



Science-based policymaking: the need to think holistically, realistically, and institutionally

Faik Kurtulmus¹ 

Received: 31 October 2024 / Accepted: 19 January 2026 / Published online: 13 March 2026
© The Author(s) 2026

Abstract

This paper identifies two challenges for science-based policymaking in representative democracies and argues that both require institutional rather than individual solutions. The first arises when seeking to align value-laden choices embedded in scientific research with democratically endorsed values. The value-laden choices in different pieces of research interact once combined and can compound or offset each other. Therefore, alignment with democratically endorsed values must be assessed holistically at the level of the total body of evidence before policymakers. Since individual scientists cannot foresee or coordinate these interactions, we need institutional mechanisms to manage values at the point where advisory bodies integrate evidence from multiple sources. The second challenge emerges from the incentives of advisory body members and policymakers to shape the evidence informing policy in ways that undermine democratic accountability. Addressing this requires a realistic approach to advisory process design that accounts for actors' incentives and constraints. Taken together, these challenges show that integrating scientific advice into policymaking in ways that uphold democratic commitments requires an institutional approach.

Keywords Science-based policymaking · Science and values · Political philosophy of science · Scientific advisory bodies

1 Introduction

Science-advisory bodies like the Scientific Advisory Group for Emergencies (SAGE) in the UK are responsible for providing evidence to inform policymakers. They are tasked with distilling complex bodies of evidence into a form policymakers can use

✉ Faik Kurtulmus
faik.kurtulmus@sabanciuniv.edu

¹ Faculty of Arts and Social Sciences, Sabanci University, Istanbul, Turkey

and, to some extent, conducting their own research. Crucially, information provided to an agent can alter their course of action even when their preferences, available options, and payoffs remain unchanged. Scientists and scientific advisory bodies also make many epistemically unforced choices during research and in presenting advice. Taken together, these two facts raise questions about whether scientific advice can be provided in a manner consistent with democratic commitments. This paper examines two challenges to ensuring that scientific research informs policymakers without undermining those commitments, and argues that addressing them requires an institutional approach.

The first challenge concerns the management of values in science. Given the ineliminable role of non-epistemic values in scientific research and its presentation (Brown, 2013; Douglas, 2009; Elliott, 2017; Elabbar, n.d.); Longino, 1990), many argue that if science-based policy is not to encroach on the public's authority and serve the public's welfare as the public understands it, it should be guided by democratically determined values (Alexandrova & Fabian, 2022; Intemann, 2015; Lusk, 2021; Pamuk, 2021; Schroeder, 2021). I argue that we should pursue this ideal *holistically*: alignment between democratically endorsed values and the values embedded in research should be sought at the level of the total evidence base available to policymakers, because the value-laden choices in individual studies can compound or offset one another. Yet even when politicians and scientists fully embrace democratic aims and the public's values are known, scientists cannot realize this ideal because of coordination problems. Individual researchers cannot foresee, let alone coordinate, how the value-laden choices in their own work will interact with those made in other studies on which policymakers depend. Consequently, institutional mechanisms are needed to centrally assess and manage value-laden decisions across the whole body of evidence presented to policymakers.

The second challenge arises when we introduce some realism regarding the motivations of politicians and members of scientific advisory bodies. When we allow that politicians and scientists can pursue their own interests, or the public interest as they rather than the public see it, they can have incentives to shape the evidence before political decision-makers. By doing so, they can make a course of action that the public would oppose appear acceptable. They can thereby undermine democratic accountability or make choices that do not pursue public welfare. Addressing this problem requires attending not only to the value-laden choices made in research, but also to the structure of the advisory process. We must adopt a *realistic* perspective that accounts for the incentives of scientists and policymakers, as well as the public's constraints. Fostering democratic accountability also calls for an institutional approach—one that attends to the strategic interactions among these actors, shaped by their incentives and constraints.

The considerations I present lay the groundwork for a solution, but they do not constitute one themselves. The first step toward solving a problem is to articulate it clearly, and that is the paper's primary aim. Only secondarily does it trace the contours of the solution we need. The paper aims to identify the trade-offs we face, the causal levers available to us, and clarify what kinds of solutions will not work. In doing so, the paper seeks to contribute to recent scholarship that emphasizes insti-

tutional, rather than individual, management of values in science (Biddle & Kukla, 2017; Contessa, 2021; Irzik & Kurtulmus, 2024; John, 2017; Pamuk, 2021; Steel, 2016; Wilholt, 2013).

One final point worth emphasizing: even though my discussion is framed around the democratic values ideal, the problem is general. Any account of the values that should be embedded in research that will guide policy will have to grapple with similar problems.

2 Thinking holistically: the site of democratic alignment

2.1 Values in science

The democratic values approach is a response to the fact that value-laden choices made at various stages of scientific research can affect the evidence ultimately presented to policymakers and, in turn, their decisions.

Since the role of values in science is familiar territory, I shall summarize briefly. The least controversial value influence, granted also by the proponents of the value-free ideal, is the choice of topics to research. These are influenced by our practical interests as well as by the past evolution of disciplines and the state of scientific knowledge (Kitcher, 2003: 63–82). More controversial, but by now generally accepted, is the role of values in inductive inferences (Douglas, 2009; Rudner, 1953). Because of the inductive gap between evidence and hypothesis, inductive inferences always carry the risk of rejecting a true hypothesis or accepting a false one. There is, of course, a third option: suspending judgment (Steel, 2016; Wilholt, 2016). Deciding which error to guard against, or whether to suspend judgment, inevitably reflects a value judgment about the consequences of each choice.

The role of values in science extends beyond inductive risk. The operationalization of variables, for example, deciding whether to discount future health losses or apply age weights when calculating disease burden, is also value-laden (Schroeder, 2021; Biddle & Kukla, 2017). As Alexandrova (2018) points out, thick concepts—concepts with both evaluative and descriptive dimensions, such as health, child development, and economic growth—inevitably incorporate value judgments too. Similarly, scientific models in fields like epidemiology, economics, and climate science involve value-laden choices about which entities to represent and how to represent them (Harvard & Winsberg, 2022). The curation of evidence, deciding what evidence to present and how to present it to policymakers, is also value-laden (Elabbar, n.d.).

By *value-laden choices*, I do not refer only to the conscious or unconscious choices of individual scientists. I am also referring to value-laden choices embodied in disciplinary conventions, such as measuring economic growth using GDP or adopting a p-value threshold of 0.05, which embed judgments about what counts as economic growth and about how to distribute inductive risk, respectively. Thus, value-laden choices include what John calls *as-if value-laden choices*: choices embedded in scientific research that are made as if scientists had adopted certain values (John forthcoming; see also Ward, 2021).

The following toy example illustrates how value-laden choices in scientific research can affect policymakers' decisions. Suppose policymakers are committed to shutting down factories that harm public health. There is a factory that emits gas X. Policymakers consult scientists about the harms of X.

Consider the following. (a) There is no research on the health effects of X. (b) There is some evidence that X is harmful, but researchers, who prioritize avoiding false positives over false negatives, deem it insufficiently certain and assert they don't know whether X is harmful. (c) There is good evidence that X negatively affects mental well-being, but scientists narrowly construe health, excluding mental health. (d) Scientists believe that X harms public health, but only report chemicals that harm more than 0.1% of the population; X harms a smaller portion of the population. In all these cases, policymakers who consult scientists will not shut down the factory in question.

Is there something problematic in these cases? Many think the answer depends on whether the scientists' value-laden choices align with the public's values. If, for instance, the public also does not wish action to be taken on factories that emit gases that harm less than 0.1% of the population, then nothing has gone wrong. *The democratic values ideal* builds on this intuition. If the values that guide scientific research are those of the public, then the value-laden nature of scientific research does not pose a problem. Even though scientists make value-laden decisions in their research, these choices do not undermine democratic rule because they reflect the public's values.¹

2.2 Values embedded in the total evidence before policymakers

Value-laden choices in scientific research raise democratic concerns because they influence policy decisions. However, it is rare for a single piece of research to, by itself, influence policy choices. More commonly, there are different pieces of research, often from various disciplines, that inform a policy decision. Thus, when implementing the democratic values ideal, or any other specific set of values, our concern should be whether *the values embedded in the total body of scientific evidence* available to policymakers align with those ideals.

It may seem that if all research is conducted in line with democratic values, then the resulting body of evidence available to policymakers will also be in line with those values. This is not the case: alignment at the level of individual studies is neither necessary nor sufficient to ensure that the values embedded in the total body of evidence reflect democratic values. Let us begin with the claim that individual-level alignment is not necessary.

It has been argued that Report 9, prepared by the Imperial College (ICL) COVID-19 Response Team led by Neil Ferguson, played a central role in shaping the UK's

¹ In an insightful paper, Thoma (2024) argues that even if the values guiding value-laden choices are democratically determined, democratic alignment may still fall short from a democratic perspective, because such choices can deprive some social groups of the epistemic resources they need for democratic deliberation. In other words, democratic alignment raises issues of distributive epistemic justice that are themselves democratically troubling. While I believe she identifies a powerful objection to the democratic values ideal, or its standard articulation, in this paper I take the ideal as given and focus instead on problems related to its implementation.

policy response in the early stages of the pandemic (Macedo & Lee, 2025, pp. 63–65). The model compared mitigation and suppression strategies by projecting cases, deaths, and intensive care unit demand relative to capacity (Ferguson et al., 2020). It excluded factors such as mental health effects, economic costs, and educational impacts. It did not represent race, income, or occupation, which limited its ability to assess how non-pharmaceutical interventions (NPIs) affected different groups (Winsberg & Harvard, 2024, p. 52). Furthermore, when faced with a choice that would lead to over- or underestimating the health harms of COVID-19, the modelers chose to overestimate them (Winsberg & Harvard, 2024, p. 55). Critics argue that these choices favored lockdowns by overestimating the harms to health and obscured distributive considerations (Harvard & Winsberg, 2022; Winsberg & Harvard, 2024).

Suppose Winsberg and Harvard are right about the value-laden choices in the ICL model, and suppose further that these values would not be democratically endorsed—for instance, because the public would also care about the social and economic costs of NPIs and about how their burdens and benefits are distributed. What follows from this? Whether the evidence available to policymakers failed to reflect democratic values depends on what other evidence they considered. A situation in which the ICL model was the sole input is very different from one in which policymakers also had access to analyses of the social costs of NPIs and of their differential effects across groups. The significance of the value-laden choices in the ICL model ultimately depends on the other evidence it was combined with. A model that does not, by itself, reflect democratically endorsed values can still contribute to a body of evidence that does.²

Now consider a case where each individual study aligns with democratically endorsed values, yet the values embedded in the total body of evidence fail to do so. Suppose a coastal city is deciding whether to invest in upgraded sea defenses in light of possible sea-level rise. Citizens favor moderate caution: they wish to guard against plausible risks but avoid overcommitting public funds.

The policy decision depends on the expected value of upgraded sea defenses, which in turn depends on two inputs: a probability estimate of damaging sea-level rise from climate scientists and an economic assessment of the damage, should it occur, from economists. Both sets of research attempt to honor citizens' preference for moderate caution. Climate scientists' probability estimates range from 10% to 20%, with a best central estimate of 15%. To avoid underestimating the risk, they report a probability 20% higher than this estimate, namely 18%. Economists' damage estimates range from \$90 million to \$160 million. Their best central estimate is \$125 million. To avoid underestimating potential harm, they likewise revise their estimate upward by 20% to \$150 million. Policymakers combining these inputs arrive at an expected damage estimate of \$27 million ($18\% \times \150 M). Since the underlying evidence supports expected values between \$9 million and \$32 million, this figure is more than 40% higher than the mean expected value. In this instance,

² It also matters which decisions the model was used to inform. A model may be entirely appropriate for planning hospital capacity yet inappropriate as the primary basis for overall pandemic policy. There is no such thing as democratically endorsed values *simpliciter*: what values a democratic public would want to see reflected in the evidence depends on the specific decision at hand and on the kinds of trade-offs that decision requires.

although each group faithfully implemented the same democratically endorsed value by being moderately cautious, when combined their research inadvertently violated it. While individual studies were moderately cautious, the overall body of evidence was overly cautious. Thus, alignment with democratic values at the individual level was not sufficient to ensure alignment across the total body of evidence.³

These two cases show that when pursuing alignment with democratic values, we need to go beyond the choices made in individual pieces of research and adopt a more holistic approach. Taken on its own, a single study may seem perfectly aligned with—or clearly at odds with—democratically endorsed values but become part of a body of evidence that has the opposite character. For this reason, in managing values in science bearing on policy, our target should be values embedded in the total body of evidence before policymakers who are using that research to answer a specific question.

Scientists face an immense coordination problem in ensuring that the values in the total body of scientific evidence aligns with democratically endorsed values. Since this aggregate level is what ultimately matters for policy, simply knowing which values are democratically endorsed is not enough. To foresee the final outcome, scientists would also need to know the other research with which their own will be combined, as well as the value-laden choices embedded in that research.

Given these coordination problems, the point at which different bodies of scientific research are brought together to inform policymakers is pivotal to the management of values. We need an institutional solution that provides centralized assessment and management of the values embedded in the total body of evidence at the point of use. Advisory bodies should be designed to monitor how evidence from different sources is combined and what value-laden choices emerge from that aggregate to ensure that the science informing policy aligns with democratic values. While this diagnosis does not constitute a full solution, it makes it clear that the solution cannot be decentralized.

Since John's recent discussion of the management of values in policy-relevant science touches on similar points, let me briefly address how my discussion relates to his. John has argued that proposals to align values in science with democratic values make three over-optimistic assumptions: "when non-epistemic values shape scientific claims, we can *easily identify* them, then *replace* them with different values; and that replacement does not entail any *epistemic loss*" (forthcoming, p. 14). John argues that rather than seeking alignment, we should treat scientific input from different disciplines as partial advice based on specific value commitments, not as all-things-considered advice. He offers an analogy familiar to us from everyday life. Your banker gives you advice as if your sole concern was your finances, and your doctor gives you advice as if your sole concern was your health. How you combine their advice is up to you. Similarly, instead of seeking full alignment with democratically endorsed values, policymakers should seek advice from "a wide variety of different scientific

³ Contessa (2021), who makes a similar argument, notes a related difficulty: research builds on prior research, allowing value choices to accumulate over time. For example, if researchers at different stages in the study of a particular chemical's harms each err on the side of caution, the result can be what Contessa calls "precautionary cascades," where the final assessment becomes excessively cautious (2021, p. 5).

committees, each providing advice based around different values” (forthcoming, p. 21). It would, then, be the task of policymakers to combine these partially aligned sources of advice to reach an all-things-considered judgment.

What the approach I’ve outlined here and John’s approach share is a commitment to looking not at individual pieces of evidence, but at the total body of evidence and the values embedded in it. However, unlike John, I do not propose giving up on alignment; instead, I specify where and how we should seek it. John’s proposal rests on different disciplines wearing their evaluative commitments on their sleeves. At a very general level, this may be the case. For instance, some disciplines like medicine may have a deontic commitment to non-maleficence, as John argues. However, this is unlikely to hold at a more fine-grained level when we consider specific inductive risk judgments or values embedded in the operationalization of variables.

Furthermore, as our example of deciding on a sea defense policy illustrates, the way inputs from different disciplines are combined does not always fit John’s proposal. In that example, economists do offer an estimate of economic costs, which seems to fit John’s proposal, but the true significance of economic costs depends on estimates of the likelihood of sea-level rise from climate scientists. In such a case, knowing the value orientations of different disciplines does not help much. It is their combination that matters.

John’s challenges to democratic alignment are forceful. As he argues, it is especially difficult to identify as-if value-laden choices. In fact, my previous argument reinforces his challenge: since the value-laden choices that emerge in the total body of evidence are not intended by any single researcher, they are particularly difficult to detect. Still, it is not impossible to identify these choices; after all, this is what much of the recent literature on values in science has been doing. One task of scientific advisory bodies will be to uncover and manage these value-laden choices, including as-if value-laden choices in the total body of evidence. Likewise, while it is sometimes impossible to revise value-laden choices without epistemic loss, sometimes it is. In both cases discussed in this section, the value-laden decisions can be revised or managed without epistemic loss—for example, by requiring additional models that examine the social and economic costs of different pandemic policies in the first case, and by downwardly adjusting the estimates provided in the second.

3 Thinking realistically: incentives

3.1 The framework

In this section and the next, we relax our earlier assumption that scientists and policymakers are solely committed to democratically endorsed goals. Recognizing that they may pursue their own interests helps illuminate critical challenges in policy-making informed by science. When thinking about democratic accountability for science-based policy, it is helpful to focus on the interactions among three actors: the scientific advisory body, the government, and members of the public. The scientific advisory body offers scientific advice to the government. The government, then, chooses which policies to implement based on its preferences and the evidence

presented to it. The people hold the government to account based on their assessment of the quality of the government's policy choices given the scientific advice it had received.

Explicitly stating the conditions for the democratic accountability of scientifically informed policy making can clarify how the incentives and constraints of these actors can lead to failures. The following three conditions are central:

1. The government, not scientists, makes policies.
2. When the government makes policies, it takes into account relevant considerations and does so in light of good evidence.
3. The public can hold the government accountable for its choices relative to the evidence it had at the time.

All of these conditions require some unpacking.

The first condition requires that the government have the ultimate authority and responsibility. While it is inevitable that scientific advisory bodies will have some influence over decisions and retain some discretion (Elabbar, *n.d.*), their preferences should not be decisive. At a minimum, when the advisory body's preferences are fixed, changes in the government's preferences should result in changes to the policy adopted in a wide range of cases. Adopting policies should be interpreted to include not just adopting new policies but also preserving existing policies when alternatives are under consideration.

The second condition has a normative and an epistemic component. The epistemic component requires that when the government needs to decide on policies that should be informed by scientific research, it does so and relies on sound scientific research.

The normative component concerns not the quality of scientific evidence but rather what the government has scientific evidence about and what it does with that evidence. Recall the considerations raised in our earlier discussion of the ICL model. Policies adopted in response to the pandemic have a range of impacts beyond health. Moreover, the benefits and burdens of these policies do not fall equally on everyone, raising questions of fairness. One objection to using the ICL model as the sole scientific input guiding policy is that these considerations and the evidence pertaining to them, which should factor into policymakers' decisions would be absent. The second condition would therefore not be satisfied if the government relied solely on the ICL model.

The normative component of the second condition should be interpreted relatively weakly, though this does not make it easy to realize. It does not require the government to assign the precise weight citizens might give to various considerations. Rather, it prohibits blatant misjudgments resulting from inadequate evidence collection or the neglect of important considerations. This condition can be understood as permitting a range of acceptable weights for different considerations. The government fails to meet this condition when it assigns weights outside this acceptable range. For instance, a pandemic policy that only considered health outcomes would fail to meet this condition.

The third condition concerns the ability of members of the public to evaluate the choices the government has made. This, of course, presupposes that the first condi-

tion is realized, and is a judgment about how well the second condition has been satisfied. It concerns whether the government did a good job in gathering scientific advice and made sound choices in light of it. The requirement is not that the government's policy choices precisely match those that members of the public would have made. This would be too demanding. We cannot assume that the members of the public have such a highly specific value function. Neither can we assume that they can precisely map their evaluations to a policy in light of the scientific evidence. Instead, this condition requires something more modest: that the public should be able to form a rough but informed assessment of government actions and hold it accountable. Crucially, the public's assessment is of the government's performance, given the scientific evidence the government had at the time, and not in light of facts discovered later that may bear on policy decisions. Colloquially, we may sum up this condition as follows: 'If the government messes up when it should have known better, the public should be able to see this and hold the government accountable'.

These conditions for democratic accountability are difficult to realize. If all relevant evidence could be presented to policymakers, or if the choice of what to present involved no discretion, meeting these conditions would be relatively straightforward. However, presenting *every* piece of evidence, even if feasible, would be unhelpful (Elabbar, n.d.; Pamuk, 2021, pp. 70–5). Some selection is inevitable. The evidence policymakers use is produced and filtered through a social process shaped by the incentives and values of those involved. Both scientists and governments can shape the body of evidence informing policy choices in ways that promote some options over others. Ensuring democratic accountability by realizing these three conditions is therefore a problem of institutional design. It requires attending to these actors' incentives to shape the evidence on which policy decisions rely, the means available to them for doing so, and the limits of the public's capacity to provide effective oversight. This section will present the challenges and the next will look at how they can be addressed institutionally.

3.2 Scientists

The first condition, "The government, not scientists, decides policy", is violated when scientists dictate policy without authorization. To assess when this happens, we need to understand how and why scientists might do so. We begin with the *how*.

One way scientists can persuade the government to adopt their favored policies is by providing false evidence. I shall set aside this possibility and assume that scientists are truthful. Given the substantive costs of being caught lying, this is a reasonable assumption. Furthermore, making this assumption helps us appreciate the difficulty of designing a system of science-based policymaking that attains the conditions we have identified.

Just as the values embedded in the total evidence before policymakers can differ from those embedded in its parts, the conclusions one draws from the total evidence can differ from those one draws from a subset.⁴ Thus scientists can shape policy by offering biased evidence. One mechanism is what Weatherall et al. call *selective shar-*

⁴ I am grateful to one of the referees for this journal for helping me note this similarity.

ing in their model of the Tobacco Strategy to mislead the public and policymakers about the harms of tobacco (Weatherall et al., 2020; Oreskes & Conway, 2011). Here is how selective sharing works. Suppose that whether p or not p is an open scientific question, and p is, in fact, true. Assume that this is a difficult probabilistic problem. Since empirical research on complex, probabilistic issues often produces findings that can support different conclusions, some studies may suggest that p is false, even though the overall body of evidence indicates that p is true (Weatherall et al., 2020, p. 1159). By selectively sharing only studies that suggest p is false, scientists can get the government to believe it is false or raise doubts about its truth. They can thereby get the government to enact policies that it would enact if it believed that p was false or uncertain.

A second way scientists can dictate policy while truthfully informing policymakers is through what I shall call *selective preservation of ignorance*. Selective sharing works by providing biased evidence that leads policymakers to arrive at false beliefs about a specific claim. Selective preservation of ignorance, by contrast, works by informing its target about some relevant issues while strategically choosing not to inform them about others, thereby producing a particular mix of knowledge and ignorance that pushes the target toward desired policy conclusions.

In cases of selective sharing, there is a specific question at issue—whether p or not—and evidence is presented in a way that misleads the audience about the truth value of p . By contrast, selective preservation of ignorance arises when a broader question is under consideration, a question that in turn implies further, relevant sub-questions. The audience fails to raise some of these sub-questions either because it does not see their relevance or because it wrongly assumes that it already knows the answer. Someone engages in selective preservation of ignorance by truthfully answering the questions that are raised or addressing some of the implied questions while deliberately omitting to address others. In doing so, they ensure that the audience's uncorrected ignorance, combined with the new information they receive, leads it to a particular conclusion about the broader issue.

To see how selective preservation of ignorance works, consider the following. Imagine a city government considering whether to build a new network of parks or bike lanes to improve public health. The two options offer comparable health benefits. However, unbeknownst to the city government, both also have non-health benefits. Parks offer greater economic benefits; bike lanes offer greater environmental benefits. The scientific advisors consider bike lanes the better policy and believe the government tends to undervalue environmental benefits. While they are required to report their estimates of the health benefits, they can choose whether to disclose the non-health benefits of the two options. In this case, by reporting only the environmental benefits and omitting the economic benefits of the two options, they can selectively preserve the government's ignorance and steer its decision toward bike lanes.

Another way in which scientists can shape policy is, as we have already noted, through the role of values in science. It is useful here to distinguish between value-laden choices at the point of use and those made in prior research, which are later incorporated into the evidence presented to policymakers. The value-laden choices in prior research are shaped by many researchers and are often made without direct awareness of the policy questions they will be applied to. Consequently, even though

these choices influence policy, this influence is usually unintentional and in ways that researchers can rarely foresee. This is even more so for as-if value choices. The value-laden choices at the point of use, either in new research conducted for specific policy questions, such as the ICL model, or the management of values in the total body of evidence discussed in the previous section, have a much more direct bearing on policy choices. For this reason, they can be made much more strategically.⁵

There is, therefore, an important difference between the questions raised by value-laden choices in prior research and the influence wielded at the advisory stage. This difference can be expressed using Pettit's distinction between influence and control. One has *influence* over a process if one can "make a difference in how the process evolves" (Pettit, 2014, p. 121). Control is a stronger notion. When one has *control* over a process, one not only has influence but one's influence gives direction to the outcome. Pettit offers the following example to illustrate the distinction. Someone pretending to be a traffic police officer and giving hand signals at a busy intersection would create confusion and thereby exert influence over traffic. A real officer, however, would be able to direct the cars with their hand motions; they would have control. Since it is difficult to see how one's research will be combined with other research or how it will be used, individual researchers will enjoy influence but will usually lack control. They will not be in a position to direct processes where their research is used as evidence. However, the scientific advisory body that brings together different pieces of research and oversees the values embedded in the total body of evidence before policymakers can have not only influence but also control. This is one crucial reason why advisory bodies are of particular democratic concern.

We have seen that scientists can impact government policy without lying or fabricating evidence through selective sharing of evidence, selective preservation of ignorance, and value-laden choices, particularly those at the advisory stage. The next question is: What incentives might scientists have to offer evidence favoring one policy over another?

In the first place, they may believe one policy to be better than another. When they are convinced that the government is unlikely to adopt the "correct" policy if not nudged by scientists, they may be tempted to resort to these measures. Another possibility is that scientists do not seek to push their favored policy but rather to correct what they perceive as the government's biases. If scientists think the government is prone to neglect some costs when choosing policies, they might take steps to address this perceived neglect.

A third possibility worth considering is the role of fear: fear of being held responsible for unpopular policies or for choices that go bad, and fear of going against a social consensus. Neither is negligible in emergencies such as COVID. Streicher and Broadbent put this well in their discussion of whether scientists are blameworthy for recommending lockdown policies to low-income countries:

⁵ The arguments of the previous section suggest that there is a limit to how much influence new research conducted for policy can have at this stage. Since that research will also be combined with other research, it is difficult for scientists to determine the values embedded in the total body of evidence individually. However, even if they cannot determine the values embedded and, for instance, ensure that the values in the aggregate evidence are optimally cautious by their light, they can make choices that favor a cautious policy.

“We suspect that scientists, who are humans first and foremost, will not readily resist a general sentiment in a community to which they belong... An epidemiologist who would speak up against the general sentiment at that time was a brave one indeed, and those who did have suffered, even where their scientific credentials were impeccable and where there [sic] arguments were at least as cogent as many other arguments circulating at the time.” (2022, p. 100083).

Unfortunately, the circumstances in which scientific input is most needed—those where the broader or longer-term consequences of policy are difficult to predict or observe—are also the circumstances in which fear makes it hardest for scientists to raise such concerns publicly. During COVID-19, for example, deaths were immediate and emotionally salient, while many of the harms of lockdowns were diffuse, indirect, or not yet visible. In such an atmosphere, it was difficult for scientists to speak openly about these less salient costs, even though their assessment was especially important for understanding the full range of policy consequences. Continuing with the theme that scientists are humans, subject to the same weaknesses and vices as all other human beings, one motivation for shaping policy is greed. Greed can shape the science that informs policy in its early stages by funding research that promotes specific policies, or in the advisory stage through corruption.

3.3 The government

Just as recent philosophy of science is instructive in showing how evidence can be intentionally or unintentionally shaped to favor certain policies, political science is instructive in understanding the incentives politicians have to neglect certain considerations, collect substandard evidence, or, more generally, game the evidence before them. As we shall see, these reasons are closely connected to our third condition: the public’s ability to hold the government to account.

The government can have two relevant incentives to shape the evidence before it. First, it can favor a policy the public does not support. However, if the government can shape the evidence before it to make that policy appear best supported by existing evidence, it can adopt an unpopular policy and avoid accountability for doing so.

Second, the government can seek to avoid blame for its policy choice that can have negative outcomes. In his book, *The Blame Game*, the political scientist Christopher Hood argues that blame avoidance is a forceful motive that “shapes the conduct of officeholders, the architecture of organizations, and their operating routines and policies” (Hood, 2011, p. 5). This perspective sheds light on governments’ use of scientific advice. One mechanism to avoid blame for negative outcomes is to reduce the government’s perceived or actual agency (Hood, 2011, pp. 16–22). If the government can plausibly argue that the evidence before it dictated a specific course of action, thereby reducing its perceived agency, it can avoid blame. The claim, made by many governments, that “They were following the science” makes sense from this perspective. To the extent that a government can convince the public that science dictated its policies and that it did not have much choice in the matter, it can avoid blame.

During COVID, scientists did take note of this rhetorical strategy. Some “began to fear that Ministers were shifting the accountability for hard decisions onto them”

(House of Commons Science and Technology Committee, 2021, p. 28). Such a strategy can affect the scientific advisory process because the government's attempt to shift blame interacts with scientists' fear of public backlash. When scientists believe they will be blamed for unpopular policies, they may become reluctant to provide evidence supporting such policies, even when those policies would better serve the public interest.

Thus, the government's claim to be following the science has two negative effects. First, it undermines accountability because "for accountability to work well, responsibilities must be clearly defined so that it is apparent whom to blame for a transgression" (Persson et al., 1997, p. 1168). In addition, it distorts the scientific advisory process by incentivizing scientists to provide risk-averse policy advice or to favor policies whose harms are relatively difficult to observe or attribute to specific policy choices, even when those policies may be worse overall.

The government can create an evidence base that is misleading or that omits certain considerations through its discretion over the scientific advisory process. A government able to select the experts who advise it is more likely to receive scientific advice that makes its preferred policies appear scientifically supported. Likewise, if scientists only answer the questions posed by the government and cannot provide information on their own initiative, the government can promote its favored policies by asking the right questions, thereby engaging in selective preservation of ignorance vis-à-vis the public. These two techniques work even if the entire advisory process is open and shared with the public because they function by shaping the evidence that informs policy.

3.4 Citizens

The previous discussion of how scientists can shape science-based policies and how the government can game the advisory process shows how difficult it is for citizens to hold their government accountable. To properly hold their government accountable, citizens would need to know that it is their government, not scientists, that has determined policies and that their government did not game the advisory process. However, it is difficult for citizens to verify these conditions on a case-by-case basis.

As our discussion has shown, the requirement that the public be able to hold the government to account based on its performance relative to the evidence available to it introduces a number of complications, too. Holding the government accountable in this way is informationally demanding on citizens. Why, then, insist on this? The answer is that when the public holds the government accountable without access to the evidence the government had, democratic governments tend to adopt bad policies in response to voters' misinformed preferences. Suppose a particular policy is optimal based on the evidence available to the government and likely to bring about the outcomes the public prefers. If the public does not have access to this evidence and instead relies on incomplete or inaccurate information, this policy may not appear as the best choice to them. In such cases, a government that wants to be reelected would have incentives to "pander" to voters' uninformed preferences and implement the inferior policy (Canes-Wrone et al., 2001; Maskin & Tirole, 2004). These mismatches—between

what the public would prefer if better informed and what they prefer based on existing information—are more likely to occur in cases where scientific input is central to formulating policy options and evaluating their costs and benefits. This creates a bias in policymaking. Costs and benefits of policies that are easier for the public to observe, in terms of magnitude or causal attribution, receive more weight. To counteract such pandering, citizens should have access to the government's evidence and assess its decisions relative to that evidence.

What is needed, therefore, is an institutional framework that consistently ensures these conditions for democratic accountability are met and is trusted by the citizens to do so. The scientific advisory process should be designed in such a way that scientists cannot shape policies, and the government cannot game the advice it receives. Citizens should, in other words, be able to trust that the scientific advisory process is designed in a way that makes such abuses unlikely. Their ability to evaluate the government's performance should therefore be supported by an institutional framework.

4 Thinking institutionally: causal levers and trade-offs

In this section, I discuss some institutional measures that can help address the problems we encountered in the previous section. My aim is not to argue for a specific set of measures but to explore their advantages and disadvantages and demonstrate the kind of institutional thinking we need to engage in.

The discussion that follows will largely concern measures to counteract the incentives that may lead scientists and politicians to act against the public interest. But we should not assume such incentives always prevail; many politicians and scientists are motivated to serve the public good. Institutional design should reflect both possibilities: it should include mechanisms to select public-spirited individuals⁶ and ethical codes to guide them⁷ while also incorporating safeguards against incentive problems.

An obvious solution to the problem of selective sharing and selective preservation of ignorance would be to limit scientists' discretion by requiring them only to answer questions put to them by the government. However, this solution would create a different problem. It would give the government a free hand in shaping the evidence before it through the choice of questions and the selection of evidence.

Predetermined rules governing evidence collection would simultaneously reduce the discretion scientists and the government enjoy. For instance, assessing various types of costs, such as health and economic impacts, could be made mandatory for all policies under consideration. Returning to the ICL model and the issues it raised, one could establish rules requiring models of NPIs to account for their impacts across various social groups and to model the non-health costs and benefits of different NPIs. Such predetermined rules would make selective

⁶ For a discussion of the importance of selection in addition to sanctions, and its distinctive benefits see (Mansbridge, 2009).

⁷ For a helpful account see (Schroeder, 2022a).

preservation of ignorance more difficult. The collection and presentation of disaggregated data for policy purposes, which would enable an assessment of the distribution of burdens and benefits of policies across different groups, can also be made mandatory (Cobham, 2020). By incorporating rules governing the kind and range of studies that advisory bodies should incorporate into their reports, for instance, requiring the inclusion of systematic reviews, the problem of selective sharing can be ameliorated.

As with any rule-based system, this proposal has the drawback of limited flexibility. It is difficult to determine the relevant costs and benefits of policies that need to be considered beforehand. It is also difficult to set the proper limits on which kinds of research should be incorporated into the reports of advisory bodies. However, such rules can be used to establish a baseline of evidence that can then be supplemented with additional evidence.

Introducing two different sets of actors into the advisory framework can retain flexibility while reducing the ability of either scientists or the government to shape the evidence before them. One approach would be to include the opposition in the advisory process. If opposition members also have a say in selecting scientists for advisory committees, it would be more difficult for the government to choose scientists who will toe its line. If the opposition can also ask questions to the advisory committee, they can introduce relevant considerations that the government may prefer to keep hidden. Their questions can also clarify the scope of choice that the government has in terms of policy options.

There are two drawbacks to involving the opposition in the selection of scientific advisors and the questions put to them. When we need to act fast, this can be time-consuming. Moreover, if the opposition is uncooperative, it may use its powers to influence the advisory process to make the government look bad rather than to improve decision-making. This is, it should be noted, a concern that applies more to the power to ask questions than having a say in the selection of experts. These disadvantages are minor compared to the potential benefits. That said, the overall balance depends on the specific political system and society in question.

A second set of actors we can introduce is another advisory body. Having not one but two advisory bodies can reduce scientists' discretion and ability to influence policy. In addition to limiting scientists' capacity to shape policy, two advisory bodies would make corruption more difficult to carry out. It would also make it harder for the government to shift responsibility onto scientists, since its agency in choosing which advice to rely on would be apparent. Finally, having two advisory bodies would diffuse responsibility among their members, thereby reducing the distortions to the advisory process discussed above that arise from fear.

Having two advisory bodies would also have epistemically beneficial consequences. It would create a larger body of experts, and their disagreements would expose errors. It would also act as a precaution against information cascades and groupthink. Of course, the two groups can each individually be subject to these phenomena. Still, disagreements between the bodies can reveal where group-level cognitive biases might be influencing their judgments.

Two advisory bodies would have serious drawbacks. They would increase the public's informational load and make holding the government accountable more difficult. When the two advisory bodies disagree, members of the public would need to decide not only whether the government did an adequate job given the available evidence but also which advisory body it should have relied on.

A better alternative that preserves pluralism in the advice the government receives and offers similar epistemic benefits would be to institutionalize the writing of dissenting opinions, much like the practice of the Supreme Court in the US (Pamuk, 2021, pp. 85–88, see also Dang, 2024). Dissenting opinions would not only register dissent but also provide justifications for it. They would include discussions of how sufficient, significant, and relevant the evidence is and what evidence was left out. By doing so, they would offer the government and the public epistemic resources to assess and question the majority view. Furthermore, it would mean that there is less pressure to provide a consensus view, which tends to converge on the least common denominator, thereby leaving out contested but possibly useful evidence. Since dissenting opinions would exist alongside the majority opinion, which the government has presumptive reasons to defer to, the public would not need to decide which advisory body the government should have followed.

I have argued that holding the government accountable requires the public to have access to relevant information. Transparency of the advisory process has been commonly put forth as a critical requirement for enabling such accountability. For instance, one of the persistent criticisms of the operation of the SAGE in the early days of the pandemic was a lack of transparency. The previous discussion raises two considerations that can inform our thinking about transparency and complicate this assessment.

Consider scientists' incentives. Transparency is commonly interpreted as requiring that both the evidence provided to the government and the identities of people on advisory committees be shared with the public. For instance, the Science and Technology Committee report, which aimed to evaluate the British government's handling of the pandemic, criticizes the secrecy surrounding SAGE's membership and the failure to publish SAGE's minutes in the same breath (2021, p. 24). Our discussion suggests that the two should be treated differently.

Revealing the membership of advisory committees limits the government's ability to select experts who will tend to support its favored policies. It also creates incentives for it to select competent experts. However, it has a disadvantage: It opens members of scientific advisory committees to public pressure. If, as I argued, members of scientific committees might fear being blamed by the public for the policies adopted in light of their advice, such transparency can be detrimental.

The reasons for transparency about the membership of scientific advisory bodies are weaker when the opposition also has a say in the selection process. There is no longer the worry that the government will select experts who will tend to offer evidence supporting the policies it favors. In such circumstances, a lack of transparency about membership can make it easier for scientists to voice unpopular opinions.

When the government has incentives to shift blame onto scientific advisory committees, then a lack of transparency about their membership can help counter some of the resulting distortions. Consider again fear as a motivation. Suppose one policy option has more salient and easily observable harmful effects than another, even though it is better overall. Scientists who fear that the government will shift responsibility for policy choices onto them would have reason to recommend the second, inferior option to avoid being blamed for visible harms. However, when the identities of committee members are not disclosed, scientists can recommend the better policy without fear.

This point about transparency highlights a key consideration in designing institutions for scientific advice: we should not view individual elements in isolation and instead consider how they interact. Transparency about membership is beneficial when advisors are selected solely by the government. When the opposition is also involved in the selection process, however, the advantages of keeping membership confidential outweigh the costs.

A second consideration concerns the constraints faced by members of the public. When an institutional framework that incorporates the previous considerations is in place, citizens can take it for granted that scientific evidence has not been gamed by either the government or scientists. Their task becomes evaluating the government's policy choices relative to its evidence. Transparency about the evidence presented to the government is not sufficient for accountability: members of the public cannot be expected to sift through all the evidence available to policymakers. There should be intermediary mechanisms that select salient evidence for members of the public. This process should be guided by what the public already knows or can be expected to know. I argued that the government could deliver public policies that the public would not favor if it were better informed, since some of the costs or benefits of policies are easier to observe than others. The process of informing the public should be designed to counteract such biases and take extra care to inform them of the less apparent costs and benefits.

If we focus solely on individual actors rather than institutional structures, democratic accountability becomes unattainable. No individual member of the public can consider the evidence before the government and decide for themselves whether it did a good job. Accountability becomes feasible only within an institutional framework. In addition to the press, opposition parties and parliamentary committees help members of the public assess whether the government has used the evidence effectively, thereby improving accountability.

5 Conclusion

This paper has offered two reasons for a thoroughly institutional approach to science-based policymaking in representative democracies. The first arises from the need to think holistically about the interactions among the value-laden choices embedded in different pieces of research informing policy. Because these interactions cannot be managed in a decentralized manner, centralized institutional management of values

at the point of use is necessary. The second reason becomes evident once we realistically consider the incentives and constraints facing scientific advisory bodies, governments, and the public. Only institutional solutions can help overcome or mitigate these incentive problems and make the public's informational load manageable. As we have seen, doing so requires thinking about the design of the advisory process but also attending to broader institutions such as the press.

In seeking to make scientific advice compatible with democratic ideals, we must move beyond an individualistic perspective and consider the institutional frameworks that can support these ideals. We need a politically attuned approach that accounts for the limits of individual action and addresses the structural conditions necessary for democratic accountability.⁸ This will require drawing on insights from philosophy of science, political science, and political philosophy. Even then, the ideal may prove impossible to realize, prompting us either to rethink the ideal itself or to look beyond the confines of representative democracy by incorporating elements of direct democratic participation.

Acknowledgements An earlier version of this paper was presented at the Ethics and Epistemology of Expertise workshop at Cambridge, and I am grateful to the participants for their thoughtful feedback. I would also like to thank Gürol Irzik and the anonymous referees for this journal for their valuable comments and suggestions. The paper was completed during a visit to SOCRATES at Leibniz University Hannover, which provided an excellent environment for writing and discussion, and was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project 470816212/KFG43.

Author contributions N/A

Funding Open access funding provided by the Scientific and Technological Research Council of Türkiye (TÜBİTAK). Funding: Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)– Project 470816212/KFG43.

Data Availability N/A

Ethical approval N/A

Informed consent N/A

Competing interests N/A

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

⁸ For a related discussion see Schroeder, 2022b.

References

- Alexandrova, A. (2018). Can the science of well-being be objective? *The British Journal for the Philosophy of Science*, 69(2), 421–445.
- Alexandrova, A., & Fabian, M. (2022). Democratising measurement: Or why thick concepts call for coproduction. *European Journal for Philosophy of Science*, 12(1), 1–23.
- Biddle, J. B., & Kukla, R. (2017). The Geography of Epistemic Risk. In K. C. Elliott & T. Richards (Eds.), *Exploring Inductive Risk: Case Studies of Values in Science* (pp. 215–237). Oxford University Press.
- Brown, M. J. (2013). Values in science beyond underdetermination and inductive risk. *Philosophy of Science*, 80(5), 829–839.
- Canes-Wrone, B., Herron, M. C., & Shotts, K. W. (2001). Leadership and pandering: A theory of executive policymaking. *American Journal of Political Science*, 45(3), 532.
- Cobham, A. (2020). *The Uncounted*. Polity Press.
- Contessa, G. (2021). On the mitigation of inductive risk. *European Journal for Philosophy of Science*, 11(3), Article 64.
- Dang, H. (2024). Minority reports: Registering dissent in science. *Philosophy of Science*, 91(5), 1169–1178.
- Douglas, H. E. (2009). *Science, Policy, and the Value-Free Ideal*. University of Pittsburgh Press.
- Elabbar, A. (n.d.). *The curatorial view of assessment and the ethics of scientific advice: Beyond decisional autonomy towards distributive epistemic justice*. Accessed 22 2024. http://iuhpst.org/media/pdf/Elabbar_IUHPST_Prize_2023.pdf
- Elliott, K. C. (2017). *A Tapestry of Values: An Introduction to Values in Science*. Oxford University Press. Email address faik.kurtulmus@sabanciuniv.edu.
- Ferguson, N., Laydon, D., Nedjati, G. G., Imai, N., Ainslie, K., Baguelin, M. (2020). *Report 9: Impact of Non-pharmaceutical Interventions (NPIs) To Reduce COVID19 Mortality and Healthcare Demand*. <https://doi.org/10.25561/77482>
- Harvard, S., & Winsberg, E. (2022). The epistemic risk in representation. *Kennedy Institute of Ethics Journal*, 32(1), 1–31.
- Hood, C. (2011). *The Blame Game: Spin, Bureaucracy, and Self-Preservation in Government*. Princeton University Press.
- House of Commons Science and Technology Committee (2021). *The UK Response to COVID-19: Use of Scientific Advice*. UK Parliament. <https://committees.parliament.uk/publications/5868/documents/66635/default/>. Accessed 25 October 2024.
- Intemann, K. (2015). Distinguishing between legitimate and illegitimate values in climate modeling. *European Journal for Philosophy of Science*, 5(2), 217–232.
- Irzik, G., & Kurtulmus, F. (2024). Distributive epistemic justice in science. *The British Journal for the Philosophy of Science*, 75(2), 325–345.
- John, S. (2017). From social values to P-Values: The social epistemology of the intergovernmental panel on climate change. *Journal of Applied Philosophy*, 34(2), 157–171.
- John, S. (forthcoming) (Ed.). Weber’s elephant: Rethinking scientific advice. *The British Journal for the Philosophy of Science*.
- Kitcher, P. (2003). *Science, Truth, and Democracy*. Oxford University Press.
- Longino, H. E. (1990). *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton University Press.
- Lusk, G. (2021). Does democracy require value-neutral science? Analyzing the legitimacy of scientific information in the political sphere. *Studies in History and Philosophy of Science, Part A*, 90, 102–110.
- Macedo, S., & Lee, F. E. (2025). *In Covid’s wake: How our politics failed us*. Princeton University Press.
- Mansbridge, J. (2009). A selection model of political representation. *Journal of Political Philosophy*, 17(4), 369–398.
- Maskin, E., & Tirole, J. (2004). The politician and the judge: Accountability in government. *American Economic Review*, 94(4), 1034–1054.
- Oreskes, N., & Conway, E. M. (2011). *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*. Bloomsbury.
- Pamuk, Z. (2021). *Politics and Expertise: How to Use Science in a Democratic Society*. Princeton University Press.
- Persson, T., Roland, G., & Tabellini, G. (1997). Separation of powers and political accountability. *The Quarterly Journal of Economics*, 112(4), 1163–1202.
- Pettit, P. (2014). *Just Freedom: A Moral Compass for a Complex World*. W.W. Norton & Company.

- Rudner, R. (1953). The scientist qua scientist makes value judgments. *Philosophy of Science*, 20(1), 1–6.
- Schroeder, S. A. (2021). Democratic values: A better foundation for public trust in science. *The British Journal for the Philosophy of Science*, 72(2), 545–562.
- Schroeder, S. A. (2022a). An ethical framework for presenting scientific results to policy-makers. *Kennedy Institute of Ethics Journal*, 32(1), 33–67.
- Schroeder, S. A. (2022b). Thinking about values in science: Ethical versus political approaches. *Canadian Journal of Philosophy*, 52(3), 246–255.
- Steel, D. (2016). Climate change and second-order uncertainty: Defending a Generalized, Normative, and structural argument from inductive risk. *Perspectives on Science*, 24(6), 696–721.
- Streicher, P., & Broadbent, A. (2022). Blame or reflection? Response to Winsberg. *Global Epidemiology*, 4, Article 100083.
- Thoma, J. (2024). Social science, policy and democracy. *Philosophy & Public Affairs*, 52(1), 5–41.
- Ward, Z. B. (2021). On value-laden science. *Studies in History and Philosophy of Science Part A*, 85, 54–62.
- Weatherall, J. O., O'Connor, C., & Bruner, J. P. (2020). How to beat science and influence people: Policymakers and propaganda in epistemic networks. *The British Journal for the Philosophy of Science*, 71(4), 1157–1186.
- Wilholt, T. (2013). Epistemic trust in science. *The British Journal for the Philosophy of Science*, 64(2), 233–253.
- Wilholt, T. (2016). Collaborative Research, Scientific Communities, and the Social Diffusion of Trustworthiness. In M. S. Brady & M. Fricker (Eds.), *The Epistemic Life of Groups: Essays in the Epistemology of Collectives* (pp. 218–233). Oxford University Press.
- Winsberg, E. B., & Harvard, S. (2024). *Scientific Models and Decision Making*. Cambridge University Press.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.