

How Philosophers of Science Violated Their Epistemic Duties during the SARS-CoV-2 Crisis

Philippe van Basshuysen and Lucie White

ABSTRACT. Were governments justified in imposing lockdowns to contain the spread of the COVID-19 pandemic? We argue that a convincing answer to this question is to date wanting, by critically analyzing the factual basis of a recent paper, “How Government Leaders Violated Their Epistemic Duties During the SARS-CoV-2 Crisis” (Winsberg et al. 2020). In their paper, Winsberg et al. argue that government leaders did not, at the beginning of the pandemic, meet the epistemic requirements necessitated to impose lockdowns. We focus on Winsberg et al.’s contentions that knowledge about COVID-19 resultant projections were inadequate; that epidemiologists were biased in their estimates of relevant figures; that there was insufficient evidence supporting the efficacy of lockdowns; and that lockdowns cause more harm than good. We argue that none of these claims are sufficiently supported by evidence, thus impairing their case against lockdowns, and leaving open the question of whether lockdowns were justified.

1. INTRODUCTION

In spring 2020, many governments worldwide imposed lockdowns in order to contain the spread of a novel coronavirus, SARS-CoV-2, and to prevent healthcare systems from becoming overwhelmed by an influx of severe cases of COVID-19, respiratory disease that is caused by the virus. It is estimated that, on the 5th of April, a maximum of 4.4 billion people, or 57% of the global population, were under full or partial lockdown (Bates et al. 2020); that is, people were permitted to leave their homes only under specific conditions, if at all, and these confinements were often accompanied by additional measures, such as closures of schools and businesses. Because the lockdowns restricted basic liberties and led to severe economic damage (Mandel and Veetil 2020), the question of whether governments were justified in imposing them comes to the fore. We aim to show that a convincing answer to this question is to date owing, by arguing that a recent paper by Eric Winsberg, Jason Brennan and Chris W. Surprenant (2020), according to which these lockdowns were *un*justified, is based on false factual claims and does therefore not succeed in motivating their conclusion.

In their paper, Winsberg et al. argue that “government leaders failed and have continued to fail to meet their epistemic duties” (2020, 216) when imposing lockdowns as a response to the COVID-19 pandemic. They argue for this claim by, first, endorsing the liberal view that, in order

to restrict basic liberties, states must abide by strong epistemic standards; for instance, it would not be justified for a state to imprison someone for a crime without collecting sufficient evidence that he is guilty. Second, they contend that the available evidence by the time of the first lockdowns (and perhaps even now) was simply not good enough to justify lockdowns: “states relied upon bad data and flawed models, and they lacked the other kinds of evidence they would need to justify lockdowns” (ibid.). We will focus here on the second part of their argument. Specifically, we will argue that Winsberg et al. mistakenly contend that the influential modeling studies by Imperial College London (ICL) were overly pessimistic in their projections of ICU demand and death rates (Section 2); that they fail to underpin their claim that epidemiologists are influenced by their values and external pressures to systematically produce overblown forecasts (Section 3); that they falsely maintain that there was no evidence for the efficacy of lockdowns and they make invalid charges against a study that analyzes the efficacy of the lockdown in California (Friedson et al. 2020) (Section 4); and that they erroneously claim that deaths from lockdowns may have been comparable to deaths from COVID-19 infections in the US by May 2020 (Section 5).

The implication of our factual corrections is that Winsberg et al. fail to substantiate the second part of their argument, and thus their conclusion that lockdowns were unjustified does not follow. Whether the evidence that was available when governments imposed lockdowns did indeed meet the high justificatory standards that Winsberg et al. contend are required is beyond the scope of this paper, as is the question of whether this benchmark must be met in order for democratic governments to legitimately impose lockdowns, although these highly interesting arguments are certainly worthy of further discussion.

2. PROBLEMS WITH DATA AND THE MODELS?

A central claim in Winsberg et al.’s paper is that there was significant uncertainty concerning many aspects of the disease, and sparse data, in spring 2020, when governments around the world imposed lockdown policies. In constructing epidemiological models to forecast the course of the disease, and the potential impact of various interventions, modelers must thus make unconstrained assumptions, or assumptions based on poor data. The result, the authors argue, are wildly unreliable predictions, which do not form an adequate basis for public policy. More specifically, they argue that the projections of these models are “overly pessimistic” (2020, 226). The model that forms much of their focus in substantiating these claims is the ICL model,

which, as the authors rightly note, had a significant impact on policy decisions in the UK.¹ We will thus take the same focus in our analysis of their claims.

The ICL model, as the authors note, “was used to estimate what public interventions would be needed to prevent hospital systems from being overwhelmed” (2020, 223). The model produces predictions concerning the total amount of COVID-19 deaths (in a 2-year period) that will eventuate and the peak demand for ICU beds that will arise under various combinations of non-pharmaceutical interventions (Ferguson et al. 2020). The primary claim of the authors is that it will not suffice to adopt a “mitigation”-based approach to the pandemic; the policy adopted by the UK government until that point (Boseley 2020), under which, rather than interrupting transmission completely, interventions aim to reduce the health impact of an epidemic (slowing transmission without completely suppressing it), while allowing population immunity to build up through the spread of infection until herd immunity is reached. Rather, the authors argue, in order to prevent healthcare systems from being overwhelmed, we must adopt a “suppression” strategy, in which attempts are made to “reduce the reproduction number (the number of secondary cases each case generates), R , to below 1 and hence reduce case numbers to low levels or...eliminate human-to-human transmission” (Ferguson et al. 2020, 3). This report was indeed instrumental in the UK’s abrupt about-face concerning their pandemic strategy (Boseley 2020). Winsberg et al. have two primary problems with this model – first, that it was based on unreliable assumptions, and second, that it generated overly pessimistic forecasts.² Let’s turn first, then, to the claims about the faulty nature of the data that the model takes as a basis for its predictions.

Winsberg et al. note that in order to predict the impact of different policy choices on demand for ICU beds, the model required (among other data), “inputs for expected death rate, hospitalization rate, and ICU admittance rate for each 100 people infected”. However, “[a]t the beginning of the COVID epidemic, and even now as we write this sentence,³ these magnitudes were not well estimated” (2020, 223). Here, Winsberg et al. cite the WHO’s early estimates, which were very high, because, as they note, they are subject to selection bias: these estimates were derived from case rates in China, which disproportionately include people who require care, which means that “the resulting data are biased towards more severe results”. Thus, “early WHO

¹ And, they claim, the US. As the model focuses primarily on projections for the UK, we will focus on the UK here.

² Winsberg et al. also complain about the lack of robustness of the model, and that it “can generate significantly different estimates even with the same parameters inputted” (2020, 225). But the model appears to be sufficiently robust to vindicate Ferguson et al.’s main claims (Edeling et al. 2020); and it has been shown that the ICL results are exactly reproducible (*ibid.*).

³ On p. 226, they indicate that their paper, or at least part of the section to which the above quote belongs, was written on May 19.

estimates were extremely high, with fatality rates as high as 3.4%⁴ and hospitalization rates well into the double digit percentages. The correct numbers are still unknown, but early estimates are clearly too high” (2020, 223).

This claim is certainly correct, but it’s curious that Winsberg et al. cite the early WHO figures here, when the figures actually used by the ICL modelers are explicitly stated in their report. In fact, the modelers assume an infection fatality rate of 0.9%⁵ (Ferguson et al. 2020), well within the bounds of current consensus (see Phipps et al. 2020), even given the fact that survival rates have improved over the course of the pandemic (Dennis et al. forthcoming). The hospitalization rates are similarly assumed to be much lower than the WHO figures, at 4.4%. They present projections for infection transmission rates in the absence of any mitigation measures (R_0) from 2 – 2.6, with a baseline assumption that $R_0=2.4$, which was a conservative estimate at the time (Birch 2020), and well within the bounds of current consensus (Hilton and Keeling 2020; Thiede et al. 2020).

So, it appears that these key assumptions weren’t overly pessimistic. But what about the results rendered by the model? Winsberg et al. claim that:

...the model performed poorly at anticipating ICU demand, which was at the heart of the policy recommendations that emerged from the model. Recall that the ICL scientists recommended a policy of “maximum suppression” (Ferguson et al. 2020). This was the most draconian set of policies the group imagined. They anticipated that even maximum suppression would at first barely avoid overwhelming the UK’s existing ICU and ventilator capacity, and it would then require cycling the economy on and off until a vaccine was available. Despite less than maximum suppression, this did not occur (2020, 225)

Let’s first determine what constitutes (as Winsberg et al. refer to it) the “maximum suppression” policy advocated by the ICL team, and the extent to which it resembled the policies implemented in the UK. The ICL model considers various combinations of four policies: “case isolation” (symptomatic cases stay home for 7 days); “voluntary home quarantine” (all household members remain at home for 14 days following identification of a symptomatic case in the household); “social distancing” (all households reduce contact outside household, school or workplace by 75%) and “closure of schools and universities”. All of these measures in

⁴ The estimate mentioned by the WHO here is an early estimate of the case fatality rate rather than the infection fatality rate (which forms the basis of Ferguson et al.’s projections) – on March 3, in a media briefing, the WHO director-general stated: “Globally, about 3.4% of reported COVID-19 cases have died” (WHO 2020).

⁵ This estimate is taken from a paper by Verity et al. from March, which estimates an overall infection fatality rate of 0.66% in mainland China (2020). The numbers are then adjusted in light of the older population demographics in the UK (Ferguson et al. 2020).

combination, which stop short of a full lockdown, are the most severe scenario considered by the ICL team (Ferguson et al. 2020).

In fact, immediately following the release of this report on March 16, all the above-mentioned measures (with the exception of closing of schools and universities) were adopted, and, in addition, it was recommended that those over 70 do not leave the house (Boseley 2020). On March 20, UK schools were closed (with the exception of Northern Ireland, which closed schools to students on March 18, and to staff on March 23) (BBC 2020b). Universities were never forced to close entirely but most had suspended face-to-face teaching by March 17 (Stanton and Jack 2020) in addition to implementing other measures such as ending term early and cancelling exams (BBC 2020a). So, the measures implemented were not that far off the “maximum suppression” recommendations made in the ICL report. In light of this, we can now ask, how well did the projections match with reality? Did it turn out, as Winsberg et al. claim, that this model’s predictions for “ICU and ventilator demand were overly pessimistic” (2020, 226)?

This is a bit difficult to evaluate, because the ICL model presents a range of projections for different R_0 values, and for different thresholds (for new COVID cases diagnosed in ICUs within a week) at which general social distancing and school and university closures are assumed to kick in.⁶ But let’s take the projection which most resembles what happened in the UK. On March 20, as outlined above, both school closures and general social distancing were in force, and many universities had suspended much face-to-face activity. In the week leading up to March 20, ICU cases with diagnosed COVID-19 climbed by roughly 300 (ICNARC 2020). If we assume, then, that social distancing and school closures are triggered when new diagnosed ICU cases reach this number, and assume, in addition, that $R_0=2.4$ (Ferguson et al.’s baseline assumption), then demand for ICU beds is projected to peak at 4000.⁷ The baseline capacity in the UK was 4123 ventilated beds (Mateen et al. 2020)⁸. It seems fair, then, to say that this projection suggests that “even maximum suppression would at first barely avoid overwhelming the UK’s existing ICU and ventilator capacity” (Winsberg et al. 2020, 225). However, roughly

⁶ Case isolation and voluntary home quarantine, in all projected scenarios, are assumed to be implemented in late March.

⁷ In fact, we might expect actual projections to be lower than this, as case isolation and voluntary home quarantine were implemented earlier than anticipated in the model, on March 16, along with general social distancing and voluntary isolation of the elderly. Most universities had also suspended much activity by March 17.

⁸ Although this does not take surge capacity into account, see Ferguson et al. 2020, Mateen et al. 2020.

3900 ICU beds were occupied in the UK's initial peak in April. The ICL team's projection, then, seems to be remarkably accurate, given the parameters considered.^{9 10}

Winsberg et al. further argue that “[a] strong indictment of the ICL model comes from examining what it would have predicted for Sweden, which has not implemented any lockdowns” (2020, 226). They take Gardner et al. (2020), who “ran a model very closely based on the ICL model” (ibid.), as evidence that the ICL model must vastly overestimate deaths, because Gardner et al.'s estimates for deaths in Sweden, even under the strictest suppression scenario that they consider, were much higher than the actual deaths under the public health strategy in effect by the time – as Winsberg et al. note, they predicted over 15,000 deaths in Sweden by the end of April. This would be a rather strange result, given that the ICL's most pessimistic scenario with “maximum suppression” in the UK ($R_0=2.6$, school closures and general social distancing triggered when weekly ICU cases hit 400 and cycled off when they hit 100), with a densely packed population of around 68 million compared to Sweden's 10 million, is a total of 48,000 deaths over a two-year period. What explains this apparent mismatch is that, while it is true that Gardner et al. adopt some of the ICL team's parameter assumptions, e.g. the percentage of hospitalized COVID-19 cases that require ICU, they use a different model than the ICL team, whose code is available online (see <https://github.com/kassonlab/covid19-epi>). The fact that Gardener's model predicts over 15,000 deaths for Sweden by the end of April under a strict suppression scenario thus does not entail any indictment of the ICL model.

Of course, our rebuttal of their critique of the ICL model – the main object of Winsberg et al.'s contentions – does not mean that *all* models that were used to inform policy making during the pandemic should have been used for this purpose. For instance, the statistical analyses of the Institute for Health Metrics and Evaluation (IHME) have been criticized because, unlike the ICL model, they lack an epidemiological basis (see Jewell et al. 2020), and they have been shown to be predictively weak, calling into question whether they should be used for policy making (Marchant et al. 2020). If such models had been the only ones available when many governments decided to impose lockdowns, or if a government had so decided only on the basis of such models despite the availability of others, it might well be argued that restricting citizens'

⁹ We could, of course, take different parameters as our basis – we have noted that Winsberg et al. do not identify the specific projection upon which they base their claims, but they do include a graph from the Ferguson paper which displays a more optimistic scenario; here, the virus is assumed to be slightly less infectious, with an R_0 value of 2.2, and social distancing and school/university closures implemented when new ICU cases climb by 100 within a week. The projection here for peak demand is 1600 beds; 39% of baseline capacity.

¹⁰ It might be that Winsberg et al. misinterpret the ICL projections as being too pessimistic because their presentation is mainly based on a blog post (Lemoine 2020), which also misinterprets these projections as being too pessimistic (see Winsberg et al. 2020, 237 footnote 3).

liberties was not justifiable on these grounds. But because Winsberg et al. fail to make a convincing case against the ICL model, they fail to support their general claim that “the models and data used in support of lockdowns were poor” (2020, 236) and thus their generalization that “governments have systematically failed to meet their epistemic obligations in this crisis” (*ibid.*, 237) does not follow.

3. PROBLEMS WITH MODELERS?

After arguing that the models were flawed in ways that led to overstatements concerning death rates, ICU demand, and so on, Winsberg et al. focus on the modelers, providing a socio-psychological explanation for why they allegedly generated these biased models. They contend that epidemiologists are influenced by values and external pressure to emphasize the risks of infectious diseases, rather than other ethical risks, such as harmful consequences of measures, in particular lockdowns, to fight these diseases: “The consequences to themselves, their careers, their discipline, their own sense of moral culpability will be much larger if they underpredict rather than overpredict death by disease” (2020, 230). We argued in the previous section that the ICL modeling that Winsberg et al. focus on did not in fact rely on overly pessimistic assumptions or generate overly pessimistic projections. This finding, if anything, disconfirms their contention that epidemiologists are systematically biased towards overstating the risks of infectious diseases. But they also seek to provide independent evidence for this claim, by referring to inflated estimates of infection fatality rates during previous epidemics, which they relate to epidemiologists purportedly being influenced by values and external pressures:

Given these influences, it is unsurprising to find a great deal of evidence from past experiences that epidemiologists favor a balance of inductive risks that leads to over-forecasting the severity of diseases. The infection fatality rate of Mad Cow Disease, H1N1, H5N1, H7N9, and MERS all were considerably lower than what epidemiologists predicted. And while SARS 2002 actually ended up being twice as fatal as originally predicted, its infectious spread was tiny compared to what they predicted (Yu et al. 2013; Wang et al. 2012; Lipsitch et al. 2015; Cauchemez et al. 2014). (2020, 230-1)

The cited papers suggest that there has been a tendency to over-forecast case fatality rates¹¹ at early stages of epidemics. Yet these papers provide an alternative explanation for this tendency to Winsberg et al.’s contention that it is due to socio-psychological factors. Namely, they identify

¹¹ We refer here to case fatality because their cited papers predominantly do. According to Lipsitch et al., infection fatality rate “defines a case as a person who has shown evidence of infection, either by clinical detection of the pathogen or by seroconversion or other immune responses” (2015, 2). These cases may be symptomatic or asymptomatic, whereas the symptomatic case fatality rate includes only symptomatic cases. Because asymptomatic cases may go undetected, it can be expected that infection fatality rates are at early stages of an epidemic particularly difficult to gauge.

medical factors that can affect the estimation of case fatality rates; in particular, for many diseases, positive cases that are detected first are typically those that present the most severe symptoms (e.g. cases with pneumonia and requiring ventilation in the case of COVID-19), thus leading to higher fatality rates among detected cases than among all cases (Lipsitch et al. 2015). This systematic detection bias explains early inflated case fatality estimates for H1N1, H5N1, H7N9 and MERS, according to the papers cited by Winsberg et al.¹² These papers also show for the above epidemics that, once more was learned about an epidemic (such as true infection rates), inflated case fatality estimates were corrected downwards. Furthermore, Lipsitch et al. (2015) seek to provide means for identifying and reducing biases, even in early stages of epidemics.

While, as we've seen, Winsberg et al. note that case fatality rates may be inflated due to selection bias (2020, 10), they suggest that epidemiologists' estimates are additionally influenced by values and external pressures. Of course, the fact that there are plausible medical explanations for inflated case fatality estimates in early stages of epidemics can plausibly be explained by selection bias does not imply that Winsberg et al.'s socio-psychological account is necessarily false. However, it should be pointed out here that they fail to provide evidence for their account, as the literature they cite provides an alternative, "internal" (medical) explanation for this phenomenon, and they do not cite evidence that epidemiologists' estimates are in addition biased through socio-psychological factors. It might be noted, furthermore, that an influential view in the philosophy of science holds that if an episode from the history of science can be explained internally, such an explanation is preferable to a socio-psychological account of the same episode (e.g. Lakatos 1970). Because there are plausible medical explanations of some overestimates of case fatality rates, Winsberg et al.'s socio-psychological account is thus not convincing on this view, and in any case not substantiated by the literature they cite.

Winsberg et al. then go on to criticize that "[r]epeated cases of overprediction can even be diagnosed in single individuals" (2020, 231), claiming that ICL epidemiologist Neil Ferguson¹³ "has often overestimated disease dangers" (ibid.) and taking Mad Cow Disease and bird flu as evidence for this claim. According to Winsberg et al., Ferguson's group predicted that Mad Cow Disease would kill around 136,000 people, while the actual number of deaths equals less than

¹² In their quote, Winsberg et al. also mention Mad Cow Disease and SARS. Mad Cow Disease is not mentioned in the cited papers and could not have been subject to inflated case fatality estimates as it is generally associated with fatal disease, i.e. its case fatality rate equals one (<https://www.cdc.gov/prions/vcjd/about.html>). Case fatality estimates for SARS were subject to another bias that is due to medical factors, which is caused by delayed reporting of deaths and decreases case fatality estimates (Lipsitch et al. 2015).

¹³ And the first author of the report on the ICL model which forms the focus of the previous section.

200. They cite a 2001 *New York Times* article, which, referring to Ferguson, states that “his group published estimates a year ago predicting that the number of variant C.J.D. cases [the disease caused by the agent responsible for Mad Cow Disease] *might* reach 136,000 in coming decades” (Blakeslee 2001, emphasis added). This quote most likely refers to Ghani et al. (2000), a paper that was co-authored by Neil Ferguson, which concluded that, under a wide range of assumptions concerning the incubation period of the disease, the infectivity of cattle, and efficacy of measures to reduce human exposure to infected material, “the current mortality data are consistent with between 63 and 136,000 cases” (2000, 583). Thus, contrary to Winsberg et al.’s contentions, Ferguson’s group did not predict that there would be 136,000 cases; rather, the number amounts to the upper bound of the large interval that is consistent with the mortality data that was available at the time of Ghani et al.’s writing. Concerning bird flu, Winsberg et al. claim that according to a 2005 article that appeared in *The Guardian*, Ferguson “told the BBC that the deaths from bird flu could be between 5,000,000 and 150,000,000; the actual number was around 300 (Sturcke 2005)” (2020, 231). But this claim is not true; according to the cited *Guardian* article (Sturcke 2005), it was David Nabarro, the UN official who was by the time in charge of coordinating the worldwide response to an outbreak, who “told the BBC that the ‘range of deaths could be anything between five and 150 million’” (Sturcke 2005).^{14 15} We conclude here that Winsberg et al.’s claims that epidemiologists are systematically prone to generate biased estimates due to socio-psychological factors (values and external pressures), and that this bias can be diagnosed in Ferguson’s predictions are not underwritten by the literature that they cite. We proceed now to Winsberg et al.’s claims about lockdowns.

4. ARE LOCKDOWNS EFFECTIVE?

Winsberg et al. not only argue that the models used to forecast the health consequences of the pandemic were flawed; they also call into question the effectiveness of lockdowns as a means of containing its spread. They claim that

a literature search reveals there are no published, peer-reviewed papers demonstrating the effectiveness of universal lockdown procedures to combat any epidemic...we lack empirical evidence that extensive lockdown policies or maximal suppression work *at*

¹⁴ It is not clear whether the quote refers to a range of 5 to 15,000,000 deaths, or (as Winsberg et al. claim) to a range of 5,000,000 to 15,000,000 deaths. This is however irrelevant as the quote is in any case not Ferguson’s, contrary to Winsberg et al.’s assertion.

¹⁵ Ferguson is quoted in the same article, relating bird flu to the 1918 pandemic and suggesting that if bird flu were comparably fatal the number of deaths might for today’s world population scale up to 200 million deaths. This statement should again not be read as a prediction; Ferguson likely made this statement to caution against a public health strategy that is based merely on mitigation and treatment, because earlier in the same year, Ferguson’s group had published a paper that shows, based on mathematical modeling, that emergent influenza pandemics can be eliminated through a different public health strategy that is based on antiviral prophylaxis and social distancing (Ferguson et al. 2005).

all, never mind that they are superior to other, less draconian practices. (2020, 228; emphasis in original)

Our literature search reveals that there is, by now, a growing body of working and peer-reviewed papers providing evidence for the effectiveness of lockdowns in reducing the spread of COVID-19; the following is an incomplete list: Amuedo-Dorantes (2020); Brauner et al. (forthcoming); Bonardi et al. (2020); Dave et al. (2021); Dehning et al. (2020); Fang et al. (2020); Flaxman et al. (2020); Hsiang et al. (2020); Huber and Langen et al. (2020); Juranek and Zoutman (2020); Qiu et al. (2020). We found one study that purports to show that lockdowns were ineffective (Homburg 2020), but it is subject to grave methodological errors (see Robra and Felder 2020). We grant that not all of these studies confirming the effectiveness of lockdowns may have been available by the time Winsberg et al. wrote their paper, and not all of them have been peer-reviewed even now. While peer-review might not be the right metric for evaluating studies in a fast-moving pandemic like COVID-19, it should nevertheless be noted that there were already peer-reviewed studies available by early May 2020 (e.g. Qiu et al. 2020), as well as papers available on selective working-paper servers (e.g. Fang et al. 2020), confirming the effectiveness of lockdowns as a means of decreasing transmission rates. Winsberg et al.'s claim that there was lack of empirical evidence that lockdowns "work at all" is simply false.

Winsberg et al. then briefly consider and discard a paper that assesses the March/April state-wide lockdown in California, finding that it reduced COVID-19 cases and deaths due to COVID-19 considerably in this period:

The best paper we can find defending lockdowns is a working paper by Friedson et al. (2020), but this paper has significant limitations. In particular, it counts drops in deaths five days after California's closing as evidence that lockdowns work. Since the virus takes longer than that to incubate, this drop could not have been caused by the lockdowns. (2020, 228)

To evaluate this argument, let's briefly consider the methodology used by Friedson et al. (2020). They use a synthetic control model approach, in which the development of COVID-19 in California is compared to its counterfactual development in "synthetic California". This is a model, consisting of a weighted linear combination of different US states, that is similar to California in terms of relevant characteristics (such as population density and COVID-related policies, e.g. travel restrictions or numbers of tests conducted). However, while a state-wide lockdown was imposed in California on March 19, in synthetic California no lockdown is imposed on this day. This is achieved by only using combinations of states for synthetic

California which imposed lockdowns at least 5 days after March 19, if at all.¹⁶ By comparing case numbers in California and synthetic California after March 19, the effect of the lockdown in reducing cases and averting deaths can be estimated. By investigating various models of synthetic California, made up of different linear combinations of states, Friedson et al. aim to achieve robust estimates of the net effect of the lockdown on COVID-19 cases in California. They find that the lockdown reduced the number of cases in California by 125.5 to 219.7 per 100,000 population in the investigated period, and they suggest that it prevented up to 1,661 COVID-19 related deaths in the same period.

In the cited passage, Winsberg et al. claim that Friedson et al. count drops in the Californian death rate in illegitimate ways – five days after the implementation of the lockdown, even though the incubation period is longer than that. It should first be noted that, even if this is true, this would not affect their estimate of the reduction of case numbers and it would thus not on its own establish that they fail to show that “lockdowns work” – the main aim of a lockdown is to slow the spread of infection, thus preventing healthcare services from becoming overwhelmed. Yet Winsberg et al. also fail to establish that Friedson et al. count drops in the death rate in illegitimate ways. It is true that in their first model, a gap in mortality evolves between California and synthetic California six (6) days after imposing the lockdown, and exponentially widens thereafter (see 2020, 26 and Figure 12a through 12c); but they independently also take a more conservative approach, in which they force the mortality rates to be equal in California and synthetic California for 13 days after the imposition of the lockdown (that is, for the estimated median incubation length plus the estimated median time from symptom onset to possible death) (see 2020, 27 and Appendix Figures 8a through 8c).¹⁷

Winsberg et al. nevertheless seem to believe that it is a significant limitation of the study that Friedson et al. also include the less conservative approach in their paper, in which death rates drop before the lockdown could have caused them to drop. However, this does not invalidate Friedson et al.’s methodology, because the causal mechanisms that may have led to the decline in the death rate may be complex; for instance, people might have changed their behavior, by increasing social distancing, even prior to the state-wide lockdown, which seems plausible as, for instance, a state of emergency was declared two weeks prior to the state-wide lockdown (Office of Governor Gavin Newsom 2020). Thus, for a convincing critique of

¹⁶ It should be noted that, by including states that imposed lockdowns on 24 March or later, rather than only states that did not impose lockdowns at all during the period investigated, Friedson et al. generate conservative estimates.

¹⁷ Their estimates of the median incubation period and median period from symptom onset to near-death follow Lauer et al. (2020) and Wang et al. (2020), respectively.

Friedson et al.'s methodology, Winsberg et al. would have to show – which they don't – that lockdowns had no *additional* effect on the death rate declining.¹⁸

Summing up, Winsberg et al. fail to acknowledge literature confirming the effectiveness of lockdowns as a means of combating the pandemic, which was available even by the time of their writing. Furthermore, concerning the study that they do consider and criticize as significantly limited, by selectively choosing only one of Friedson et al.'s models as the subject of their critique, and by failing to make a convincing case against this model, Winsberg et al. do not invalidate Friedson et al.'s findings.

5. DO LOCKDOWNS CAUSE MORE HARM THAN GOOD?

But even if we are to assume, to a certain extent, that lockdowns do indeed prevent harm, Winsberg et al. marshal another argument against them – that they do more harm than good, at least in the context of the United States. A central claim here is the following:

Deaths connected to layoffs that are the result of COVID-19 might already be in the same ballpark as the number of deaths caused by the virus itself (Cordle 2020) (2020, 234)

Let's delve into this claim in detail. Here, Winsberg et al. cite a short paper that appeared on LinkedIn by Vaughn Cordle (2020). In Cordle's analysis (from April 2020), he suggests that lockdown-related layoffs will lead to a jump in suicides and drug overdoses. He provides several estimates of the layoff-related death toll in the U.S., based on how long a lockdown will continue. It's difficult to gauge exactly what he means by lockdown, as the U.S. always had piecemeal lockdown provisions in place that differed from state to state and county to county, but as his analysis here is based on unemployment rates, this is not central to his contentions.

Cordle suggests, in the grimmest scenario he considers, that if the lockdown extends through May, there will be a 31% jump in unemployment (47 million people unemployed). He estimates that this will lead to “a doubling of drug overdoses (69,735) and an additional 15,137

¹⁸ It is instructive to compare this criticism to the argument Stefan Homburg makes in his above-mentioned paper, which has been criticized for containing methodological errors. Homburg argues that the COVID-19 death rates in Italy and some other countries imposing lockdowns started falling before the lockdowns could have caused them to fall, and he concludes that “the lockdown had no visible impact on fatal outcomes” (2020, 5). In their critical response, Bernt-Peter Robra and Stefan Felder (2020) point out that Homburg's conclusion does not follow because the prospect of a lockdown could have led to changes in individual behavior, even before the lockdown was implemented, and Homburg fails to show that the lockdowns had no additional effect on death rates. Winsberg et al. commit a similar fallacy as Homburg (2020); taking declines in death rates that occur “too early” as evidence that lockdowns were not effective oversimplifies the possible causal mechanisms that may have led to these declines.

suicides” (2020), thus, we should expect an estimated total of 84,872 layoff-related deaths.¹⁹ This, Cordle notes, is 63% of the projected COVID-related deaths by the IHME by August 4 2020 (2020).²⁰ The first thing to note about this claim, in light of what Winsberg et al. make of it, is that even this most drastic scenario (which Cordle does not take to be the most realistic scenario) does not suggest that layoff related deaths *will* be in the same ballpark as COVID-related deaths, let alone that they were at the time that the article was likely submitted (May 2020). The second is that unemployment rates never reached anything near these levels. Although, at the time that Cordle wrote, as unemployment reached an alarming peak of 23.1 million, it was perhaps justifiable to expect that it might spiral further, it was clear by May that unemployment was nowhere near these rates.

It should also be noted that Cordle’s actual estimate at the time of writing, given the expectation of the easing of measures and a resultant economic recovery from mid-May, was that 15 million people would be unemployed by the end of 2020, which he projects will lead to 41,067 layoff related deaths (31% of the IHME’s projections for the COVID-related deaths as of the 4th of August) (2020). With the benefit of hindsight, we can also see that these projections may have turned out to be overblown – it’s too early to say what the lockdown meant for suicide rates, but early indications are promising (John 2020; Wilson 2020). Deaths from drug overdoses, which have been increasing in the US since 2019, did appear to accelerate during March to May 2020, but not to the degree predicted (CDC 2020a).

As Winsberg et al. claim, there are other factors that could lead to an increase in lockdown related deaths – the long term effects of unemployment, and untreated illnesses such as cancer, to name a few. But there are also positive externalities as the result of lockdowns – premature deaths, for example, were avoided by a reduction in pollution levels (Venter et al. 2020). The E.U. saw a 36% reduction of motor vehicle fatalities in April 2020 (compared to April 2019) (ETSC 2020) and the total number of motor vehicle fatalities in the U.S. from March to May appears to have declined as well (NHTSA). Of course, we might question the idea that such positive externalities can play a role in justifying lockdown, but this is relevant to the claim that lockdown might cause more deaths than the virus.

¹⁹ It’s not completely clear when we should expect these numbers to eventualize, but as he compares them to the August 4 projections for COVID-related deaths, presumably by that date.

²⁰ As an aside, this estimate of COVID-related deaths turned out to be a bit conservative, if we compare it with the CDC’s figures for the 4th of August 2020 – they recorded a total of 156,311 cumulative deaths in the US by that date (2020b). It should be noted, however, that Cordle claims that the CDC, at least as of April, overestimates death rates from COVID-19, but as he doesn’t give much of an indication of why and by how much, it’s difficult to take this into account here.

6. CONCLUSION

Contrary to the contentions of Winsberg et al., it seems that the modelling that steered policy interventions in the UK towards restrictions made quite accurate projections, and they fail to underpin their claim that experts were systematically biased towards overprediction. There was indeed evidence that lockdowns can provide an effective means of reducing virus transmission, and the claims they make about the degree of harm caused by lockdowns are vastly overblown. In showing that many of the most fundamental assumptions made in their case against lockdowns are false, we have not touched on their other contentions concerning the degree and standard of evidence required in order to implement such restrictive policies in a liberal democracy. But in showing that the evidence in favor of the efficacy of lockdowns is far better than they suggest, it might be argued that the available evidence comes a lot closer to even the high standard that they require. And of course, the question of whether their proposed epistemic standard is one that must be met in order to justify such interventions is open to debate.

REFERENCES

- Amuedo-Dorantes, Catalina, Cristina Borra, Noelia Rivera-Garrido, and Almudena Sevilla. 2020. "Timing is Everything when Fighting a Pandemic: COVID-19 Mortality in Spain." Institute of Labor Economics (IZA) Discussion Paper No. 13316. <http://ftp.iza.org/dp13316.pdf>
- Bates, Amanda E., Richard B. Primack, Paula Moraga, and Carlos M. Duarte. 2020. "COVID-19 pandemic and associated lockdown as a "Global Human Confinement Experiment" to investigate biodiversity conservation." *Biological Conservation* 248: 108665. <https://doi.org/10.1016/j.biocon.2020.108665>
- BBC. 2020a. "Coronavirus: More Universities Halt Teaching and Exams." March 13. Accessed December 30, 2020. <https://www.bbc.com/news/education-51880355>
- BBC. 2020b. "Coronavirus: UK Schools, Colleges and Nurseries to Close from Friday." March 18. Accessed December 30, 2020. <https://www.bbc.com/news/uk-51952314>
- Birch, Jonathan. 2020 "Science and Policy in Extremis: What We Can Learn from the UK's initial Response to COVID-19?" *London School of Economics and Political Science Blog*, December 22. Accessed December 30, 2020. <https://www.lse.ac.uk/philosophy/blog/2020/12/22/science-and-policy-in-extremis/>
- Blakeslee, Sandra. 2001. "Estimates of Future Human Death Toll from Mad Cow Disease Vary Widely." *The New York Times*, October 30. <https://www.nytimes.com/2001/10/30/health/estimates-of-future-human-death-toll-from-mad-cow-disease-vary-widely.html>

Bonardi, Jean-Philippe, Quentin Gallea, Dimitrija Kalanoski, and Rafael Lalive. 2020. “Fast and local: how did lockdown policies affect the spread and severity of COVID-19?” *CEPR COVID Economics* 23. <https://cepr.org/sites/default/files/news/CovidEconomics23.pdf>

Boseley, Sarah. 2020. “New Data, New Policy: Why UK’s Coronavirus Strategy Changed.” *The Guardian*, March 16. Accessed December 30, 2020. <https://www.theguardian.com/world/2020/mar/16/new-data-new-policy-why-uks-coronavirus-strategy-has-changed>

Brauner, Jan M., Sören Mindermann, Mrinank Sharma, et al. forthcoming. “Inferring the effectiveness of government interventions against COVID-19.” *Science*. DOI: 10.1126/science.abd9338

Cauchemez, Simon, Christophe Fraser, Maria D. Van Kerkhove, et al. 2014. “Middle East respiratory syndrome coronavirus: quantification of the extend of the epidemic, surveillance biases, and transmissibility.” *Lancet Infectious Diseases* 14: 50-56.

Center for Disease Control (CDC). 2020a. “Increase in Fatal Drug Overdoses Across the United States Driven by Synthetic Opioids Before and During the COVID-19 Pandemic.” CDCHAN-00438. Accessed December 30, 2020. <https://emergency.cdc.gov/han/2020/han00438.asp>

Center for Disease Control (CDC). 2020b. “Trends in Number of COVID-19 Cases and Deaths in the US Reported to CDC, by State/Territory.” Accessed December 30, 2020. https://covid.cdc.gov/covid-data-tracker/#trends_totalandratedeaths

Cordle, Vaughn. 2020. “Pandemic 2020: Layoff-related Deaths.” April 23. Accessed December 24, 2020. <https://www.linkedin.com/pulse/pandemic-2020-layoff-related-deaths-increase-covid-19-cordle-fa/?fbclid=IwAR18TeBQA3Lna61kdMPzPcWyrBnj70n6EP9qF7sFIL10JZIPdVeX9Mnt6Ok>

Dave, Dhaval, Andrew I. Friedson, Kyutaro Matsuzawa, and Joseph J. Sabia. 2021. “When do Shelter-in-Place Orders Fight COVID-19 Best? Policy Heterogeneity Across States and Adoption Time.” *Economic Inquiry* 59(1):29-52. doi:10.1111/ecin.12944

Dehning, Jonas, Johannes Zierenburg, F. Paul Spitzner, et al. 2020. “Inferring change points in the spread of COVID-19 reveals the effectiveness of interventions.” *Science* 369 (6500): eabb9789. DOI: 10.1126/science.abb9789

Dennis, John M., Andrew P. McGovern, Sebastian A. Vollmer, and Bilal A. Mateen. forthcoming. “Improving Survival of Critical Care Patients With Coronavirus Disease 2019 in England: A National Cohort Study, March to June 2020.” *Critical Care Medicine*. DOI: 10.1097/CCM.0000000000004747

Edeling, Wouter, Hamid Arabnejad, Robert C. Sinclair, et al. 2020. Model uncertainty and decision making: Predicting the Impact of COVID-19 Using the CovidSim Epidemiological Code. Accessed December 31, 2021. https://assets.researchsquare.com/files/rs-82122/v3_stamped.pdf

European Transport Safety Council (ETSC). 2020. “The Impact of COVID-19 Lockdowns on Road Deaths in April 2020.” July. Accessed December 30, 2021. https://etsc.eu/wp-content/uploads/PIN-Corona-Briefing_final.pdf

Fang, Hanming, Long Wang, and Yang Yang. 2020. “Human Mobility Restrictions and the Spread of the Novel Coronavirus (2019-NCOV) in China.” *NBER Working Paper* 26906. <http://www.nber.org/papers/w26906>

Ferguson, Neil M., Derek A. T. Cummings, Simon Cauchemez, et al. 2005. Strategies for containing an emerging influenza pandemic in Southeast Asia. *Nature* 437(8):209-14. doi:10.1038/nature04017

Ferguson, Neil M., Daniel Laydon, Gemma Nedjati-Gilani, et al. 2020. “Report 9: Impact of Non-Pharmaceutical Interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand.” Imperial College COVID Response Team. March 16. DOI: 10.25561/77482

Flaxman, Seth, Swapnil Mishra, Axel Gandy, et al. 2020. “Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe.” *Nature* 584: 257–261. DOI: 10.1038/s41586-020-2405-7

Friedson, Andrew I., Drew McNichols, Joseph J. Sabia, and Dhaval Dave. 2020. “Did California’s Shelter-In-Place Order Work? Early Coronavirus-Related Public Health Effects.” *NBER Working Paper* 26992. <http://www.nber.org/papers/w26992>

Gardner, Jasmine M., Lander Willem, Wouter Van Der Wijngaart, et al. 2020. “Intervention strategies against COVID-19 and their estimated impact on Swedish healthcare capacity.” *medRxiv*, <https://doi.org/10.1101/2020.04.11.20062133>

Ghani, Azra C., Neil M. Ferguson, Christl A. Donnelly, and Roy M. Anderson. 2000. “Predicted vCJD mortality in Great Britain.” *Nature* 406:583-4.

Hilton, Joe, and Matt J. Keeling. 2020. “Estimation of Country-Level Basic Reproductive Rates for Novel Coronavirus (SARS-CoV-2/COVID-19) Using Synthetic Contact Matrices.” *PLoS Computational Biology* 16 (7): e1008031. DOI: 10.1371/journal.pcbi.1008031

Homburg, Stefan. 2020. “Effectiveness of Corona Lockdowns: Evidence for a Number of Countries.” *Economists’ Voice* 17 (1): 20200010. DOI: 10.1515/ev-2020-0010

Hsiang, Solomon, Daniel Allen, Sébastien Annan-Phan, et al. 2020. “The effect of large-scale anti-contagion policies on the COVID-19 pandemic.” *Nature* 584: 262-267. DOI: 10.1038/s41586-020-2404-8

Huber, Martin, and Henrika Langen. 2020. “The Impact of Response Measures on COVID-19-Related Hospitalization and Death Rates in Germany and Switzerland.” *arXiv preprint*. DOI: arXiv:2005.11278v3

IHME (2020). Forecasting the Impact of the First Wave of the COVID-19 Pandemic on Hospital Demand and Deaths for the USA and European Economic Area Countries. *medRxiv*. <https://doi.org/10.1101/2020.04.21.20074732>

Intensive Care National Audit and Research Centre (ICNARC). 2020. “ICNARC Report on COVID-19 in Critical Care: 27 March 2020.” *ICNARC COVID-19 Study Case Mix Programme Database*. Accessed December 30, 2020. www.icnarc.org/DataServices/Attachments/Download/b5f59585-5870-ea11-9124-00505601089b

Jewell, Nicholas P., Joseph A. Lewnard, and Britta L. Jewell. 2020. "Caution Warranted: Using the Institute for Health Metrics and Evaluation Model for Predicting the Course of the COVID-19 Pandemic." *Annals of Internal Medicine*. DOI: 10.7326/M20-1565

John, Ann, Jane Pirkis, David Gunnell, Louis Appleby, and Jacqui Morrissey. 2020. "Trends in suicide during the COVID-19 pandemic." *British Medical Journal* 371: m4352.

Juranek, Steffen, and Floris T. Zoutman. 2020. "The Effect of Social Distancing Measures on Intensive Care Occupancy: Evidence on COVID-19 in Scandinavia." FOR Discussion Paper 2/20, NHH Norwegian School of Economics.

Lakatos, Imre. 1970. "History of Science and Its Rational Reconstructions." In R.C. Buck and R.S. Cohen (eds.), *PSA 1970: Boston Studies in the Philosophy of Science*, 8, Dordrecht: Reidel, pp. 91–135.

Lauer, Stephen A., Kyra H. Grantz, Qifang Bi, et al. 2020. "The Incubation Period of Coronavirus Disease 2019 (COVID-19) From Publicly Reported Confirmed Cases: Estimation and Application." *Annals of Internal Medicine* 172:577-582. <https://doi.org/10.7326/M20-0504>

Lemoine, Philippe. 2020. "Are We Headed toward an Unprecedented Public Health Disaster?" *Nec Pluribus Impar*. March 21. Accessed January 5, 2021. <https://necpluribusimpar.net/are-we-headed-toward-an-unprecedented-public-health-disaster/>

Lipsitch, Marc, Christl A. Donnelly, Christophe Fraser, et al. 2015. "Potential Biases in Estimating Absolute and Relative Case-Fatality Risks during Outbreaks." *PLoS Neglected Tropical Diseases* 9 (7): e0003846. DOI: 10.1371/journal.pntd.0003846

Mandel, Antoine, and Vipin P. Veetil. 2020. "The Economic Cost of COVID Lockdowns: An Out-of-Equilibrium Analysis." *Economics of Disasters and Climate Change* 4:431-451. <https://doi.org/10.1007/s41885-020-00066-z>

Marchant, Roman, Noelle I. Samia, Ori Rosen, et al. 2020. "Learning as We Go – An Examination of the Statistical Accuracy of COVID-19 Daily Death Count Predictions." *arXiv*:2004.04734.

Mateen, Bilal A., Harrison Wilde, John M. Dennis, et al. 2020. A geotemporal survey of hospital bed saturation across England during the first wave of the COVID-19 Pandemic. *medRxiv*. <https://doi.org/10.1101/2020.06.24.20139048>

National Highway Traffic Safety Administration (NHTSA). 2020. "Early Estimate of Motor Vehicle Traffic Fatalities for the First Half (Jan-Jun) of 2020." October. Accessed December 30, 2020. <https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/813004>

Office of Governor Gavin Newsom. 2020. "Governor Newsom Declares State of Emergency to Help State Prepare for Broader Spread of COVID-19." 4 March. Accessed December 26, 2020. <https://www.gov.ca.gov/2020/03/04/governor-newsom-declares-state-of-emergency-to-help-state-prepare-for-broader-spread-of-covid-19/>

Phipps, Steven, R. Quentin Grafton, and Tom Kompas. 2020. "Robust Estimates of the True (Population) Infection Rate for COVID-19: A Backcasting Approach." *Royal Society Open Science* 7: 200909. DOI: 0.1098/rsos.200909

- Qiu, Yun, Xi Chen, and Wei Shi. (2020). "Impacts of social and economic factors on the transmission of coronavirus disease 2019 (COVID-19) in China." *Journal of Population Economics* 33: 1127-1172. DOI: 10.1007/s00148-020-00778-2
- Robra, Bernt-Peter, and Stefan Felder. 2020. "Homburg's Lockdown Analysis: Conclusions without Data and an Appropriate Estimation Model." *Economists' Voice* 17 (1): 20200015. DOI: 10.1515/ev-2020-0015
- Stanton, Bethan, and Andrew Jack. 2020. "UK Universities Suspend Face-to-Face Teaching." *Financial Times*, March 17. Accessed December 30, 2020.
<https://www.ft.com/content/f325ed7e-6862-11ea-800d-da70cff6e4d3>
- Sturcke, James. 2005. "Bird Flu Pandemic 'Could Kill 150m'." *The Guardian*, September 30.
<https://www.theguardian.com/world/2005/sep/30/birdflu.jamessturcke>
- Thiede, Renate, Nada Abdelatif, Inger Fabris-Rotelli, et al. 2020. "Spatial variation in the basic reproduction number of COVID-19: A systematic review". *arXiv preprint*. DOI: arXiv:2012.06301v1
- Venter, Zander, Kristin Aunan, Sourangsu Chowdhury, and Jos Leliveld. 2020. "COVID-19 lockdowns cause global air pollution declines with implications for public health risk." *Proceedings of the National Academy of Sciences of the United States of America* 117 (32): 18984-18990.
- Verity, Robert, Lucy C. Okell, Ilaria Dorigatti et al. 2020. "Estimates of the severity of COVID-19 disease." *medRxiv preprint*. DOI: 10.1101/2020.03.09.20033357
- Wang, Dawei, Bo Hu, Chang Hu, et al. 2020. "Clinical Characteristics of 138 Hospitalized Patients With 2019 Novel Coronavirus-Infected Pneumonia in Wuhan, China." *Journal of the American Medical Association* 323 (11): 1061-1069. DOI: 10.1001/jama.2020.1585
- Wang, Taia T., Michael K. Parides, and Peter Palese. 2012. "Seroevidence for H5N1 Influenza Infections in Humans: Meta-Analysis." *Science* 335(6075): 1463. DOI:10.1126/science.1218888
- Wilson, Clare. 2020. "We need to be more careful when talking about suicide and the pandemic." *New Scientist*, November 13. Accessed December 30, 2020.
<https://www.newscientist.com/article/2259889-we-need-to-be-more-careful-when-talking-about-suicide-and-the-pandemic/>
- Winsberg, Eric, Jason Brennan, and Chris W. Surprenant. 2020. "How Government Leaders Violated Their Epistemic Duties during the SARS-CoV-2 Crisis." *Kennedy Institute of Ethics Journal* 30 (3-4): 215-242.
- World Health Organization (WHO). 2020. "WHO Director-General's opening remarks at the media briefing on COVID-19 - 3 March 2020." Accessed January 4 2021.
<https://www.who.int/director-general/speeches/detail/who-director-general-s-opening-remarks-at-the-media-briefing-on-covid-19---3-march-2020>
- Yu, Hongjie, Benjamin J. Cowling, Luzhao Feng, et al. 2013. "Human infection with avian influenza A H7N9 virus: an assessment of clinical severity." *The Lancet* 382: 138-45.