they used for their sensitivity analysis.

So what assumptions should Gaymard et al. have made about the distribution of the generation time? There is no easy answer to that question, but one thing we can say for sure is that they shouldn't have assumed a mean of 6.5 days and a standard deviation of 4 days. Indeed, given the range of estimates in the literature, both the mean and the standard deviation seem way too high to be used as a best estimate to model transmission. Even a mean of 5.5 days, as they used in a sensitivity analysis, seems too high, especially when you keep in mind what I said about how the generation time in France right now is probably much shorter than during the period when the data used to derive the estimates in the literature were collected. One thing we could do is use the estimate of the generation time distribution found in [Challen et al. \(2020\)](#), a meta-analysis I already mentioned above, which ultimately is based on what they considered a best estimate of the incubation period distribution and their meta-analytical estimate of the serial interval distribution. The resulting distribution has a mean of 4.8 days and a standard deviation of 1.7 days.

No thanks to Gaymard et al., who as usual not only did not publish their code but also didn't reply to my emails asking if they could send it to me, I was able to replicate their results by writing a model similar to what they described in the supplementary materials of their paper. (I think everyone should publish their code along with their papers and that people who don't should be systematically shamed, but it's particularly unacceptable when the authors are paid with my taxes and don't even reply to emails in which I politely ask for it.) In particular, when I assume a R of 1 for the historical lineage, a mean of 6.5 days and a standard deviation of 4 days for the generation time (as in their central scenario), I estimate B.1.1.7's transmissibility advantage at approximately 59%, the same thing Gaymard et al. report in their paper. Moreover, I also get the same results when I run the same sensitivity analyses, so I'm pretty confident that there is no major difference between the model I put together and the one they used. Now what happens when I use the same model but assume a mean of 4.8 days and a standard deviation of 1.7 days for the generation time? Under those assumptions, keeping a R of 1 for the historical lineage, B.1.1.7's transmissibility advantage drops to 44%. (By making a few simplifying assumptions, it's easy to prove analytically that B.1.1.7's transmissibility advantage is positively related to the mean of the generation time distribution, even if that distribution is assumed to be identical for B.1.1.7 and the historical lineage. I sketch the proof in the comments of my code on GitHub if you're interested.) Thus, as you can see, making more realistic assumptions about the generation time distribution resulted in a substantial reduction of the estimate of B.1.1.7's transmissibility advantage.

But while it's better than what Gaymard et al. have done, I think it's still hardly ideal to use Challen et al.'s estimate of the generation time distribution, because ultimately it's based on published estimates of the serial interval and the incubation period. Now, as we have seen, not only are they probably not very good even for the early phase of the pandemic in China (where most of the data used to obtain those estimate came from), but there are very good reasons to think that, even if they were, they would not apply to the situation in France right now. The truth is that we have no idea what the generation time distribution is right now in France or in most other places for that matter. The mean and standard deviation could be anywhere in the range of published estimates and even outside of it, we just don't know and even a meta-analysis is not particularly useful, because it pools estimates based on data collected during the early phase of the pandemic in China and a handful of other places. Given this uncertainty, even if Gaymard et al. hadn't made assumptions that are on the very high end of the published estimates, it would have been fundamentally unserious to try only 2 specifications of the generation time distribution. Unfortunately, Gaymard et al. are hardly exceptional in that respect, since it's more or less what most epidemiologists do to model transmission. But the fact that most epidemiologists do that doesn't mean that it's not problematic, it just means that the epidemiological literature on SARS-CoV-2 systematically underestimates true uncertainty.

In my opinion, given the uncertainty about the generation time distribution that applied in France at the beginning of 2021, the only reasonable option is to do sensitivity analyses in a far more systematical way than Gaymard et al. and try a wide range of possible values for the parameters of that distribution. So what I did is run $3 \times 41 \times 16 = 1968$ different specifications of their model. I tried every value between 3 days and 7 days for the mean of the generation time distribution by increments of 0.1 day, every value between 1.5 days and 3 days for the standard deviation of that same distribution by increment of 0.1 day and each of the 3 values of R for the historical lineage that Gaymard et al. considered. I created a histogram, on which I superimposed a density plot, that summarizes the results:

As you can see, based on this analysis, B.1.1.7's transmissibility estimate could be anywhere between 21% and 72%. This is much wider than the 50% to 70% range that Gaymard et al. reported, not just in their paper but also in the media. If I was able to do that on my 2014 MacBook Pro, surely they could have done it on the far more powerful computers they have in their research institutions, so there is absolutely no excuse for doing the botched sensitivity analysis they and the reviewers of their paper considered sufficient.

As we have seen, Gaymard et al. used a simple exponential growth model, which means that both the historical lineage and B.1.1.7 were assumed to grow at a constant rate during the period of study. This may have been approximately true in January, although even that is doubtful, but it has definitely not been true since then. One way to see that is to plot the model's prediction for the prevalence of B.1.1.7 over time and use the latest data about variants to compare this prediction to what actually happened:

As you can see, the model's prediction closely matched reality for a while, but eventually this ceased to be true because the prevalence of B.1.1.7 didn't increase as fast as predicted by the model in March.

It's even more clear if you plot the incidence predicted by the model and compare it to the actual incidence during the same period, meaning the total number of cases in France. Here is what it looks like when I make the same assumptions as Gaymard et al. in their central scenario:

As you can see, while the model predicted the rise in prevalence of B.1.1.7 pretty well until March, predicted incidence started diverging from reality at the beginning of February and ended up being completely off. Indeed, the model predicted that incidence would grow exponentially and that by April 1 there would be more than 500,000 cases a day, which clearly is not what happened.

In part, that's because Gaymard et al. assumed a particular value for the reproduction number of the historical lineage instead of estimating it from the data and they effectively fitted their model to the prevalence of B.1.1.7 instead of its incidence, but for the most part it's because neither B.1.1.7 nor the historical lineage grew at a constant rate during that period. We know that because, by using a relation that holds approximately between the growth rate of incidence and the effective reproduction number of the epidemic, it's possible to estimate the effective reproduction number of B.1.1.7 at various points and see how it has changed as this lineage rose in prevalence:

The particular value of R depends on the assumption we make about the mean generation time, but as long as the generation time distribution remained approximately the same in France since the beginning of the year (which is a reasonable assumption), we'll find a downward slope no matter what assumption we make about the generation time distribution. Indeed, even accounting for the fact that data quality is not great and that we don't really know what the mean generation time is, it's clear that B.1.1.7's effective reproduction number has gone down significantly in France since the beginning of the year.

Had this not been the case, since the prevalence of B.1.1.7 rose from ~3% at the beginning of the year to more than 80% now, the overall effective reproduction number would have increased massively. Indeed, as panel E in the figure I reproduced above shows, Gaymard et al.'s model predicted that it would have increased by more than 50% between the beginning of year and now. But as you can see, that's not what actually happened:

Instead of monotonously increasing as B.1.1.7 rose in prevalence, R hovered around 1 during most of that period and, while it has increased recently, it's no higher than it was at the beginning of January and it's certainly nowhere as high as Gaymard et al.'s model predicted.

Unfortunately, because it's been drummed into people's heads that epidemics always grow exponentially, most people are unaware of that fact and assume that B.1.1.7 has been growing at the same rate since the beginning of the year, which was balanced out for a while by the fall of the historical lineage, but is about to result in a catastrophic rise of incidence. (This belief is hardly unique to French commentators, but is also widespread in the US, where articles such as [this](#) warning about the impending doom that B.1.1.7 and other variants will bring are perhaps the only thing growing exponentially right now.) What is more surprising is that even some professional epidemiologists, such as Dominique Costagliola, [seem](#) to believe that and indeed are largely responsible for the ubiquitousness of this narrative in the media. (Costagliola is not just any epidemiologist, but was awarded the Great Prize by Inserm, France's main research institution for medical research, in 2020. Interestingly, she didn't reply when I offered to bet money on how well this model would predict incidence, so perhaps she is not as clueless as she seems. Nevertheless, it's incredible that she has no qualms about demanding that 67 million people be locked down every time she talks to the media, but won't even bet a few hundreds euros that the model she apparently uses to make that recommendation will not be massively wrong.) Yet as we have seen, not only is this narrative false, but it only takes a few minutes to show that it's wrong, since you just have to plot B.1.1.7's growth rate over time.

Ironically, if B.1.1.7 had really grown exponentially in France since the beginning of the year, it would mean that it has a much smaller transmissibility advantage than estimated by Gaymard et al. have concluded by using only data from January. Indeed, if I assume a simple exponential growth model but fit it to incidence data by variant up until the end of March and estimate R for the historical lineage from the data instead of making a gratuitous assumption about it, I find that B.1.1.7 is only 27% more transmissible than the historical lineage with a mean generation time of 4.8 days and a standard deviation of 1.7. Again, we don't really know what the generation time distribution is and even this meta-analytic estimate is not very reliable, so here is what the range of estimates look like when I run $41 \times 16 = 656$ different specifications of the model spanning the same parameter space as before:

As you can see, depending on what assumptions you make about the parameters of the generation time distribution, B.1.1.7's transmissibility advantage could be anywhere between 16% and 42% if we assume a simple exponential growth model. So although the people who insist that B.1.1.7 has grown exponentially since the beginning of the year clearly don't realize it, this claim is actually inconsistent with the view they also repeat all day long that it's between 50% and 70% more transmissible than the historical lineage. Of course, this model is false and would perform terribly if used to predict the course of the epidemic, but this is what you should say about B.1.1.7's transmissibility advantage if you believe it has grown exponentially since the beginning of the year.

Most epidemiologists probably realize that B.1.1.7 has not grown exponentially, but they insist the disaster they predicted was only postponed thanks to the fact that the curfew was advanced to 6pm from 8pm on January 16 or they ascribe the failure of their predictions to come true to some other *deus ex machina*, such as the school holidays in February. Here is what Simon Cauchemez, one of the co-authors of Gaymard et al. (2021), told Le Monde during the [interview](#) I mentioned above:

How can the current epidemic plateau be explained?

One of the persistent difficulties for modellers is to anticipate the effects on transmission rates of measures that have never been taken before, such as the current curfews. We are forced to reason by analogy with measures taken previously, such as the first and second lockdowns and the effects they had on transmissibility. And even this reasoning by analogy has its limits, because compliance by the population is an important factor that can change over time.

For example, in January, when the first curfews were put in place, the data we had on mobility showed that these measures did not seem to change much. At least not as much as during the October 2020 curfews. So we were worried about a small impact. We now see that the curfew that started in January has reduced the transmission rate of the historical virus quite significantly, leading to a plateau in hospitalizations in the second half of January. We even saw a slight decrease in hospitalizations in the first half of February. This suggests an even greater reduction in transmission rates over this period, which was a surprise, as control measures did not really change at this time.

How to explain it?

Perhaps by the combined effect of the curfew, the vacations, a "severe cold" effect or a resumption of precautions and a readjustment of behavior, because there was a lot of talk at the time of "hard" lockdown. Unfortunately, transmission rates have since increased again.

As you can see, at no point does he question the conclusion of their study (i. e. that B.1.1.7 has a constant transmissibility advantage over the historical lineage of 50% to 70%), he just comes up with a bunch of *ad hoc* explanations for why the explosion they predicted didn't materialize and suggests that the current rise in incidence, which is nothing like what epidemiologists [were](#) predicting at the end of January, proves they were right after all.

Epidemiologists don't just use this line of argument when they talk to the media, but also in their scientific publications. For instance, in a [paper](#) whose conclusions were presented by one of the authors during a press conference organized by the French government on February 18, another team at Inserm argued that both the curfew in January and the school holidays in February prevented the explosion, but that it was only a temporary respite since the effective reproduction number would quickly rise as schools reopened between the end of February and the beginning of March:

Facing B.1.1.7 variant, social distancing was strengthened in France in January 2021. Using a 2-strain mathematical model calibrated on genomic surveillance, we estimated that curfew measures allowed hospitalizations to plateau, by decreasing transmission of the historical strain while B.1.1.7 continued to grow. School holidays appear to have further slowed down progression in February. Without progressively strengthened social distancing, a rapid surge of hospitalizations is expected, despite the foreseen increase in vaccination rhythm.

Many likely believe this paper shows that, even though advancing the curfew to 6pm on January 16 was not associated with any sudden reduction of R , it did have a very large effect on transmission and so did the school holidays in February. But it does no such thing.

It's hard to know exactly what they did, since as usual they didn't publish their code and, when I emailed the corresponding author to ask if she could send it to me, she didn't reply. (If this is starting to sound like a pattern among French epidemiologists paid on the public dime, that's because it is and it has been like that since the beginning of the pandemic.) However, the model is apparently based on that used in another [publication](#) from last year whose code is available, so it's possible to get a rough sense of what they did. I won't discuss the details of their model because it doesn't matter, but here is what you need to understand. First, the model assumes that B.1.1.7 has a constant transmissibility advantage over the historical lineage of 60% (they also tried 50% and 70% for sensitivity), which they justify by citing Gaymard et al. (2021) and previous studies based on British data. Thus, B.1.1.7's transmissibility advantage is not estimated from the data, but assumed to have a particular value at the outset. Moreover, the model assumes that only the curfew or school holidays could have affected transmission, so basically what they did is use a model that rests on those assumptions to estimate what effect the curfew and the school holidays must have had in order to fit the data as well as possible. Unsurprisingly, since the explosion predicted by the same team at the end of January didn't occur, the model concluded that both the school holidays in February and even more so advancing the curfew to 6pm on January 16 had a very large effect on transmission, because that's the only way the model could possibly have fitted the data!

I was able to replicate this result with a simpler model that makes the same assumptions, but instead of assuming B.1.1.7's was 60% more transmissible at the outset, I asked the model to estimate its transmissibility advantage from the data. (I won't go into the details of the model, but I basically used a discrete version of a SIR model where the curfew and school holidays are assumed to immediately affect R , which I estimated by MCMC with weakly informative priors on the parameters of interest. I didn't try to model the impact of vaccination on transmission, but it shouldn't make a big difference because unfortunately France has been vaccinating at a very slow pace anyway. I also didn't model the effect of immunity induced by natural infection, but this wouldn't have made a large difference in a model with homogenous population mixing. If you want to know more, I invite you to check the code.) Here is a figure that summarizes the results when I assume the generation time distribution has a mean of 4.8 days and a standard deviation of 1.7 days:

As you can see, the model finds that advancing the curfew to 6pm had a massive impact on transmission, while school holidays had a significant though smaller effect. (French schools are divided into 3 administrative zones that don't go on holidays at the same time, which is why the effect of school holidays in February doesn't occur all at once.) The model estimates that B.1.1.7 has a transmissibility advantage of 35%, which again is significantly less than Gaymard et al.'s 59% estimate. The curfew is estimated to have reduced transmission by 25%, while the model finds that school holidays in February reduced it by less than 10% when schools were closed in all 3 administrative zones.

To be clear, although epidemiologists routinely present that kind of analysis with a straight face, it's completely ludicrous and I only indulged in that exercise to show that, even if I accept their framework, the estimate of B.1.1.7's transmissibility advantage is lower than what Gaymard et al. found based on fitting a simple exponential growth model on 2 data points from January. If you really believe that advancing the curfew from 8pm to 6pm reduced transmission by 25%, I hope that you will come back to reality soon, because we miss you here. As I explained in my [critique of Flaxman et al. \(2020\)](#) and more recently in my [case against lockdowns](#), the fundamental problem with that kind of model is that it assumes that, in addition to immunity, only government interventions or events like school holidays affect transmission, but we know that's false. Transmission is affected by many other things beside government interventions and school holidays, such as people's voluntary behavior modifications in response to changes in epidemic conditions, which as I have been arguing for months we have very good reasons to believe have a much larger effect than government interventions. The result is that, if you tell a model that only government interventions affect transmission and incidence doesn't continue to grow exponentially until the herd immunity threshold is reached (which it never does), the model is going to ascribe that to government interventions because there is nothing else it could ascribe it to, but this obviously doesn't prove anything and the fact that epidemiologists take that sort of things seriously speaks volumes about what a joke their field is.

In addition, this model assumes that B.1.1.7's transmissibility advantage is constant, as we can see by plotting the effective reproduction number of B.1.1.7 and that of the other variants:

The green and red curves are parallel, which indicates that B.1.1.7 has a constant transmissibility advantage over the other variants, but as we shall see this has not been true.

I didn't do a thorough sensitivity analysis by trying hundreds of specifications that make different assumptions about the mean and standard deviation of the generation time, because this model is more computationally-intensive to train than a simple exponential growth model and it would have taken several days on my laptop, but I don't need to in order to figure out the range of estimates since B.1.1.7's transmissibility advantage monotonously increases with the mean of the generation time distribution and monotonously decreases with the standard deviation. Thus, by estimating the model with a mean of 3 days and a standard deviation of 3 days for the generation time distribution on the one hand and a mean of 7 days and a standard deviation of 1.5 days on the other hand, I can tell that if I had run a sensitivity analysis spanning the same parameter space as before I would have found that B.1.1.7's transmissibility advantage is somewhere between 22% and 53% according to this model. Again, the assumptions you make about the generation time distribution have a substantial effect on the estimate of B.1.1.7's transmissibility advantage, but under almost every specification of the model the estimate is significantly lower than even the lower bound of the range given by Gaymard et al.

We could try to estimate the parameters of the generation time distribution directly from the data instead of making assumptions about them, but it would be a waste of electricity since, as I will argue shortly, the whole exercise is completely pointless because it rests on the assumption that B.1.1.7 had a constant transmissibility advantage over the historical lineage and that assumption is false. Before I show that, however, I want to try another approach to estimate B.1.1.7's transmissibility advantage on the assumption that it's constant. As we have just seen, epidemiologists do so by modeling transmission, which requires that they make strong mechanistic assumptions beyond the assumption of constant transmissibility advantage and the assumptions they make about the generation time distribution. But another possibility is to use econometric methods to estimate B.1.1.7's transmissibility advantage by looking at the

correlation between the prevalence of B.1.1.7 in a department and the effective reproduction number in that department. There are 96 departments in metropolitan France and we have 6 weeks of data on variants. So we know the average prevalence of B.1.1.7 by department for each week and we can easily compute the growth of incidence in a department from one week to the next. From those weekly rates of growth, by assuming a particular value for the mean of the generation time distribution, we can obtain an approximation of R in every department during each period in each department.

Here is what it looks like when you plot R on the prevalence of B.1.1.7:

As you can see, R in a department is positively associated with the prevalence of B.1.1.7 in that department, which is what you'd expect if B.1.1.7 had a transmissibility advantage over the other variants.

However, this could easily be misleading, so one can't draw any conclusion from that. For instance, this association could be due to the fact that both R and the prevalence of B.1.1.7 have increased over time, even if there is no association or even a negative association during any given period. In fact, if you do the same thing as above but look at each period separately, the result suggests this might be a real concern:

Except for the period from week 7 to week 8, there doesn't seem to be any association between R and the prevalence of B.1.1.7, which is not for lack of variation in the data since there is plenty of it for each period.

In the hopes of sorting this out, we can estimate variations of the following model:

$$\log(R_{i,t}) = \beta p_{i,t} + \Delta_t + \gamma_i + \epsilon_{i,t}$$

where $R_{i,t}$ is the effective reproduction number in department i during period t , $p_{i,t}$ is the prevalence of B.1.1.7 in department i during period t , Δ_t is a period fixed effect, γ_i is a department fixed effect and $\epsilon_{i,t}$ is a random error term. Don't worry if you don't understand what this means, it's less complicated than it seems. First, this model assumes that $R_{i,t}$ is a linear function of $p_{i,t}$, which means that every time the prevalence of B.1.1.7 increases by 1 percentage point in a department the effective reproduction number in that department should increase by approximately β and this relation doesn't change as the prevalence of B.1.1.7 rises. This is effectively the same as assuming that B.1.1.7 has a constant transmissibility advantage over the historical lineage. As for Δ_t and γ_i , they are here to capture the effects specific to one period or one department. For instance, if the reproduction number increases across departments during a period because it was cold so people spent more time inside, it will be captured by Δ_t . Similarly, if something about a department makes the reproduction number higher than in others at the same time even if the prevalence of B.1.1.7 is the same, e. g. because it's more densely populated, this will be captured by γ_i . Finally, $\epsilon_{i,t}$ will on average be zero and is just here to capture everything else that might affect $R_{i,t}$, as long as it's not correlated with the prevalence of B.1.1.7 in a department.

Without going into the details, it's not easy to figure out whether it's best to include Δ_t and γ_i in the model, just one of them or neither of them. For instance, if we include a period fixed effect, we are assuming that any shock to the effective reproduction number associated with a particular week is the same across departments, but this might not be true if that shock is due to the weather, because not all departments are subject to the same changes in weather. Similarly, if we include a department fixed effect, we are assuming that anything that makes the effective reproduction number higher or lower in a department is constant over time. Obviously, this will be true if the factor in question is population density, but it might not be true for other factors. If any of those assumptions is false, our estimate of the effect of the prevalence of B.1.1.7 will be biased (to be clear, bias is not the only issue with fixed effects, which among other things also reduce statistical power), so I think it's best to try every specification of the model and see the results we get with each of them, because we don't really know which specification is the best.

Here is a table that summarizes the results of fitting each specification of this simple econometric model:

The prevalence of B.1.1.7 has a statistically significant effect on the effective reproduction number in all specifications of the model except when only a period fixed effect is included.

We can convert those coefficients into estimates of B.1.1.7's transmissibility advantage and add confidence intervals, which makes them more readily interpretable:

The point estimates range from 10% to 41%, which is much lower than Gaymard et al.'s 50% to 70% range. However, they are very imprecisely estimated and, if you look at the confidence intervals, B.1.1.7 could be anywhere between 9% less transmissible and 65% more transmissible. This range would be even wider if I tried different assumptions about the generation time distribution to convert weekly growth rate into effective reproduction numbers. For instance, if I assume that it has a mean of 6.5 days (the value used by Gaymard et al. in their central scenario), the point estimates range from 14% to 60% and the confidence intervals put B.1.1.7's transmissibility advantage somewhere between -12% and 98%. If I assume a mean of 3.5 days, point estimates range from 7% to 29%, while the confidence intervals imply that B.1.1.7 could be anywhere between 7% less transmissible and 44% more transmissible.

Now let's take stock and summarize what we have found so far. Gaymard et al. use a simple exponential growth model they fit on data about variants from January and estimate that B.1.1.7 is 50% to 70% more transmissible than the historical lineage. However, even if we use the same type of model on the same early data, we find that B.1.1.7's transmissibility estimate could be anywhere between 21% and 72% depending on what assumptions we make about the generation time distribution. If we continue to assume a simple exponential growth model but use up to date data, which as we have seen is a mistake because despite what many people think B.1.1.7 has not grown at a constant rate since the beginning of the year, we find that it could be anywhere between 16% and 42%. If we try to model the effect that advancing the curfew to 6pm in January and school holidays in February had on transmission, in a way similar to what the epidemiologists who advise the French government do (except that instead of assuming that B.1.1.7 is 50% to 70% more transmissible I estimate this advantage from the data), we find that B.1.1.7's transmissibility advantage could be anywhere between 22% and 53% depending on what assumptions we make about the generation time distribution, which is a wide range but is still significantly lower than Gaymard et al.'s 50% to 70% range.

As we have just seen, estimates of B.1.1.7's transmissibility advantage are even more all over the place if, instead of making strong mechanistic assumptions to model transmission, we use the more agnostic econometric approach. Indeed, depending on what model you're using and what assumptions you make about the generation time distribution to convert weekly growth rate into effective reproduction numbers, you can't even rule out that B.1.1.7 is *less* transmissible than the historical lineage. (To be clear, I don't believe for a second that it is, I'm just pointing out that the data are sufficiently ambiguous that you can't formally exclude it.) In short, even if we assume that B.1.1.7 has a *constant* transmissibility advantage over the historical lineage, we have no idea what this advantage really is because estimates vary wildly depending on what kind of models we use and we have no way of knowing which model is best. But this doesn't stop epidemiologists from going in the media and claiming as if this were established fact that B.1.1.7 is 50% to 70% more transmissible than the historical lineage, when in fact they reached that conclusion by assuming a ridiculous model and fitting it on 2 data points from January. I don't know if they are knowingly misrepresenting how much uncertainty there is or if they are just incompetent, but I don't really care because either way it's clear that my taxes should not be used to pay them.

This is already bad enough, but don't go anywhere, because it's about to get a lot worse. Every single approach we have used so far to estimate B.1.1.7's transmissibility advantage was based on the assumption that it's constant. Although epidemiologists rarely make that explicit, this assumption underlies virtual every study on B.1.1.7's transmissibility. However, despite what many people seem to think (which is not surprising given how epidemiologists present their findings in the media), this is not something they found in the data. They just assume it's true and estimate B.1.1.7's transmissibility advantage under that assumption. But is this assumption true? If it's not, even putting aside the fact that as we have seen they are not robust to reasonable changes in specification of the model, estimates such as those in Gaymard et al. (2021) are completely worthless. Indeed, Gaymard et al.'s estimates are based on the growth of B.1.1.7 in January when it made up a very small share of cases, so they are bound to be misleading if B.1.1.7's transmissibility advantage has gone down since then. It's actually very easy to check if the data are consistent with it, so you'd think that epidemiologists would do it before going in the media and claiming that B.1.1.7 is 50% to 70% more transmissible, but apparently they couldn't be bothered.

As we have seen, B.1.1.7's effective reproduction number has fallen a lot in France since January, but this doesn't mean that its transmissibility advantage didn't remain constant as long as the effective reproduction number of the historical lineage has fallen in the same proportion. However, if you actually compute B.1.1.7's transmissibility advantage from the data, this is not what you observe:

As you can see, at the national level, B.1.1.7's transmissibility advantage as measured directly from the growth rate in cases by variant has gone from a high of 57% to 11% today. If you make different assumptions about the mean of the generation time distribution, you will find somewhat different results, but the overall pattern will be the same.

This pattern is also very clear if, instead of aggregating data at the national level, you look at B.1.1.7's transmissibility advantage at the department level:

It would probably be even clearer if we had data on variants by department before week 7, but unfortunately if they exist the French health authorities have not published them. In many departments, such as Paris (where incidence is higher than almost anywhere else in the country), B.1.1.7 is now growing *less* rapidly than the other variants and there are plenty of cases so this is not a sampling artifact. How do people who claim that it's 50% to 70% more transmissible explain that? They should at least grapple with those facts, but instead they completely ignore them.

I don't know why B.1.1.7's transmissibility advantage has gone down over time, but whatever the underlying mechanism, it's clear that it has. One possible hypothesis that has been [proposed](#) is that different groups are not equally susceptible to all variants and that some groups that are relatively less susceptible to previous strains are relatively more susceptible to B.1.1.7, so upon being introduced in the population B.1.1.7 would find it easier to spread than the historical lineage because the people who are most susceptible to it would have been relatively spared up until that point, but this would progressively cease to be true as more of them get infected. This is just one of several possible explanations that have been proposed, none of which I find completely satisfactory, but whatever the underlying mechanism, it's clear that B.1.1.7's transmissibility advantage has gone down as it rose in prevalence. (As Stephen McIntyre [noted](#) on Twitter, the same thing is apparently true in the UK, a fact that the authors of the unpublished analysis he cites don't seem to be in a hurry to advertise.) This is presumably why the projections that predicted a catastrophic explosion of incidence have failed to come true over and over everywhere.

Epidemiologists should admit that fact and try to figure out what the underlying mechanism is, but instead they just pretend it didn't happen and continue to churn out apocalyptic predictions they use to call for more stringent restrictions. The way in which French epidemiologists have talked about B.1.1.7 since the beginning of the year is a perfect illustration of that kind of behavior. Even though B.1.1.7's transmissibility advantage, as measured from weekly growth rates, has been less than 20% for more than a month and was only 11% based on the latest weekly growth (again the precise figure will vary depending on what assumptions you make about the generation time distribution), epidemiologists have totally ignored that and continue to assume that it's between 50% and 70% for their projections, a range that was obtained by fitting a ridiculous model to 2 data points in January. Of course, since B.1.1.7 is not really 50% to 70% more transmissible than the historical lineage or at least not anymore, their projections have repeatedly proven completely wrong. But instead of admitting that and trying to figure out why B.1.1.7's transmissibility advantage has gone down, they come up with *ad hoc* explanations for why their projections failed to come true even though they were totally right.

This is very easy because, as I have explained above, you can always retrospectively fit a model that will "show" that policies adopted by the government after they made their initial projections had a massive effect on transmission because under the assumption that B.1.1.7 is 50% to 70% more transmissible than the historical lineage it's the only way this model can explain why incidence didn't explode, even if nobody with half a brain can seriously believe that such policies had such a massive effect. Of course, the projections based on this model will also prove completely wrong, but it won't be a problem for epidemiologists because they'll just fit another model on more recent data and it will confirm that we only avoided a disaster thanks to the policies adopted by the government on their advice. For instance, the French government recently closed most small businesses, where according to [contact tracing data](#) less than 0.07% of all infections take place. (Actually, this figure is for all businesses included those that are not affected by this decision, but on the other hand it's only based on cases whose source of contamination is known and it's only about customers.) This will cost a lot of money and it obviously won't have any effect on transmission, but you can be certain that a month from now epidemiologists will "show" that it had a massive effect, which is why there isn't 500,000 cases a day. In short, epidemiologists are very good at predicting the past, but they're having a harder time with the future.

It's genuinely very difficult to predict the course of the epidemic, so I don't fault them for not being able to do it. I fault them for pretending they can, failing to admit when they were wrong and, precisely because they don't admit it, not updating on new evidence as it comes in. Again, it's really mind-blowing that, even as the data now indicate that B.1.1.7 is only about 11% more transmissible than the other variants, epidemiologists continue to use estimates based on the early expansion of that lineage in January and consequently still assume it's 50% to 70% more transmissible for the projections they present to the government. As I already said, I'm not sure if that's because they're incompetent or dishonest, but I suspect it's both. They are clearly biased toward pessimistic predictions because they support more stringent restrictions and they know that apocalyptic predictions will pressure the government into tightening restrictions, but I think they are clueless enough to buy their own hype. Unfortunately, most people don't understand how epidemic projections work and treat epidemiological models like magic, so epidemiologists have no difficulty imposing their narrative in the media. They have been so successful that now every time there is a flare-up of the epidemic somewhere, everyone assumes it's because of B.1.1.7 even though it's clearly not as transmissible as initially thought.

I have focused on France but, as I noted above, this conclusion doesn't crucially hinge on French data. B.1.1.7 now dominates everywhere in Europe, but it didn't result in the apocalyptic explosion of incidence anywhere. People often bring up the third wave in the UK or Ireland as proof that B.1.1.7 is far more transmissible than the historical lineage, but incidence fell even more rapidly than during the first lockdown, even though mobility data suggests people's

behavior changed considerably less. Even taking into account vaccination and natural immunity, this is not what you'd expect if B.1.1.7 were really 50% to 70%. Meanwhile, Spain also had a massive wave in January back when the prevalence of B.1.1.7 was still very low, but incidence fell very quickly after that despite the lack of lockdown and remains very low today even though B.1.1.7 is now the dominant strain in every region and the country is even more open than it was in January. It's true that in several countries, such as France, incidence started to increase again, but there is no need to assume that B.1.1.7 is super-transmissible to explain that. Indeed, previous waves were worse in most of them, yet the prevalence of immunity, whether acquired by natural infection or vaccination, couldn't explain that if B.1.1.7 really were 50% to 70% more transmissible. Such resurgences of the epidemic have happened several times in the past before B.1.1.7 started to expand and, as long as not enough people have acquired immunity through vaccination or natural infection, there is no reason why it could not happen again.

One comment

- Pingback: [Covid 4/22: Crisis in India | Don't Worry About the Vase](#)

Comments are closed.

[previous](#)[The Case against Lockdowns](#)

[next](#)[New Mission at CSPI](#)

Copyright by CSPI 2020. All rights reserved.