

Draft paper, 4 March 2021

**Science and policy in extremis:
the UK's initial response to COVID-19**

Jonathan Birch

Department of Philosophy, Logic and Scientific Method,
London School of Economics and Political Science,
Houghton Street, London, WC2A 2AE, UK.

j.birch2@lse.ac.uk

<http://personal.lse.ac.uk/birchj1>

Abstract

Drawing on the SAGE minutes and other documents, I consider the wider lessons for norms of scientific advising that can be learned from the UK's initial response to coronavirus in the period January-March 2020. I highlight three key issues: the normative force of scientific advice, the role of reasonable worst-case scenarios, and the limits of independence and neutrality. A recurring theme is the difference between normal scientific advising and scientific advising in extremis, when a significant fraction of a country's population face immediate peril and there is no adequate pre-existing plan for managing the risk. Through reflection on these issues, I formulate seven proposals for effective scientific advising in extremis. In an epilogue, I consider what the implications of these proposals would have been for the UK's response to the "second wave" of late 2020.

1. Table studies

The COVID-19 pandemic has presented governments around the world with extraordinarily difficult decisions. Governments have generally sought to base these decisions on scientific advice, but there are many ways in which a decision can be “based on scientific advice”. My focus here is on the United Kingdom, and on the UK's initial response to COVID-19 in the early months of 2020. My aim is to analyse the advisory and decision-making process.¹

One might wonder: what distinctive contribution can philosophy of science make to debates about the response to COVID-19? We are not a public inquiry; it is not our role to apportion blame or credit, to demand apologies, or to make recommendations that are specific to a single national context. We are not journalists; it is not our role to tell gripping narratives.²

We have not been at the table when decisions are made, and there is little point in us telling the world what decisions we would have supported, had we been at the table. What we can do is study the table. We can analyse, with a philosophical eye, the dynamics of the relationship between government and advisor, and that between science and values, in unprecedented and dire circumstances, in the hope of extracting valuable lessons. That is my project here.

Why do this? I have four interrelated goals in view. One is to arrive at generalizable norms for effective scientific advising: norms that may usefully generalize to other national contexts and to other major crises, including future pandemics. Another is to better understand how

¹ In analysing this case, I have been particularly influenced by analyses of the interplay of science and values in other cases by Douglas (2009), Steele (2012), John (2015), Lewens (2018) and the contributors to Elliott and Richards (2017). Atkinson et al. (2020) have analysed the same events using different evidence, namely interviews with panels of witnesses, and the themes they highlight are somewhat different from my themes.

² For a good narrative account of the period under discussion here, see Grey and MacAskill (2020).

normal advising differs from advising *in extremis*, when the lives of a significant fraction of a country's citizens are in immediate peril and there is no adequate pre-existing plan for managing the risk. I will argue that there is a normative difference between these two contexts of scientific advising (that is, different norms apply in *extremis*) and I aim to clarify the nature of the difference. A third is to better understand the government-advisor relationship in *extremis*, and in particular the sense in which advice should be *independent*. A fourth is to better understand the relation between science and values, and in particular the role that non-epistemic (ethical, social, political) value judgements may, in *extremis*, play in scientific advising.

The process of scientific advising in the UK has been impressively transparent, providing us with a rich set of resources on which to draw. The most senior advisory group guiding the UK's initial response was SAGE, the Scientific Advisory Group for Emergencies. The group met 74 times in 2020. Since late May 2020, the minutes have been made publicly available, usually within one month of the meeting. SAGE is represented on COBR (also known as Cobra), the UK's primary decision-making body for civil contingencies, which is traditionally (but not always) chaired by the Prime Minister. Although minutes from COBR are classified, it is reasonable to assume that the SAGE minutes, along with other research papers and memos released by SAGE, provide a fair reflection of the scientific advice being provided to COBR at the time in question.

Another advisory group, NERVTAG (the New and Emerging Respiratory Virus Threats Advisory Group), formally advises the Department of Health and Social Care, but throughout 2020 has collaborated closely with SAGE. This group has met more than 30 times, providing advice which feeds in to SAGE, and its minutes have also been made publicly available.

SAGE also has two important subgroups: SPI-M-O (Scientific Pandemic Influenza Group on Modelling, usually known by its earlier name of SPI-M) and SPI-B (Scientific Pandemic Insights Group on Behaviour). While the minutes of these groups are not published, SPI-M-O has regularly produced "consensus statements" that are intended to convey to SAGE the consensus view within the group, and these are publicly available. A further source of evidence is public testimony to the Health and Social Care Select Committee and the Science and Technology Select Committee.

The analysis that follows will be structured around three topics: (1) the normative force of advice, (2) the role of reasonable worst-case scenarios, and (3) the limits of independence and neutrality. In each case, I will draw on the available resources to highlight key features of the advisory process, leading me to propose one or more generalizable norms for scientific advising in extremis. I will conclude with an epilogue that shifts the focus to September 2020. I will ask whether some of the lessons that could have been learned from the experience in the spring were in fact learned.

2. The normative force of advice: evidence

The slogan "advisers advise and ministers decide", coined by Margaret Thatcher in 1989, is a popular saying in UK government circles. The Chief Medical Officer, Chris Whitty, has himself used it to describe how he sees the relationship between SAGE and government (HSCC 2020, Q646). But what does it mean in practice? Should scientific advisers limit themselves to advising on what means would be effective in relation to which ends, without endorsing any particular end? Should they provide menus of options from which ministers choose? Or should they issue imperatives, which ministers can either follow or ignore?

It will be helpful, for my purposes, to introduce three broad forms a piece of scientific advice can take, corresponding to the above options:

No unconditional recommendations (NUR): NUR-advice declines to endorse any end, focusing instead on conditionals and means-end relationships. It may take the form "If your goal is this, then this would be an effective means. If, on the other hand, your goal is this..." or the form "If you do this, then we expect this to happen..."³

Disjunctive unconditional recommendation (DUR): DUR-advice endorses an end but leaves open various substantively different means, to leave meaningful choices open to the policymakers following the advice. It takes the form: "Do this or this or this or..."

Single unconditional recommendation (SUR): SUR-advice endorses a single end and a single means. It takes the form: "Do this."

The shift from NUR to either DUR or SUR is a normatively significant one, because it involves endorsing a goal, and this requires non-epistemic (ethical, social, political) value judgements about the relative merits of different goals. I will henceforth refer to non-epistemic value judgements simply as "value judgements". The shift from DUR to SUR may involve additional value judgements, if the narrowing down of means to a single option is based on ethical, social or political considerations.

³ We could split NUR into two sub-categories: advice that emphasizes means-end relationships and advice that merely states conditionals about the effects of possible actions. For my purposes, I will lump these together.

This three-way distinction is helpful for understanding what happened in the UK in early 2020. When we look at the evidence, we see a broad trend: SAGE starts with NUR, then shifts to DUR (on 9 March) and then finally to SUR (on 18 March). My contention (in light of this trend) will be that the case reveals the limitations of NUR and DUR advice in extremis—and the distinctive value of SUR advice.

There is a draft paper, dated 26 February and discussed at a SAGE meeting on 27 February 2020, that I take to capture the consensus view of SAGE at that moment (SAGE 2020b). It steers carefully clear of any unconditional recommendations. The paper presents various mitigation options, including social distancing and shielding, and various ways of combining them. For example, the paper notes that "Implementing a subset of measures (e.g. the first three) would be expected to have a more moderate impact – still substantially reducing peak incidence, while making a second wave of infection in Autumn less likely. This might be the preferred outcome for the NHS." (SAGE 2020b). This is an example of NUR-advice.

Different options are presented and related to different goals, without endorsement.

SAGE adds that "It is a political decision to consider whether it is preferable to enact stricter measures at first, lifting them gradually as required, or to start with fewer measures and add further measures if required" (SAGE 2020b). This is striking because the decision in question is not *purely* political. The right decision depends partly on normative/evaluative considerations (How bad is it to compromise civil and economic liberties more than is necessary? How bad is it to expose citizens to more risk than necessary?), but it also depends partly on scientific facts about the comparative effectiveness of initially strict measures which are then relaxed as opposed to initially lax measures which are then tightened. When a judgement inherently involves evaluative/normative considerations, and yet also depends

sensitively on scientific facts, I will call it a *mixed judgement*.⁴ SAGE’s approach, at this time, is to pass mixed judgements over to political leaders.

An updated draft, dated 4 March, is discussed on 3 March (SAGE 2020c). In this draft the advisers add a comment that explicitly adopts a NUR approach: “SAGE has not provided a recommendation of which interventions, or package of interventions, that Government may choose to apply” (SAGE 2020c, their underlining). In the early days of the crisis, SAGE has already seen modelling results clearly showing the potential for a public health catastrophe, but they do not believe their role involves making unconditional recommendations.

In another updated draft, dated 9 March and discussed on 10 March (SAGE 2020d), there is a detectable change of approach. The comment about “implementing a subset of measures” is now modified to include an explicit endorsement: “A combination of these measures is expected to have a greater impact: *implementing a subset of measures would be ideal*. Whilst this would have a more moderate impact it would be much less likely to result in a second wave” (SAGE 2020d, italics added).

What SAGE explicitly endorses here is *not* maximally aggressive suppression of community transmission. What it endorses is the implementation of a “subset” of the measures listed in the paper. The prevailing at this time is that the costs of maximally aggressive suppression would exceed the benefits, because suppressing transmission completely would lead to a catastrophic, unmitigated epidemic when the measures were relaxed. An important epidemiological point here is that “flatter” epidemics spread out over a longer period are, in

⁴ This term is inspired by Plutynski’s (2017) discussion of mixed judgements in cancer screening, and Alexandrova’s (2018) discussion of mixed claims in the science of well-being.

theory, less severe than fast epidemics with sharp peaks, both because the healthcare system is less overstretched at any given time *and* because the total number of people infected is likely to be lower.⁵ This received wisdom is captured by a sketch of a graph in the 4 March and 9 March draft papers, which appears designed to illustrate the superiority of “high transmission reduction” over “very high transmission reduction, later lifted” (**Figure 1**).

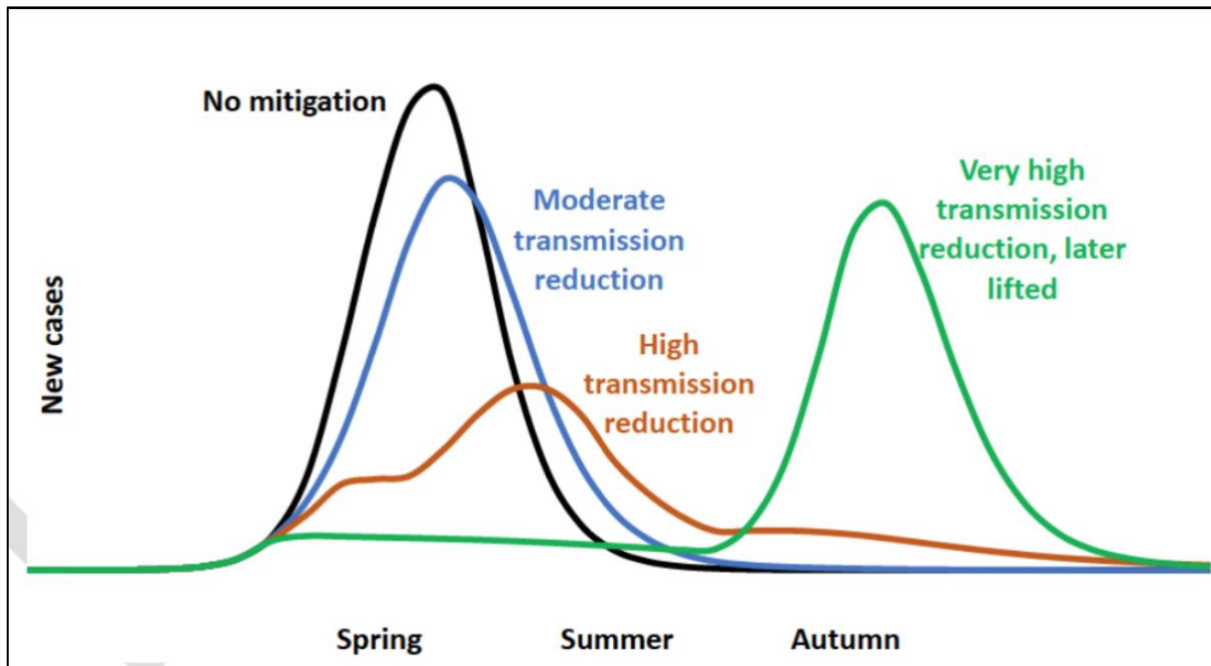


Figure 1: A sketch of a graph from SAGE 2020d, dated 9 March 2020.

The advice given is an example of DUR-advice. The recommendation is disjunctive: the government is advised to implement a “subset” of measures, but the question of which subset, from the many possible subsets, is deliberately left to ministers. A subset could range

⁵ Suppose (for example) a minimum of 60% of the population must be infected for herd immunity to be achieved. In a fast epidemic with a sharp peak, this percentage is likely to be overshot, potentially by a large amount.

from all but one of the suggested measures to just one or two.⁶ DUR-advice is a compromise: it involves making a value-laden judgement about ends (“high transmission reduction” is the goal), while remaining as neutral as possible about means to those ends.

There is another change of approach on 16 March. On that day, SAGE discusses several new modelling papers, including one from the Imperial College COVID-19 Response Team (Ferguson et al. 2020). This paper explicitly contrasts mitigation strategies that aim *not* to suppress transmission completely (i.e. the type of strategy endorsed by SAGE in the 9 March paper) with a strategy of maximally aggressive suppression, including school closures. The paper contains the first example (chronologically) in all the documents I have reviewed of a single unconditional recommendation: "We therefore conclude that epidemic suppression is the only viable strategy at the current time" (Ferguson et al. 2020).

The reasoning that leads to this conclusion is value-laden, but I do not mean this as a criticism. For some time, models discussed by SAGE had been showing that, without maximally aggressive suppression, the demand on critical care beds would be enormous. SAGE had already written on 26 February that “In the reasonable worst-case scenario, demand on beds is likely to overtake supply well before the peak is reached” (SAGE 2020b). This was not news, to SAGE, on 16 March. But prior to 16 March, the prevailing view nonetheless opposed maximally aggressive suppression, for the reasons noted above.

So what was the news on 16 March? For one thing, reliable data about critical care capacity was plotted on the same graph as “reasonable worst-case scenario” predictions regarding

⁶ Strictly speaking, even the empty set is a subset, but in context it is clear the advisers mean *non-empty* subset.

demand for critical care, revealing the size of the mismatch (**Figure 2**). It is plausible that some members of SAGE, while aware of the mismatch, were not aware of its dramatic size. Even with mitigation measures in place, the model predicted that demand was likely to exceed supply by at least a factor of eight.




Figure 2: Projected demand for critical care and critical care capacity, from Ferguson et al. (2020).

To move from these grim forecasts to the conclusion that maximally aggressive suppression is the only viable strategy, value judgements are required. This is another example of a mixed judgement. The role of value judgements becomes particularly clear when we note that the strategy the authors are actually endorsing involves not just *short-term* school closures, but *regular, sustained school closures* until a vaccine or effective treatment is implemented or until herd immunity is achieved. The forecast is that school closures will be needed again whenever cases start to rise sharply, and that this will have to be done for roughly two thirds

of the time for at least 18 months (**Figure 3**). Prior to 16 March, I can find no evidence of this having been considered by SAGE as a realistic possibility: the expectation was that measures would be one-shot, time-limited, and the epidemic would return when they were lifted. The new possibility, which Imperial now describes as the only viable option, involves a level of school closure that was previously unthinkable. There is a value judgement involved in regarding this as a genuine option, and further value judgements involved in regarding it as so clearly preferable to an overwhelmed healthcare system that there is a normative imperative to pursue it. The value judgements here are *momentous*, in the sense that they involve weighing large-scale harms and benefits (of various kinds) to millions of people.




Figure 3: A graph from Ferguson et al. (2020), in which the blue-outlined blocks represent periods of school closure.

The Imperial paper directly informs the advice given by SPI-M-O to SAGE 16 March, in which SPI-M-O writes:

It was agreed that the addition of both general social distancing and school closures to case isolation, household isolation and social distancing of vulnerable groups would be likely to control the epidemic when kept in place for a long period. *SPI-M-O agreed that this strategy should be followed as soon as practical, at least in the first instance.* (SPI-M-O 2020b, italics added)

This is another example of SUR-advice (a single unconditional recommendation) and it is an explicit endorsement of Ferguson and colleagues' recommendation.

SAGE, at this moment, continues to prefer disjunctive recommendations. The minutes of the 16 March meeting record that "'SAGE advises that there is clear evidence to support *additional social distancing measures* be introduced as soon as possible." (SAGE 2020e, italics added). The nature of the additional measures is left open, and school closures are not yet recommended. Two days later, this changes. On 18 March, SAGE advises that "available evidence now supports implementing school closures on a national level, as soon as practicable to prevent NHS intensive care capacity being exceeded."

I read this as an unconditional recommendation. The language admittedly leaves some room for ambiguity: the language is still about what the "evidence supports", not about what the government should do. But an end and a specific means to that end are being endorsed. This is the first and, to my knowledge, only case of SAGE making a single unconditional recommendation. It is no longer avoiding recommendations altogether or presenting menus of options. It is no longer seeking to avoid, as far as possible, value judgements. It is

endorsing a specific political choice. SAGE is saying: do this now. The recommendation was followed, and schools closed on 20 March.

3. The normative force of advice: reflections and a proposal

What can we learn from this case? What it shows, I suggest, is the limitations of NUR-advice and DUR-advice in extremis. When I say “in extremis”, I mean that the situation is one in which the lives of a significant fraction of a country's citizens are in immediate peril, and no pre-existing plan for managing this situation can be straightforwardly implemented.⁷ These situations call for *momentous mixed judgements*: judgements that involve an evaluative weighing of large-scale harms and benefits to millions of people, and yet also depend sensitively on scientific facts. The decision to close schools (or not) in the interests of public health is an obvious example.

Momentous mixed judgements are not unique to situations of immediate peril. However, in normal situations, there is much more time to make them. This allows the possibility of a slow process in which we settle the relevant scientific issues to an acceptable level of certainty, carry out cost-benefit analyses of different options, and then communicate all this information (along with the residual uncertainty) to political leaders and to the public. In normal times, decision-makers can consult widely, and potentially seek a democratic mandate, before making such judgements. On this model, it seems entirely reasonable for advisers to avoid unconditional recommendations, as SAGE initially did.

⁷ I don't think precise thresholds for “significant” or “immediate” are possible. These facts may be sources of debate, and normative considerations may be relevant to the judgement. This is discussed later in this section.

By contrast, scientific advising in extremis does not have the luxury of time. The danger of NUR-advice is that political leaders will not see any imperative to act immediately, since no such imperative has been communicated to them, and they will lack the expertise to infer such an imperative from epidemiological data. If nothing has been recommended, the option of doing nothing (and waiting for more information) is likely to be seen as an attractive one. It is very hard to imagine a political leader taking an action as drastic as closing schools without this action first having been recommended by a scientific adviser.

The problem with DUR-advice is that it is an awkward compromise. It satisfies the demand for recommendations, but leaves some important mixed judgements open, and the upshot is that some judgements that require scientific expertise are still outsourced to ministers who lack that expertise. For example, if you call for a “subset” of social distancing measures, while leaving open the nature of the subset, ministers are left to choose which subset to implement. Yet the optimal subset will depend not just on evaluative/normative considerations but also on epidemiological considerations. Similarly, if you leave it to ministers to decide whether to go for initially strict measures that can be relaxed later, or initially lax measures that can be tightened later, they will be drawn towards the latter option for political reasons, but will lack the scientific expertise necessary to evaluate which of these options makes more epidemiological sense.

Single unconditional recommendations (like “close schools now!”) raise serious concerns of a different kind. They require momentous value judgements, and those value judgements (on the SUR approach) are made by unelected scientific advisers. They are made by modellers at Imperial, by SPI-M-O, and by SAGE. Ministers simply receive an imperative which they can accept or reject. When such an imperative arrives, in the middle of a serious crisis, it carries a

great deal of normative force and would be politically difficult for any leader to ignore—the recommendation is close to being a decision. This leads to the worry that there is no true democratic accountability for the value judgments. There is no particular reason to think the value judgements reflect the values of voters, and no way for voters to replace scientific advisers when they disagree with their value judgements. I see this concern as having a great deal of force *in normal times*. It leads to an accountability gap: no one is democratically accountable for value judgements that shape people’s lives in dramatic ways (ask any parent – a school closure is dramatic). However, I don’t see this as a decisive reason to avoid SUR *in extremis*.

That leads me to my first proposal: I suggest *there is normative difference between normal scientific advising and scientific advising in extremis*. In the background to this proposal is Michael Walzer’s work on the “problem of dirty hands” (Walzer 1973, 1977, 2004). Walzer argues that there is a normative difference between the norms of political leadership in normal times and the norms of political leadership in extremis. In normal times, a political leader should adhere to the moral norms of the community they lead. But in extremis, when the moral community itself is in immediate peril, different norms apply, and it may be reasonable for a political leader to do things that violate the moral norms of the community. For example, perhaps a leader could permissibly divert enemy bombing away from strategically important city centres on to working class residential areas, as the UK did in World War Two (Burri 2018). I am proposing that we extend this from political leaders to scientific advisers: I suggest that different norms apply to scientific advisers in extremis, and that forms of advice that may be inappropriate in normal times can, in extremis, be permissible and praiseworthy.

In particular, I suggest that *the norm against accountability gaps, although reasonable in normal times, may be appropriately suspended in extremis*. Some accountability gaps may be tolerated. The ability to make momentous mixed judgements may be temporarily handed over to scientific advisers, so that the advisers are asked to make a single unconditional recommendation, which political leaders may either implement or veto.

A complication here is that the judgement of when exactly the shift from normal times to immediate peril has occurred is *itself* a momentous mixed judgement. The borderline will be vague, and the precise moment to make the shift will not be fixed by the scientific facts, but will also have an evaluative component (cf. Graff Fara 2000). So who should make *this* mixed judgement? It should, ideally, be the government, so that there is at least accountability for the initial creation of accountability gaps. I propose this as an ideal, not something that can always be achieved, and it was probably not achieved in the case under discussion. In this case, the shift from DUR to SUR appears to have happened in a bottom-up way, starting with Imperial and proceeding through SPI-M-O to SAGE. In effect, the modellers at Imperial took it upon themselves to make a single unconditional recommendation on 16 March. Ideally, the government would decide that a SUR is appropriate in the circumstances, and explicitly instruct its advisers to make one.

According to one credible estimate (Birrell et al. 2021), the cumulative number of COVID-19 infections in England was 14,600 on 2 March, 891,000 on 20 March (when schools closed) and 1.74m on 23 March (when a lockdown began). It is tempting to say: if SAGE had unambiguously recommended maximally aggressive suppression (including school closures) on 2 March, when SPI-M-O first reported that “It is highly likely that there is sustained transmission of COVID-19 in the UK at present” (SPI-M-O 2020a), then many lives would

have been saved, the epidemic would have taken a different course, and the school closures themselves would have been much briefer than they ultimately were. But I need to be clear here: my proposal in this section concerns the *form* of advice, not its content. While SAGE could certainly have issued SURs before 18 March, there is no particular reason to think (in light of the evidence covered in Section 2) that they would have recommended maximally aggressive suppression, and some reason to doubt this. The question of how the *content* of the advice could have been improved will be taken up in the next section.

4. Reasonable worst-case scenarios: evidence

The concept of a “reasonable worst-case scenario” (RWCS) has been at the core of SAGE’s approach to the pandemic from the beginning. An RWCS is a set of assumptions that reflect one way in which the epidemic in the UK may unfold. The set of assumptions is “reasonable” in the sense of being regarded by SAGE as a serious possibility. It is “worst-case” in the sense of being at the pessimistic end of the range of serious possibilities.

The RWCS concept has been prominent in civil contingencies planning in the UK since at least 2009, when it played an important role in the response to the H1N1 (“swine flu”) pandemic. A succinct explanation of the concept was presented by SAGE to the Science and Technology Select Committee in 2010:

The reasonable worst case is a concept developed for emergency planning in the UK. This concept is designed to exclude theoretically possible scenarios, which have so little probability of occurring that planning for them would lead to a disproportionate use of resources. They are not predictions of what will happen but of the worst that might realistically happen, and therefore we would expect

most pandemics to be less severe and less widespread than the reasonable worst case. By planning for the reasonable worst case planners are assured that they have a high probability of meeting the demands posed by the hazard should it occur. (STSC 2010, Appendix A)

What does it mean to plan for the reasonable worst case? We can distinguish two types of role an RWCS could play in pandemic planning, a *strategic* role and an *operational* role:

Strategic role: the RWCS feeds into decisions about high-level strategy, such as which social/behavioural interventions to implement and when.

Operational role: Once the high-level strategy is fixed, the RWCS is circulated to the various agencies charged with implementing the strategy, to guide finer details of planning (e.g. how large a stock of a given item should be purchased).

These can also be conceptualized as “upstream” (feeding *into* top-level decision-making) and “downstream” roles (feeding down *from* top-level decision-making to the implementation mechanisms), respectively. What we see, in our case, is that the RWCS plays both roles. My contention will be that the case reveals the limitations of allowing RWCS assumptions to dominate strategic, and not just operational, planning. The principle behind the use of RWCSs seems to be this: if you assume you are in the RWCS, and plan accordingly, then you will be as well prepared as possible for less severe scenarios. This principle is plausible when the RWCS is used in an operational role, but questionable when the RWCS is used to set high-level strategy.

What was the RWCS, in the UK's initial response to COVID-19? It was, in some respects, highly pessimistic. SAGE's RWCS planning assumptions, adapted from an existing RWCS for pandemic influenza, set out a scenario in which 80% of the population gets COVID-19 over a period of about nine weeks, with 50% displaying symptoms, and an infection fatality rate (IFR) of 1%. The result would have been around 520,000 excess deaths within three months. In a draft dated 4 March 2020, it is also explicitly assumed that no effective treatments or vaccines will become available either before or during the epidemic. This line is deleted from the version dated 6 March, but seems to tacitly guide strategic planning.

SAGE also makes some highly pessimistic background assumptions that, although not formally part of the RWCS, are important for understanding it. First, it is assumed that contact tracing would cease to be effective once there were more than 50 cases per week. Second, it is assumed that interventions (such as social distancing and shielding) could be sustained for a maximum of 13 weeks. Some modelling was done of a scenario in which they are sustained for 26 weeks, but 13 weeks is the assumption in SAGE's written advice. Third, partial compliance is assumed. For example, it is assumed that 50% will comply with household quarantine. SAGE describes this in its 4 March paper as "high levels of compliance" (SAGE 2020c).

As we saw in Section 3, the modelling at this time robustly indicates that maximally aggressive suppression merely postpones the epidemic to a time when measures are relaxed (see Figure 1). The prevailing view within SAGE is that one should avoid suppressing transmission too much, so as to flatten the curve over a longer period of time. What the graph sketch in Figure 1 does not show is that all the depicted scenarios are ones in which

healthcare system (NHS) capacity is at some point vastly exceeded, either sooner or later, regardless of mitigation strategies.

There is a line in SAGE's 26 February draft advice that says "in the event of a severe epidemic, without action, the NHS will be unable to meet all demands placed on it. In the reasonable worst case scenario, demand for beds is likely to overtake supply well before the peak is reached" (SAGE 2020b). The words "without action" are deleted between the 26 February draft and the 3 March draft (SAGE 2020c) to take account of the new modelling that has become available on 2 March. The new modelling suggests that "without action" is misleading, because none of the actions being considered at that time would be enough to prevent demand for healthcare outstripping supply.

Yet in other respects, the RWCS assumptions are excessively optimistic. SAGE (in the 4 March draft) assumes that R_0 , the basic reproductive rate of the virus in the absence of mitigation, is 2.4, leading to a doubling time of 4-6 days. Estimates varied a great deal at the time, and still do, but this was, even then, towards the lower end of serious estimates for R_0 . A study published in *The Lancet* on 31 January had estimated R_0 at 2.7 (Wu et al. 2020). On 11 February, researchers at the Theoretical Biology and Biophysics unit at the Los Alamos National Laboratory, USA, released a preprint estimating that R_0 was between 4.7 and 6.6 (Sanche et al. 2020). SAGE's line about R_0 is deleted in the 6 March draft, in which no estimate of R_0 is included.

There is also another optimistic background assumption, introduced on 25 February, that surveillance "should provide evidence of an epidemic around 9- 11 weeks before its peak" (SAGE 2020g). Just as R_0 being higher than expected is not part of the RWCS, surveillance

being poorer than expected is also not part of the RWCS. In short, the RWCS assumptions are a mix of bleak pessimism and excessive optimism.

What were the consequences of these choices? The costly delays between 2 March and 23 March, when cumulative infections surged from about 17,600 to about 1.74m (to use the MRC-BSU estimates; Birrell et al. 2021), can be seen as consequences of planning assumptions made around the beginning of this period. The combination of an optimistic estimate for R_0 , optimism about surveillance, and an assumption that 13 weeks was the realistic maximum duration for any intervention, made it seem optimal to delay the introduction of social-behavioural interventions that were in fact urgently needed.

This problem was compounded by continued scepticism, up to at least 16 March, about the idea of maximally aggressive suppression. As we have seen, for a time, between 9 March and 16 March, SAGE's advice was explicitly *against* maximally aggressive suppression and in favour of more moderate measures. There was a lot of talk, at the time, about "herd immunity". SAGE was perceived by its critics to be following a "herd immunity strategy", whereas the Chief Scientific Adviser (Sir Patrick Vallance) vehemently denied that there was any such strategy in private emails subsequently released to the BBC (Kermani 2020). We can see that the recommended strategy between 9 and 16 March was a "high transmission reduction strategy" in the sense of Figure 1, which aimed to flatten the curve without suppressing it completely. This is a strategy in which herd immunity is achieved, but slowly and in a managed way, so as to lower the peak of maximum healthcare demand (hence the resistance to the term "herd immunity strategy").

This initially sceptical attitude towards maximally aggressive suppression can also be traced to the RWCS planning assumptions. In a reasonable worst-case, no effective treatment, contact tracing system or vaccine becomes available before measures are relaxed, and no long-term behavioural changes are instilled, so total suppression leads inevitably to the epidemic returning with unmitigated force in the autumn, infecting 80% of the population and overwhelming the health service. This corresponds to the green line in Figure 1. In that bleak scenario, we come to bitterly regret the aggressive measures adopted in the spring. This is not in fact come to pass. We have indeed experienced a large second wave, but not on the scale of the original RWCS portrayed in Figure 1, which is close in size to a completely unmitigated epidemic. A contract tracing system has been implemented, though its effectiveness is a source of debate. More importantly, multiple vaccines have become available and (at the time of writing) are proving effective.

5. Reasonable worst-case scenarios: reflections and proposals

Let us return to the principle guiding the use of RWCSs: if you assume you are in the RWCS, and plan accordingly, then you will be as well prepared as possible for less severe scenarios. I will call this the **RWCS principle**. This case, I suggest, shows us some important exceptions to the RWCS principle.

The principle is not true if your RWCS is pessimistic in *most* relevant respects but not *all* of them. I will introduce the term “globally pessimistic” for a scenario that is pessimistic in all relevant respects. If the RWCS is not globally pessimistic (e.g. its estimate for R_0 is too optimistic), then there is a serious risk that reality will be worse than the RWCS in those specific respects. This is what happened in relation to R, which was almost certainly above

2.4 before the spring lockdown in England, according to the MRC-BSU, despite substantial attempts at mitigation (Birrell et al. 2020).

Of course, a globally pessimistic scenario may well seem unreasonable (it is tempting to think: it is very unlikely that we will be unlucky in *all* respects!), and this brings out a tension in the concept of a RWCS. To avoid reality catching you out, the RWCS must be globally pessimistic, and may therefore look unreasonable when viewed as a whole, but what matters is that *each individual assumption* used to specify the scenario represents a serious possibility. To be clear, these assumptions should still be at the population level, and may still include estimates of statistical variables (we should not think: any individual could die from COVID-19, therefore the RWCS should include a 100% fatality rate). This leads me to a proposal: *RWCSs, if used in any form of planning, should be globally pessimistic, which is to say at the pessimistic end of scientific opinion regarding all of their planning assumptions, not just some.*

More subtly, there can be circumstances in which assuming you are in the RWCS justifies actions, delays, or omissions that will impair your response significantly if you are in a less severe scenario. For example, it makes sense to say that, *in a reasonable worst case*, no effective treatment or vaccine will become available, contact tracing will never become effective, and no long-term behavioural changes will be instilled, even if you delay the epidemic by several months. In this case, maximally aggressive suppression of transmission is likely to make things worse in the long run (Figure 1). But suppose you are in fact in a less severe scenario, in which one of these three pessimistic assumptions is false. In such a scenario, maximally aggressive suppression is likely to be far superior to more moderate action, in terms of both its public health consequences and its long-term economic

consequences. If you plan for the RWCS, and are thereby led to adopt a strategy that involves aiming for the brown/orange line in Figure 1, you will not be as well prepared as possible for a less severe scenario.

This leads me another proposal concerning the use of RWCSs in scientific advice: although they are a useful guide for operational planning, *RWCSs should not dominate strategic planning*. For strategic planning, it is important to consider a *wide range of possible scenarios*, including but not limited to the reasonable worst case (which, if globally pessimistic, will be unlikely). The apparent inevitability of a large wave as soon as measures were relaxed was sensitive to a specific set of worst-case planning assumptions, which assumed that measures could not be sustained until an effective treatment, vaccine, or contact tracing system was implemented. These assumptions dominated strategic planning up to 16 March. Yet as soon as modellers dared to relax one of those pessimistic assumptions, as in Ferguson et al. (2020), the strategic picture dramatically changed.

It might be objected that an imperative to consider a wide range of possible scenarios, including more optimistic ones, makes more sense in normal times than in extremis. Would considering a wider range of scenarios not simply have led to even more delays, and potentially to unfounded optimism? I accept that the scientific advisers were not in a position to assign precise probabilities or values to different scenarios, and were therefore not in a position to do an expected utility calculation. However, I do not think this removes the obligation to consider a wide range of scenarios.

Even without assigning probabilities to scenarios, we can look for policies that deliver acceptable outcomes *robustly* (under a wide range of parameter values), as opposed to

policies that deliver acceptable outcomes only under very specific circumstances. One natural suggestion, then, is that advisers should only recommend actions that robustly lead to acceptable outcomes across all realistic scenarios (see e.g. Bradley and Bright 2020). Yet this may be asking too much in this context. From SAGE’s point of view, there was a wide range of scenarios in which maximally aggressive suppression leads to an unacceptable outcome, i.e. a devastating epidemic when measures were relaxed. This happens robustly in models in which there is no place for a vaccine, an effective treatment or effective contact tracing. The problem is not that this result was particularly fragile with respect to the variables actually being modelled, but that it was dependent on potential game-changers (i.e. the possibility of a vaccine, effective treatment or effective contact tracing becoming available) that were not modelled at all—and could not realistically be modelled except as sudden, exogenous changes to the infection fatality rate (IFR) or reproductive rate (R) of the virus.

It may be tempting to say: in that case, what is needed is robustness across a wide range of models that *includes* potential game-changers. The problem with this suggestion is that game-changers can cut both ways. There are also scenarios in which the IFR or R rate is suddenly increased by new variants of the virus evolving, and in some of those unfortunate scenarios we would regret our failure to build up herd immunity when we had the chance.⁸ There is no principled reason for including positive game-changers in the model while excluding

⁸ Avian influenza viruses often evolve in a way that *increases* their pathogenicity. As Akhtar (2012, p. 106) notes, “All avian influenzas start off as mild, low-pathogenic viruses. However, once they enter domestic bird populations, they can rapidly mutate into highly pathogenic viruses.” Once a respiratory virus enters a large population of close-packed individuals, it is under selection for reproducing fast and being emitted fast by its host—not selection for sparing its host’s life. For discussion of the relevant selection pressures, see Sober (2020).

negative game-changers. There is a deep problem here: robustness across game-changer-free scenarios is not enough, but asking for policies that are robustly effective across all game-changer scenarios is asking the impossible. A choice has to be made as to which game-changers we want our planning to be robust against and which we do not, and here too it seems that momentous mixed judgements are unavoidable. These judgements, when made by advisers, should be communicated as clearly as possible to decision-makers so that they can be challenged if necessary.

This leads me to a fifth proposal: *scientific advisers should highlight, as part of their advice, the scenarios in which their recommended actions would lead to serious regret. In other words, they should communicate which potential game-changers they have chosen to set aside when looking for robustly acceptable policies.* For example, a single unconditional recommendation to close schools should highlight that, if the fraction of asymptomatic infections turns out to be very high, such that herd immunity has already been reached, children would suffer a great harm for no public health benefit. Such warnings should then be contextualized, with (for example) an explanation of why betting on the fraction of asymptomatic infections being so high would be an extraordinary bet. Likewise, when SAGE recommended a “high transmission reduction” strategy over maximally aggressive suppression, it should have highlighted, in a contextualized way, the potential for this strategy to lead to serious regret—in the form of tens of thousands of excess deaths which could in retrospect have been prevented by pursuing a more aggressive suppression strategy at an earlier stage. Decision makers who are given a clear recommendation, together with a contextualized explanation of scenarios in which the recommendation would lead to serious regret, can make a decision about whether the risk of serious regret is acceptable to them or not.

6. Independence and neutrality

The neutrality and independence of scientific advice are not the same thing. Here is how I propose we distinguish them:

Neutrality: Scientific advice remains neutral on politically contested value judgements, such as the economic cost that should be accepted to save one life.

Independence: Scientific advice is formulated without undue constraint or influence from political actors.

It would be helpful if we could spell out precisely what qualifies as “politically contested” or “undue constraint or influence”. However, we cannot do this from the armchair. It is through reflecting on the details of particular cases, such as this one, that we can develop more precise norms regarding neutrality and independence.

I take it that *some* constraint and influence is both desirable and unavoidable. Political leaders have a legitimate role in posing questions to advisers, in establishing advisory bodies such as SAGE, and in setting their terms of reference. They can also legitimately communicate constraints regarding what is or is not politically feasible, in their view. But I also take it that some types of influence and constraint are undue and compromise an adviser’s independence. To give an obvious example, bribing an adviser to give favourable advice would be undue influence. In SAGE’s 68 page guidance document (Cabinet Office 2012), neither the neutrality nor the independence of SAGE is discussed, although there is some discussion of how subordinate bodies like SPI-M-O relate to SAGE.

The SAGE minutes and papers suggest a broad commitment to neutrality. For example, all three drafts of the advice on social and behavioural interventions drawn up in late February and early March (SAGE 2020b,c,d) state that “the note does not cover economic, operational or policy considerations.” While it is not clear what is meant by “policy considerations”, this may be read as encompassing politically contested value judgements. A desire to be neutral may be part of the explanation for the reluctance to make recommendations in late February and early March, although the lack of any estimates regarding the economic costs of the measures under consideration is probably also part of the explanation.

Occasionally, SAGE does violate neutrality in subtle and perhaps unintended ways. For example, on 3 February 2020 SAGE noted that “to prevent imported infections ... would require draconian and coordinated measures” (SAGE 2020h) and that merely stopping direct flights from China would be insufficient. I do not read “draconian” as a fully neutral term; it suggests that to close borders would be excessive. On 11 February 2020, the SAGE minutes record that “it is not possible for the UK to accelerate diagnostic capacity to include COVID testing alongside regular flu testing in time for the onset of winter flu season 2020-21” (SAGE 2020f). This is framed as a fact about operational capacity, but there is an implicit political judgement involved in the use of the term “impossible”. Tacit assumptions about what is and is not *politically* possible are in play, and in particular it is assumed that it is not politically possible to make COVID testing a major national priority and create substantial new infrastructure to support it. This was in fact politically possible—and actually happened. So, while SAGE appears to aspire to neutrality, it does not always succeed.

Yet it is not clear that neutrality is even desirable in extremis. I take it to be desirable in normal times, at least when the judgements at issue are momentous, due to the worries about democratic accountability for momentous value judgements that we considered earlier. But I have already argued that, in extremis, it may be appropriate for scientific advisers to make such judgements. In their paper of 16 March, the Imperial modellers took it upon themselves to make momentous political judgements about which strategies were “viable” or not, and they judged that averting healthcare system collapse was more important than keeping schools open. They did not act unreasonably in doing that. What strikes me as more desirable in extremis is that departures from neutrality are *communicated* as such. When something is judged politically (im)possible or (un)viable, this judgement should be explicitly conveyed, so that it can be challenged if necessary.

Let us turn to independence. Government employees are inevitably present at SAGE meetings, because the Chief Scientific Adviser (Sir Patrick Vallance) and the Chief Medical Office (Chris Whitty) are themselves government employees. Moreover, a secretariat must be on hand. I do not think this could plausibly be called undue influence. More controversially, however, advisers directly employed by the Office of the Prime Minister, 10 Downing Street, were present at the vast majority of SAGE meetings from the 20 February 2020 onwards. Ben Warner from 10 Downing Street was a regular attendee. Dominic Cummings, at that time the chief adviser to the Prime Minister, is registered as attending two meetings. This was a source of controversy at the time. It was criticized by Sir David King, a former Chief Scientific Adviser who went on to found a group called “The Independent Sage”. This group has offered independent advice on the pandemic—stressing that it, unlike SAGE, is genuinely “independent”.

Were these complaints reasonable? When the SAGE minutes were published, a standard addendum was added to all the minutes containing the following statement, emphasizing that only “scientific experts” provide evidence and advice:

There are three categories of attendee. Scientific experts provide evidence and advice as part of the SAGE process. HMG [government] attendees listen to this discussion, to help inform policy work, and are able to provide the scientific experts with context on the work of government where appropriate. The secretariat attends in an organisational capacity (SAGE 2020f).

We should make allowances for the fact that, in a pandemic, advice may need to be implemented with exceptional speed. Bringing in a recommended measure the same day, rather than several days later, could save thousands of lives. It seems reasonable that political leaders or their representatives should be on hand to hear the advice as soon as it is agreed.

Yet it is still a problem, I contend, if political leaders or their political advisers (as opposed to other officials, such as secretaries) are present throughout the *whole* meeting in which advice is discussed, formulated and agreed. I see three main reasons for this. First, there may be circumstances in which a scientific adviser wants to make a point salient *to other scientific advisers* without thereby making the point salient to decision-makers. They might, for example, want to highlight a source of uncertainty about the effectiveness of some possible action, without giving decision-makers an excuse to delay or avoid that action.

Second, scientific advisers need to be free to express *dissenting opinions* without giving decision-makers the impression that there is a lack of scientific consensus on important

points. One can easily imagine a discussion in which a dissenting view is expressed but then met with counterarguments, so that a consensus view gradually emerges. If a political leader is present or directly represented in the room, they will be privy to all the disagreement and discord that leads to the consensus—making advisers less likely to express dissenting views in the first place.

Third, advisers need to be able to hold frank discussions about how best to *communicate* their advice to political leaders. For example, they need to be able to discuss the form of advice (Should we make an unconditional recommendation?), the way in which uncertainty is communicated, and the likely consequences of different choices. There needs to be space for remarks such as: “We need to send a simple, direct message to decision-makers, because what we have now is too nuanced and complicated.” Such discussions are likely to be inhibited if the decision-maker who is the intended target of the advice is either present or directly represented in the room.

Note that the above three points still apply *even if the political advisers present in the room say nothing*. The concern is not only that they will exert undue influence by speaking, although this is obviously possible. Even if they do not speak at all, their mere presence gives scientific advisers an incentive to remain quiet on certain topics. They have an incentive not to express dissenting views, or to raise sources of uncertainty, in a way that might impede the take-up of advice.

Is there a way to accommodate these three points while also recognising the need for urgent implementation of the advice? Here is a proposal: *at least part* of any meeting of a group of scientific advisers should take place without political leaders or their political advisers

present in the room, to allow time for dissenting opinions to be expressed, for agreement to be reached, and for frank discussion of how to communicate the agreed advice (including sources of uncertainty) to decision-makers. It may still be reasonable to dedicate part of the meeting to communicating the agreed advice to political leaders and their advisers.

7. Summary of proposals

Before the epilogue, let us pause to review the proposals of the preceding sections:

1. The norms of scientific advising in extremis (when a significant fraction of a country's population faces immediate peril) are different from the norms of scientific advising in normal times. Ideally, political leaders (rather than advisers) should decide precisely when this shift has occurred.
2. In extremis, scientific advisers should make single unconditional recommendations, rather than presenting political leaders with menus of options or sets of means-end relationships. This may create an accountability gap, because no one is democratically accountable for the value judgements that lead to the recommendation, but accountability gaps may be tolerated in extremis.
3. Reasonable worst-case scenarios (RWCSs), if used at all, should be globally pessimistic: at the pessimistic end of scientific opinion regarding all of their planning assumptions, not just some or most of them.
4. RWCSs should not dominate strategic planning. Strategic planning should consider a wide range of possible scenarios and look for policies that lead to robustly acceptable outcomes, even if it is impossible to assign probabilities to these scenarios.
5. When making recommendations, scientific advisers should communicate (with context) information about scenarios in which acting on the recommendation would

lead to serious regret. In other words, they should communicate which potential game-changers they have chosen to set aside when looking for robustly acceptable policies.

6. In extremis, the norm of neutrality (i.e. remaining neutral on politically contested value judgements) may reasonably be suspended, even if the value judgements are momentous, but advisers should communicate (as well as they can) the value judgements that have shaped their advice.
7. In extremis, the norm of independence (i.e. formulating advice without undue constraint or influence from political actors) remains important and should not be suspended. It may be reasonable to include political leaders and their political advisers in meetings, to allow the rapid communication of advice. However, scientific advisers should hold at least part of any meeting without political actors in the room, to allow time for dissenting opinions to be expressed, for agreement to be reached, and for frank discussion of how to communicate the agreed advice.

The proposals are intended to be at the right level of generality to be, I hope, helpful in the management of future crises. To evaluate whether they do generalize in a useful way, I want to move forward in time—to September 2020.

8. Epilogue

In the autumn of 2020 and the subsequent winter, the UK has experienced a second major epidemic. One could argue that the virus never went away, but there was a marked lull in the summer months, followed by a resurgence. This gives us a chance to revisit the themes of this article—the normative force of advice, the role of reasonable worst-case scenarios, and norms of neutrality and independence—in a new context.

Let us first consider the normative force of advice. In September, SAGE is once again reluctant to make single unconditional recommendations. It continues to prefer disjunctive recommendations, i.e. menus of options. On 21 September, as cases rise, SAGE offers the government a “shortlist” of non-pharmaceutical interventions (SAGE 2020a). The shortlist consists of “a circuit-breaker (short period of lockdown) to return incidence to low levels”; “advice to work from home for all those that can”; “banning all contact within the home with members of other households (except members of a support bubble)”; “closure of all bars, restaurants, cafes, indoor gyms, and personal services (e.g. hairdressers)”; and “university and college teaching to be online unless face-to-face teaching is absolutely essential” (SAGE 2020a). SAGE comments that “a package of interventions will need to be adopted to reverse this exponential rise in cases” and adds that “a consistent package of measures should be adopted which do not promote, or appear to promote, contradictory goals” (SAGE 2020a). This is a clear example of DUR-advice. There are many substantially different ways to follow the recommendation.

One item from the shortlist (advice to work from home) was implemented (see Cabinet Office 2020). A package of measures was put together, but none of the other elements of the package were drawn from SAGE’s shortlist, and they were all clearly less radical interventions (e.g. attendance at weddings was limited to 15 people, down from 30). So, the government did indeed implement “a consistent package”, but the consistent package largely overlooked SAGE’s suggestions, and did not succeed in bringing R below 1.

I contend that a single unconditional recommendation would have been helpful here, and more appropriate to the seriousness of the developing situation. Recommending one or more

specific measures from the shortlist (such as the circuit-breaker lockdown) would have left the government with a simple choice: implement the recommended measures, or be seen to manifestly ignore its own scientific advisers. By framing their advice in a disjunctive way, SAGE allowed the government to escape this choice by implementing a “consistent package” that was, foreseeably, insufficient to bring R below 1.

RWCSs continue to play a dominant role in both operational and strategic planning. A new RWCS is drawn up on 30 July and is confidential until leaked to *The Spectator* on 29 October (SAGE 2020i). The new RWCS is not globally pessimistic, and is in some respects strikingly optimistic. In particular:

The scenario modelled incidence continuing as per current trends until the end of July 2020, with all non-household contacts assumed to be constant with current levels. Incidence is then assumed to double once by the end of August 2020, and double again during the first two weeks of September. *At this point, social contacts are reduced that reduce R to approximately 1, keeping infection levels steady until the end of October.* Two-week doubling times return throughout November (i.e. incidence quadruples through November), after which policy measures are put in place to reduce non-household contacts to half of their normal pre-March 2020 lockdown levels, while all schools contacts are assumed to be maintained. These measures are sustained until the end of March 2021. (SAGE 2020i, italics added)

In other words, the RWCS assumes that unspecified but highly effective measures will be taken by the government in mid-September to bring R to approximately 1. It is extraordinary

to see such an assumption feature in a reasonable *worst-case* scenario (a point made by Aaron Bell MP in a hearing of the Science and Technology Select Committee, STSC 2020b). Did scientific advisers really regard this as a reasonable worst-case?

It seems that, since the spring, a significant change has occurred in the process by which the RWCS is constructed. In the spring, SAGE had set the assumptions of the RWCS. In the summer, by contrast, the RWCS, is the result of negotiation with ministers:

This profile of increasing incidence to the end of November 2020, was agreed by SPI-M-O co-chair in collaboration with SAGE and Cabinet Office Civil Contingencies Secretariat and COVID-19 Taskforce. No specific assumptions as to what these measures may be were made. (SAGE 2020i)

This comment leaves the role of ministers rather opaque. This is partly because the command structure in the UK response, which was reasonably clear in the spring, has itself become rather opaque. The role of COBR appears to be much reduced, with two new committees, COVID-S and COVID-O, responsible for decisions regarding strategic and operational planning, respectively (this is explained by the Secretary of State for Health in STSC 2020a). I take it that at least one of these new committees was involved in agreeing the RWCS, and a SPI-M-O consensus statement from 16 September notes that “the RWCS *agreed with ministers* assumed that policy interventions would be made in mid-September to halt the rise in infections” (SPI-M-O 2020b, italics added). On 3 November, the Chief Scientific Adviser Sir Patrick Vallance explained the process to the Science and Technology Select Committee in a way that seemed to further marginalize the role of scientific expertise: “We model what the Civil Contingencies

Secretariat sees as a reasonable worst case and that is then modelled by the SPI-M modellers” (STSC 2020b, Q1510).

So the government once again adopted a RWCS that was excessively optimistic in some respects—specifically, it was excessively optimistic about the government’s own actions. Effective measures were *not* taken in mid-September and cases continued to rise throughout October, leading to a national lockdown at the end of October. Although preferable to an unmitigated epidemic, I see this as a sad outcome: the immense amount of resources invested in tracking early warning signs of a major epidemic (via an immense mass testing operation) did not translate into a swift and effective response when those warning signs were observed.

There is a lesson here about RWCSs, and a lesson about independence. An effective pandemic response is hindered if planning assumptions are negotiated between scientific advisers and political actors, with political actors pushing for unrealistic assumptions and influencing the predicted values of fundamental epidemiological parameters, such as R . If political leaders insist that strategic planning assumptions for a “worst-case” build in optimistic assumptions about their own actions, that is (I suggest) an example of undue influence. And it is *especially* important that strategic planning is not dominated by a RWCS if there is no robust way to ensure that the RWCS is free of undue political influence. Note, though, that once high-level strategic decisions are made, it is not undue influence for operational (“downstream”) planning to be informed by those decisions.

This case provides further support for some of the proposals I put forward earlier. It underlines the importance of clear, direct recommendations that are not overly disjunctive, the importance of considering multiple possible scenarios for planning purposes, and the

importance of robust mechanisms for maintaining the independence of scientific advice, which is very likely to come under threat in dire circumstances. My hope is that these proposals can generalize to other pandemics, and other major crises, and that, if enacted, they would lead to better advisory and decision-making processes.

Acknowledgements

I am very grateful to Richard Bradley, Liam Kofi Bright, Peter Godfrey-Smith, Tim Lewens, Nick Makins and Joe Roussos for their comments on a draft and discussion of those comments.

References

- Alexandrova, A. (2018). Can the science of well-being be objective? *British Journal for the Philosophy of Science*, 69, 421-445.
- Akhtar, A. (2012) *Animals and public health: Why treating animals better is critical to human welfare*. Basingstoke: Palgrave Macmillan.
- Atkinson, P., Gobat, N., Lant, S., Mableson, H., Pilbeam, C. Solomon, T., Tonkin-Crine, S., & Sheard, S. (2020). Understanding the policy dynamics of COVID-19 in the UK: Early findings from interviews with policy makers and health care professionals. *Social Science & Medicine*, 266, 113423. <https://doi.org/10.1016/j.socscimed.2020.113423>
- Birrell, P., Blake, J., van Leeuwen, E., MRC Biostatistics Unit COVID-19 Working Group, & De Angelis, D. (2021). COVID-19: Nowcast and forecast, 2021-01-14. Retrieved from: <https://www.mrc-bsu.cam.ac.uk/nowcasting-and-forecasting-12th-february-2021>
- Bradley, R., & Bright, L. K. (2020) Public health decisions when the science is uncertainty. *fifteeneightyfour* Blog. Retrieved from: <http://www.cambridgeblog.org/2020/05/public-health-decisions-when-the-science-is-uncertain/>

Burri, S. (2020). Why moral theorizing needs real cases: The redirection of V-Weapons during the Second World War. *Journal of Political Philosophy*, 28 (2), 247 - 269.

Ferguson, N. M., Laydon, D., Nedjati-Gilani, G., Imai, N., Ainslie, K., Baguelin, M., Bhatia S., Boonyasiri, A., Cucunubá, Z., Cuomo-Dannenburg, G., Dighe, A., Dorigatti, I., Fu, H., Gaythorpe, K., Green, W., Hamlet, A., Hinsley, W., Okell, L. C., van Elsland, S., Thompson, H., Verity, R., Volz, E., Wang, H., Wang, Y., Walker, P. G. T., Walters, C., Winskill, P., Whittaker, C., Donnelly, C. A., Riley, S., & Ghani, A. C. (2020). Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. Retrieved from:

<https://www.gov.uk/government/publications/impact-of-non-pharmaceutical-interventions-npis-to-reduce-covid-19-mortality-and-healthcare-demand-16-march-2020>

Cabinet Office (2012). Enhanced SAGE guidance. Retrieved from:

<https://www.gov.uk/government/publications/scientific-advisory-group-for-emergencies-sage>

Cabinet Office (2020). Coronavirus (COVID-19): What has changed – 22 September.

Retrieved from: <https://www.gov.uk/government/news/coronavirus-covid-19-what-has-changed-22-september>

Douglas, H. (2009). *Science, policy and the value-free ideal*. Pittsburgh: University of Pittsburgh Press.

Elliott, K. & Richards, T. (Eds.) (2017). *Exploring inductive risk: Case studies in values in science*. Oxford: Oxford University Press.

Graff Fara, D. (2000). Shifting sands: An interest relative theory of vagueness. *Philosophical Topics*, 28(1), 45-81.

Grey, S. & MacAskill, A. (2020, 7 April). Special Report: Johnson listened to his scientists about coronavirus - but they were slow to sound the alarm. *Reuters*. Retrieved from:

<https://www.reuters.com/article/us-health-coronavirus-britain-path-speci-idUSKBN21P1VF>

Health and Social Care Committee [HSCC]. (2020). Oral evidence: Management of the coronavirus outbreak, HC 36. Retrieved from:

<https://committees.parliament.uk/work/81/management-of-the-coronavirus-outbreak/publications/>

Kermani, S. (2020, 23 September). Coronavirus: Whitty and Vallance faced 'herd immunity'

backlash, emails show. *BBC News*. Retrieved from: <https://www.bbc.co.uk/news/uk-politics-54252272>

John, S. (2015). Inductive risk and the contexts of communication. *Synthese*, 192, 79–96.

<https://doi.org/10.1007/s11229-014-0554-7>

Lewens, T. (2019). The division of advisory labour: the case of 'mitochondrial donation'.

European Journal for Philosophy of Science, 9, 10. <https://doi.org/10.1007/s13194-018-0235-3>

Plutynski, A. (2017). Safe, or sorry? Cancer screening and inductive risk. In Elliott and Richards (2017), pp. 149-169.

Sanche, S., Lin, Y. T., Xu, C., Romero-Severson, E., Hengartner, N., Ke, R. (2020). High contagiousness and rapid spread of severe acute respiratory syndrome coronavirus 2.

Emerging Infectious Diseases, 26, 1470-1477. <http://doi.org/10.3201/eid2607.200282>

Scientific Advisory Group for Emergencies [SAGE]. (2020a). Fifty-eighth SAGE meeting on Covid-19, 21 September 2020. Retrieved from:

<https://www.gov.uk/government/publications/fifty-eighth-sage-meeting-on-covid-19-21-september-2020>

Scientific Advisory Group for Emergencies [SAGE]. (2020b). Potential effect of non-pharmaceutical interventions (NPIs) on a COVID-19 epidemic in the UK, 26 February 2020. Retrieved from: <https://www.gov.uk/government/publications/potential-effect-of-non-pharmaceutical-interventions-npis-on-a-covid-19-epidemic-in-the-uk-26-february-2020>

Scientific Advisory Group for Emergencies [SAGE]. (2020c). Potential effect of non-pharmaceutical interventions (NPIs) on a COVID-19 epidemic in the UK, 4 March 2020. Retrieved from: <https://www.gov.uk/government/publications/potential-impact-of-behavioural-and-social-interventions-on-an-epidemic-of-covid-19-in-the-uk-4-march-2020>

Scientific Advisory Group for Emergencies [SAGE]. (2020d). Potential effect of non-pharmaceutical interventions (NPIs) on a COVID-19 epidemic in the UK, 9 March 2020. Retrieved from: <https://www.gov.uk/government/publications/potential-impact-of-behavioural-and-social-interventions-on-an-epidemic-of-covid-19-in-the-uk-9-march-2020>

Scientific Advisory Group for Emergencies [SAGE]. (2020e). Seventeenth SAGE meeting on Covid-19, 18 March 2020. Retrieved from: <https://www.gov.uk/government/publications/sage-minutes-coronavirus-covid-19-response-18-march-2020>

Scientific Advisory Group for Emergencies [SAGE]. (2020f). Sixth SAGE meeting on Covid-19, 11 February 2020. Retrieved from: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/888774/S0374_Sixth_SAGE_meeting_on_Wuhan_Coronavirus_Covid-19_.pdf

Scientific Advisory Group for Emergencies [SAGE]. (2020g). Tenth SAGE meeting on Covid-19, 25 February 2020. Retrieved from:

<https://www.gov.uk/government/publications/sage-minutes-coronavirus-covid-19-response-25-february-2020>

Scientific Advisory Group for Emergencies [SAGE]. (2020h). Third SAGE meeting on Covid-19, 3 February 2020. Retrieved from:

<https://www.gov.uk/government/publications/sage-minutes-coronavirus-covid-19-response-3-february-2020>

Scientific Advisory Group for Emergencies [SAGE]. (2020i, 29 October). Covid in winter 2020, a worst-case scenario. *The Spectator*. Retrieved from:

<https://www.spectator.co.uk/article/classified-covid-in-winter-2020-a-worst-case-scenario>

Scientific Pandemic Influenza Group on Modelling [SPI-M-O]. (2020a). Consensus statement on 2019 novel coronavirus (COVID-19), 2 March 2020. Retrieved from:

<https://www.gov.uk/government/publications/spi-m-o-consensus-statement-on-2019-novel-coronavirus-covid-19-2-march-2020>

Scientific Pandemic Influenza Group on Modelling [SPI-M-O]. (2020b). Consensus view on behavioural and social interventions, 16 March 2020. Retrieved from:

<https://www.gov.uk/government/publications/spi-m-o-consensus-view-on-behavioural-and-social-interventions-16-march-2020>

Scientific Pandemic Influenza Group on Modelling [SPI-M-O]. (2020c). Consensus statement on COVID-19, 17 September 2020. Retrieved from:

<https://www.gov.uk/government/publications/spi-m-o-consensus-statement-on-covid-19-17-september-2020>

- Science and Technology Select Committee [STSC]. (2010). Memorandum submitted by the Government Office for Science and the Cabinet Office (SAGE 00). Retrieved from: <https://publications.parliament.uk/pa/cm201011/cmselect/cmsctech/498/498we02.htm#a1>
- Science and Technology Select Committee [STSC]. (2020a). Oral evidence: UK science, research and technology capability and influence in global disease outbreaks, HC 136, Tuesday 21 July 2020. Retrieved from: <https://committees.parliament.uk/oralevidence/761/pdf/>
- Science and Technology Select Committee [STSC]. (2020b). Oral evidence: UK Science, Research and Technology Capability and Influence in Global Disease Outbreaks, HC 136, Tuesday 3 November 2020. Retrieved from: <https://committees.parliament.uk/oralevidence/1122/default/>
- Sober, E. (2020). Infectious diseases and the evolution of virulence. Auxiliary Hypotheses Blog. Retrieved from: <http://www.thebsps.org/auxhyp/evolution-of-virulence-sober/>
- Steele, K. (2012). The scientist qua policy advisor makes value judgments. *Philosophy of Science*, 79, 893–904. <https://doi.org/10.1086/667842>
- Walzer, M. (1973). Political action: The problem of dirty hands. *Philosophy and Public Affairs*, 2(2), 160–180.
- Walzer, M. (1977). *Just and unjust wars: A moral argument with historical illustrations*. New York: Basic Books.
- Walzer, M. (2004). *Arguing about war*. New Haven: Yale University Press.
- Wu, J. T., Leung, K., & Leung, G. M. (2020). Nowcasting and forecasting the potential domestic and international spread of the 2019-nCoV outbreak originating in Wuhan, China: a modelling study. *The Lancet*, 395, 689-697. [https://doi.org/10.1016/S0140-6736\(20\)30260-9](https://doi.org/10.1016/S0140-6736(20)30260-9)