




Science advice: making credences accurate

Simon Blessenohl¹ · Deniz Sarikaya^{2,3} 

Received: 29 June 2021 / Accepted: 15 February 2022 / Published online: 15 April 2022
© The Author(s) 2022

Abstract

Policy-makers often rely on scientists to inform their decisions. When advising policy-makers, what should scientists say? One view says that scientists ought to say what they have a high credence in. Another view says that scientists ought to say what they expect to lead to good policy outcomes. We explore a third view: scientists ought to say what they expect to make the policy-makers' credences accurate.

Keywords Scientific policy advice · Uncertainty · Values in science · General philosophy of science

1 Introduction

Many policy decisions involve difficult empirical questions. For example, the decision whether to regulate the emission of a certain pollutant involves the empirical question whether the pollutant is toxic. Similarly, the decision whether to prohibit the construction of new nuclear power plants involves the empirical question how safe nuclear reactors are. Since policy-makers need to make such decisions but cannot be experts in all relevant domains, states have developed mechanisms to bring policy-makers in contact with scientists. For example, a policy-maker might invite a scientist to an expert briefing in which the scientist reports on the toxicity of a pollutant.

This paper explores views about what scientists should say when they advise policy-makers. One view holds that scientists should say what they have a high credence in. For example, they should *not* assert 'the pollutant is toxic' if they only have a credence of 0.7 that the pollutant is toxic. Rather, they should make their uncertainty explicit to say something weaker that they have a high credence in, such as 'it is likely that the

✉ Deniz Sarikaya
Deniz.Sarikaya@uni-hamburg.de ; Deniz_Sarikaya@live.de

¹ USC School of Philosophy, 3709 Trousdale Pkwy., MHP 113, Los Angeles 90089, CA, USA

² Philosophisches Seminar, University of Hamburg, Hamburg, Germany

³ Department for History, Archeology, Arts, Philosophy and Ethics (HARP), Centre for Logic and Philosophy of Science, Vrije Universiteit Brussel, Pleinlaan 2, Room 5B425, Brussels 1050, Belgium

pollutant is toxic'.¹ Another view holds that scientists should say what they expect to have the best policy consequences.² For example, scientists should not make their uncertainty explicit if they know that this would cause policy-makers to leave the emission of the pollutant unregulated and scientists think that regulation would be desirable given the likelihood of the pollutant being toxic and the social costs of failing to regulate emission if it is toxic.

Neither of these two views is fundamentally concerned with the policy-maker. The first view is fundamentally concerned with the scientist communicating precisely what they have a high confidence in. The second view is fundamentally concerned with policy consequences and hence only derivatively concerned with policy-makers, insofar as it is through them that science advice exerts a causal influence on policy decisions. While proponents of these views discuss the role and characteristics of policy-makers, it is somewhat surprising that their views are not fundamentally concerned with the policy-maker. After all, it is natural to think that the point of science advice is to make policy-makers more informed. This suggests considering a third view, which is fundamentally concerned with the epistemic benefit of the policy-maker. According to this view, scientists should say what they expect to make the policy-makers' credences most accurate. That is, if a scientist has a credence of 0.7 in a pollutant being toxic, taking that to be the accuracy-maximizing credence, then they should say whatever brings the policy-maker's credence close to 0.7. If this requires not making uncertainty explicit or saying things they know will bring about suboptimal policy consequences, so be it.

For ease of reference, let us state the three views as follows:

- (Precision)** Advising scientists ought to say what they have a high credence in.
- (Policy)** Advising scientists ought to say what maximizes the expected value of the policy consequences of their utterance.
- (Accuracy)** Advising scientists ought to say what maximizes the expected accuracy of their addressees' credences.

Before we start discussing these views, we must clarify some central concepts. First, *advising scientists* are scientists who communicate their findings to policy-makers.³ For the purposes of this paper, a *policy-maker* is someone with the formal power to make laws. For instance, a policy-maker might be an elected member of the legislative branch, such as a member of the U.S. Congress. Alternatively, a policy-maker might be an employee who writes secondary legislation at an executive agency, such as the

¹ See Betz (2013) and further references in footnote 28.

² See Steele (2012) and further references in footnote 14.

³ Science advice can occur in various institutional contexts. We focus on cases in which policy-makers commission scientists to report on particular questions. Another type of science advice occurs when organizations such as the Parliamentary Office of Science and Technology of the U.K. proactively research important technologies and write reports to raise awareness for certain issues among policy-makers. The goals and norms might be different for this form of science advice than for the form we discuss. For an overview of the various institutions that facilitate science advice, see Weingart and Lentsch (2009). Note also that there are processes other than science advice through which scientists might influence policy-making. Christiano (2012), for example, construes the role of scientific experts as providing a pool of theories on the basis of which policy-makers can develop legislation.

Environmental Protection Agency. To keep our discussion focused, we only consider norms for science advice in democracies.⁴

Second, the *expectations* in (Accuracy) and (Policy) are taken with respect to the scientist's credences. Thus, (Accuracy) says that scientists ought to say what makes their addressees' credences accurate according to an assessment of accuracy based on the advising scientist's credences.

Third, what do we mean by 'ought'? Like other activities, science advice might be governed by different kinds of norms which employ different senses of 'ought'. Maybe there is a moral norm, which specifies what scientists morally ought to say, and an epistemic norm, which specifies what scientists epistemically ought to say. Then, the above norms might not be mutually exclusive: they might simply employ different senses of 'ought'. But we read all these norms in a specific sense of 'ought': the moral sense. We ask what scientists morally ought to say when they communicate their findings to policy-makers.

The three views are simplified versions of more plausible views. In particular, the views say that advising scientists should *always* be guided by one particular concern. This is implausible. To give a simple example, in a case in which the scientist would be killed if they followed one of the three norms, it is likely to be false that the scientist ought to follow any of the three norms. But instead of making the views more complicated by introducing qualifications to deal with pathological cases, we hereby restrict the scope of our discussion to everyday cases of science advice that lack such exotic features. We are interested in what should guide scientists in such cases: precision, policy, or accuracy.⁵

Scholars have proposed a wide variety of norms for science advice. We will say more about norms similar to (Policy) and (Precision) in Sects. 3.1 and 3.2, respectively. Here, we want to briefly position (Accuracy) relative to some other views that are also fundamentally concerned with the policy-maker.

The underlying idea of focusing on the epistemic benefit of science advice to the policy-maker is shared by John (2018). In particular, he argues that, at least in some cases, scientists ought to make those "assertion which would secure uptake of some claim which is in the [addressees'] epistemic interest to believe" (p. 83). This claim is very much in line with (Accuracy). That said, John does not appeal to the notion of accurate credences and might not endorse making the maximization of such a notion of epistemic benefit the sole focus of advising scientists, as (Accuracy) does.

The norm (Accuracy) is also closely related to norms that are formulated in terms of a (partially) non-epistemic benefit to the recipients of science advice. In particular, some scholars argue that scientists ought to promote the autonomy of their addressees. For example, Elliott (2006, p. 639) suggests that scientists should "promote the ability of those who use scientific information to make autonomous decisions that accord with

⁴ For a discussion of science advice in non-democratic societies, see, for instance, John (2019).

⁵ The norm (Precision) is simplified in a further way: having a high credence in p cannot be a sufficient condition for it to be the case that one ought to say p . Otherwise, the norm would recommend saying trivial things such as 'the pollutant is toxic or not'. Since our discussion of this norm in the rest of the paper only criticizes high credences as a necessary condition for being permitted to say something, we did not complicate matters by turning it into a plausible sufficient condition.

their own beliefs and values”. In the context of advising the public, Resnik (2001, p. 147) argues

that HCEs [health-care experts] should follow neither a purely objective approach nor a paternalistic approach to communications with the public. [...] HCEs should provide lay people with the information and advice they need to make sound decisions.

While these views might not focus on advice to policy-makers and appeal to different notions of benefit, the focus on benefiting recipients of science advice suggests to think of them as being in the same family of norms as (Accuracy).

There is a further family of views that is in a different way concerned with the recipients of science advice. These views hold that scientists should defer to their addressees’ moral judgments when such judgments are needed to decide how to communicate empirical findings. For instance, Irzik and Kurtulmus (2019), Parker and Lusk (2019), and Schroeder (2020) argue, respectively, that scientists should consult the values of the public, the values of the decision-makers, and ‘democratic’ values held by the public and its representatives when they need to make value judgments to decide how to communicate scientific results.⁶ Such views share an important property with (Accuracy): they recommend that scientists adapt how they communicate to the specific addressees that they advise.⁷ However, those views are in another respect dissimilar to (Accuracy): they assume that scientists must rely on value judgments—albeit not their own—to decide how to communicate their findings. In contrast, (Accuracy) denies that scientists need to rely on such value judgments for communication decisions.⁸ Instead, scientists can usually rely on empirical judgments about what makes their addressees’ credences accurate, as we explain in more detail in Sect. 3.1.

The main aim of this paper is to provide a detailed discussion of the merits and demerits of (Accuracy) relative to (Policy) and (Precision). Some of the arguments we discuss would also apply if we considered a different recipient-focused view than (Accuracy), in particular if that view also centered on the epistemic benefit to the policy-maker. Hence, even if one is skeptical of (Accuracy), we hope that one will still find the subsequent discussion useful as an evaluation of the general idea of letting science advice be guided by a concern for the epistemic benefit to the policy-maker.

In the next section, we outline some ways in which (Accuracy) is not fully specified yet and could be fleshed out further. In Sect. 3, we raise an objection to (Policy) and an objection to (Precision) and explain how (Accuracy) avoids those objections. In contrast to (Policy), (Accuracy) does not permit scientists to skew their advice based on their moral assessment of policies. This is an advantage because such influence

⁶ In a similar vein, (Boulicault and Schroeder 2021, p. 113) argue in the context of a discussion of trust between the public and scientists that

[f]or scientists’ claims on a matter of great public importance to be trustworthy to the public, it would seem reasonable for the public to insist on having input into the level of certainty the scientists employed.

⁷ We will elaborate on this point in Sect. 3.2.

⁸ While (Accuracy) denies that scientists need to make such value judgments to decide how to communicate their findings, it leaves open whether scientists need to make value judgments to decide, for example, which hypotheses to investigate. We will elaborate on this point in Sect. 3.1.

would be procedurally unjust. In contrast to (Precision), (Accuracy) allows scientists to assert simplified claims that they have a low confidence in if doing so would make their addressees' credences accurate. This is an advantage because, in many such cases, it seems implausible that scientists have to stick to saying what they have a high credence in even if that would be epistemically unhelpful for their addressees. In Sect. 4, we turn to a problem for (Accuracy). It seems to require scientists to tell lies if doing so happens to maximize the expected accuracy of the policy-makers' credences. In Sect. 5, we assess actual science advice from the perspective of (Accuracy) to give a better idea of what (Accuracy) would amount to in practice.

2 Precifications of (Accuracy)

As stated, (Accuracy) is underspecified in various ways. This section will outline the different choice points one encounters when one attempts to make (Accuracy) more precise. Many arguments in this paper are largely independent from how one chooses at those choice points. Nevertheless, discussing the choice points will convey a better understanding of (Accuracy), and it will make it easier to later flag arguments that depend on particular ways of making (Accuracy) more precise.⁹

2.1 The accuracy of a credence

The first choice point concerns the notion of the accuracy of a credence. Various ways of defining this notion have been proposed in the literature.¹⁰ For example, one could measure the accuracy of a credence $Cr(p)$ in some proposition p by the squared distance between $Cr(p)$ and 1 if p is true or 0 if p is false. Thus, for example, if p is true, then the agent's credence is maximally accurate if it is 1, fairly inaccurate if it is 0.5, and maximally inaccurate if it is 0. Many accuracy measures have the property that if you have a credence of, say, 0.7 in p , and you satisfy certain requirements of rationality, then you consider 0.7 to have the highest expected accuracy among all possible credences in p . On such a notion of accuracy, if you try to make someone else's credences accurate, you will try to make their credences close to your own credences. In the rest of this paper, we will assume that (Accuracy) is fleshed out using such a notion of accuracy and hence implies that scientists ought to say what makes their addressees' credences close to their own credences.

2.2 Aggregating over propositions

The remaining three choice points concern how (Accuracy) aggregates the effects of what the advising scientist says on the credences in different propositions of different people at different times. The most straightforward choice regarding how to aggregate

⁹ Moreover, other norms of science advice proposed in the literature are often underspecified in similar ways. Drawing attention to the ways in which (Accuracy) is underspecified might therefore also be helpful for thinking about these norms.

¹⁰ See Pettigrew (2016) for an overview of different measures of accuracy.

over propositions would be that accurate credences in all propositions are equally important for determining what the scientist ought to say. But the resulting version of (Accuracy) would be implausible. Suppose that the scientist is asked by a policy-maker to advise them on the toxicity of some pollutant. Suppose further that the scientist knows that the policy-maker has much more inaccurate views about the effects of fracking on biodiversity, which they could easily correct if they changed the topic of the conversation. The suggested understanding of (Accuracy) would imply that the scientist should try to inform on fracking because that would realize larger accuracy gains. But it seems more plausible that if a scientist is asked to advise on the health effects of a pollutant, they should say things that make the policy-makers' credences in propositions about the pollutant accurate. They should not disregard the topic of the conversation merely because there are unrelated topics that they could better inform about.

A version of (Accuracy) that captures these verdicts would weight accuracy gains according to how relevant propositions are in the given conversational context rather than giving all accuracy gains equal weight. This version entails that if a scientist is asked to advise about the toxicity of a certain pollutant, they should say what makes credences in propositions about the toxicity of that pollutant accurate because these propositions are relevant in the given conversational context. It also entails that if two propositions are equally relevant and the scientist has to choose between either saying something that makes one credence very accurate or another credence a bit accurate, they should choose the former.

It is important to be clear about the notion of relevance at play. We intend to use the word in the sense in which it is used in the linguistics literature.¹¹ Relevance, in that sense, is always relative to a particular conversational context. How relevant p is depends on the extent to which p relates to what the conversation is about. Therefore, the relevance of p in a conversational context depends on what has been said before in the conversation (did someone ask about p ?) and on mental states of the conversation's participants (what are my addressee's goals for this conversation and is p important to those goals?).¹²

For example, the accuracy of the proposition that the pollutant is toxic for children and that it is toxic for adults might receive roughly equal weight because they are equally relevant in a conversational context in which the speaker is asked whether the pollutant is toxic. In contrast, if regulation of the pollutant for adults would be impossible due to political reasons, then the proposition that the pollutant is toxic for adults would be less relevant in the context, assuming that the addressees' goal for the conversation is to acquire action-relevant information. Finally, suppose that the policy-makers communicated that they are interested primarily in the effects of the pollutant on children. Then, according to the view we describe, the scientist should

¹¹ For example, Grice's (1975) Maxim of Relation requires speakers to 'be relevant'. It is relevance in that linguistic sense which we have in mind (see also Wilson and Sperber 2012).

¹² What is relevant can come apart from the literal meaning of the questions that were uttered. For example, if you were bitten by a snake and asked a biologist 'is this snake *poisonous*?', what is relevant is whether the snake is *venomous*. That is the proposition you desire to know the truth value of, even though you failed to utter words that express that proposition.

prioritize the accuracy of the credence in the proposition that the pollutant is toxic for children over that of the proposition that the pollutant is toxic for adults.

A different choice one could make concerning aggregation across propositions is to let the weight of a proposition depend partly or wholly on how morally important it is for the policy-maker to have accurate credences in this proposition. For example, if effects on children are more morally important than effects on adults, then the advising scientist should prioritize making credences about the effects on children more accurate.

In the rest of the paper, we will focus on the purely relevance-weighted version of (Accuracy). We do so partly for dialectical reasons. We find it instructive to illuminate the tenability of a norm that focuses on epistemic benefit and does not make what advising scientists ought to say depend on difficult moral questions such as the proper trade-off between harm to children and harm to adults. But we also think the purely relevance-weighted version of (Accuracy) is independently defensible.

Either the addressee is more interested in the more morally important proposition or not. In the first case, the more morally important proposition is also more relevant, on our notion of relevance. This could be because the addressee explicitly communicated that they are more interested in the more morally relevant proposition. But even if they did not explicitly do so, the mere fact that they are more interested in it is enough to make it relevant. As an analogy, if someone asks you for directions to a gas station, the proposition that the gas station is closed is more relevant than the proposition that it is two blocks south from here. This is because the interlocutor is more interested in that proposition, even if she has not explicitly said so. While the notion of relevance that we appeal to is non-moral—to be relevant is to relate to what is at issue in the current conversational context—the moral views of the policy-maker might thus indirectly influence what is relevant in a conversational context, as they determine which propositions they desire to know more about.

In the second case, the addressee does not have a stronger interest in the more morally important proposition, for instance because of a mistaken moral judgment. Then, the relevance-weighted view says that the scientist should not focus on propositions that are actually more morally important to have accurate credences in, for instance, propositions about the effects on children. It seems defensible that the scientist informs policy-makers about what the policy-makers are interested in, even if that is not always what is morally most important. To take a concrete example: policy-makers have, on average, greatly underappreciated the moral importance of being well-informed about pandemics. But nevertheless, scientists should not have tried to make credences about pandemics more accurate in contexts in which they were asked to inform about, say, biodiversity. When scientists write op-eds, position papers, or books, they may and arguably should choose topics where more accurate credences would be morally important. But when they are called upon by policy-makers to inform about a specific issue, it is plausible to say that what they should say depends on accuracy effects on propositions about the specific issue, regardless of its moral importance.

As a further defense of aggregating over propositions purely based on relevance, note that it is consistent with saying that, in some situations, scientists should refuse to continue to act as science advisors. For example, if a policy-maker has despicable

moral views and only asks about the effects of a pollutant on white men, then, in the resulting conversational context, such propositions are more relevant—in the linguistic sense—than the effects on Hispanic women. Confronted with such a situation, refusing to continue to act as a science advisor and possibly expressing disagreement with the moral presuppositions underlying such racially biased questions strikes us as a morally defensible course of action. The relevance-weighted version of (Accuracy) is consistent with this verdict; it is a norm about what to do if one acts as a science advisor, not a norm about when to agree to act as science advisor.

In addition to the question whether empirical propositions should get extra weight based on their moral importance, there is also the question whether the accuracy of credences in propositions *with moral content* should receive any weight. For example, should scientists also worry about making accurate the credence in the proposition that legislation that regulates emissions of the pollutant is all-things-considered morally required? In this paper, we will explore the view that scientists should purely focus on making credences about empirical matters accurate. Again, this choice is partly motivated by a desire to explore a view that allows scientists to advise while focusing on non-moral matters: the accuracy of credences in empirical propositions. While we think that this view is independently defensible, we acknowledge that there is room for discussion about the extent to which scientists should try to inform not only on the empirical matters that underlie moral judgments about what policies should be adopted, but also directly on those moral questions.

2.3 Aggregating over people

If an advising scientist says something, they influence the credences of policy-makers to whom they directly communicate, such as the readers of their reports or the person at the other end of phone. But they also indirectly influence the credences of other policy-makers. If a policy-maker is told by a scientist that some pollutant is toxic, then the policy-maker might repeat the claim in the presence of other policy-makers, thereby influencing their credences. Scientists might also indirectly influence the credences of the general public, for instance when the science advice is recorded and made publicly available, as is the case with Congressional briefings.

Again, the most straightforward way of fleshing out (Accuracy) would be that the accuracy of all people's credences matters equally in determining what the scientist should say. But such a strictly impartial understanding of (Accuracy) is implausible. Suppose that a scientist advises a member of a party who, in contrast to two other members of the party, has an unreasonably low credence in some pollutant causing lung cancer. The scientist might know that to make that person's credences accurate, they should emphasize the evidence for the pollutant causing lung cancer. But they also know that the policy-maker would eventually use the scientist's claims in their arguments and speeches, likely making the credences of the two other members who are already overconfident even less accurate. In that case, a purely impartial construal of (Accuracy) would implausibly entail that the scientist should make the credences of the person who sought their advice inaccurate. Rather, it seems that in virtue of that policy-maker having asked the scientist to inform them, scientists should prioritize making that particular policy-maker informed.

This suggests understanding (Accuracy) so that the scientist ought to say what makes the (immediate) addressees' credences accurate. Alternatively, one might say that while immediate addressees matter much more than other policy-makers and the general public, these other groups receive some non-zero weight, too. It is consistent with both views that the accuracy effects on the credences of immediate addressees' are the determining factor in most cases of science advice. This is all we will assume about how (Accuracy) aggregates over persons in the rest of the paper.

2.4 Aggregating over time

In many cases, scientists do not need to choose between making credences accurate in the near future or in the far future. Suppose that making an assertion would give a policy-maker more accurate credences five minutes from now than making some other assertion. Then, it should usually be expected that the former assertion would also give the policy-maker more accurate—or at least equally accurate—credences a week from now. All sensible ways of aggregating accuracy over different future times would prefer such Pareto-superior outcomes.

There might be some cases in which scientists do face an intertemporal trade-off. They could either say something that is simple and easy to remember and would therefore make credences accurate for a long time. Or they could say something that is hard to remember but would make the addressees' credences very accurate for a short time. At what rate should accurate credences in the future be discounted relative to the present? While we will not discuss such questions of intertemporal aggregation in detail, let us briefly note that the conversational context might often contain a relatively clear mandate as to how one should aggregate over time. If a policy-maker asks for advice regarding a particular policy decision, scientists should plausibly focus on ensuring accurate credences at the time of the decision. In contrast, if scientists are asked to provide some general advice on issues that regularly arise in a committee their addressees are long-standing members of, they should plausibly give more weight to accurate credences at times further away in the future. But, in any case, all we need to assume for our arguments is that (Accuracy) is understood so that credences in the future matter somewhat and that the mode of aggregation does not require scientists to make moral judgments to decide what empirical claims to assert.

In summary, there are several dimensions along which (Accuracy) can be fleshed out in different ways. For the purposes of this paper, we chose to be opinionated on the question how (Accuracy) trades off accuracy effects on credences in different propositions. We will understand (Accuracy) as considering accuracy effects on credences in empirical propositions, weighted by how much they relate to what is at issue in the conversational context. But our discussion does not require making strong assumptions about exactly which notion of accuracy is used or exactly how one aggregates across accuracy effects on different people at different times. In fact, many of our arguments can be discussed without explicitly considering more than one addressee, proposition, or point in time. After having discussed how to deal with such complexities, we will now keep things simple where we can.

3 The case for (Accuracy)

3.1 Only epistemic values should guide science advice

Various writers argue that a scientist's moral judgments ought to at least partially guide their decisions about what empirical claims to make. For example, Katie Steele, reviving a well-known argument by Richard Rudner, considers a scientist who decides whether to report that they have medium confidence or high confidence in the claim that the state of the world S obtains.¹³ She writes that¹⁴

They would need to guess what policy decisions advice of high confidence or medium confidence would lead to and consider the impacts if, in either case, the state S obtains or does not obtain. [...] Scientists [in making a decision about what to report] commit to an evaluation of the desirability of outcomes.

This is how (Policy) would tell the scientist to decide whether to report medium or high confidence that the state S obtains: assess the moral value of the possible policy consequences of saying various things and then say what maximizes expected value. In contrast, (Accuracy) would tell the scientist to report medium or high confidence depending on which of the two will give the policy-maker more accurate credences in the proposition that the state S obtains. The scientist should not consult their moral judgments about the desirability of having this or that policy in place.

A problem with (Policy) becomes apparent once we appreciate that what we morally ought to do when we act within a political system is sometimes not what has the best policy consequences. This is because, in addition to bringing about good policy consequences, it also matters that we follow just procedures. For example, you might find yourself in a situation in which you could bring about better policy consequences by conducting electoral fraud. Nevertheless, except in extraordinary circumstances, you should not conduct electoral fraud because doing so would undermine procedural justice.¹⁵

The problem with (Policy) is that it tells scientists to act in a way which undermines procedural justice, under a broadly democratic understanding of this ideal. For exam-

¹³ See Steele (2012) and Rudner (1953). See also Churchman (1948), Shrader-Frechette (1994), and Franco (2017) for related views. Heather Douglas (2000, 2008, 2009) also defends a view on which moral judgments about policy consequences should influence a scientist's decision what to say. However, her view appears to be that only the scientist's moral assessment of the consequences of saying something false but not of the consequences of saying something true should guide their decisions (see Douglas 2008, p. 14, in particular). We suspect that none of these writers would endorse (Policy) in full generality. For example, they would probably hold that the set of potential statements among which scientists should choose based on their moral judgments is restricted on the basis of considerations other than policy consequences (see footnote 45). Pace (Policy), scientists should not just say whatever maximizes expected policy consequences, possibly disregarding truth entirely. However, such differences do not matter for the arguments in this section. The arguments apply to any view that holds that scientists should in standard cases of science advice decide what to say based on their moral judgments about policy consequences, whether that decision is made from a set of options that was already narrowed down or not.

¹⁴ Steele (2012, pp. 898–899).

¹⁵ There is a large literature on why following just—and, in particular, democratic—procedures matters in addition to arriving at just outcomes. For some recent work, see Christiano (2008), Estlund (2009), and Kolodny (2014).

ple, a policy-maker is more likely to pass a policy regulating emissions of a pollutant if a scientist says ‘the pollutant is toxic’ than if they do not say that. If an advising scientist is more likely to say ‘the pollutant is toxic’ if they have certain moral views—as (Policy) implies—then a policy-maker will be more likely to regulate emissions of the pollutant if the advising scientist has certain moral views. But it is procedurally unjust that the advising scientist’s moral views influence policy decisions in this way.¹⁶

Why should one think that this influence is procedurally unjust? After all, the mere fact that some non-elected person’s moral judgments exercise a large influence on a policy-maker’s decision might not be procedurally unjust. Suppose that a member of an NGO advocates for a particular policy position in the presence of the policy-maker. In that case, the moral views of a non-elected person—the member of an NGO—exert a strong influence on the decisions of a policy-maker. But this, one might say, is entirely unproblematic from the point of view of procedural justice. It is therefore unclear why the influence of scientists’ moral judgments should be problematic.

We agree that procedural justice might often be consistent with a non-elected person’s moral judgments having a large influence on policy decisions. But we contend that procedural justice requires that, in such situations, policy-makers can viably deny that person’s and any other person’s moral judgment such an influence. The policy-maker’s decision to give weight to a non-elected person’s moral judgments must not be due to the policy-maker being put in a situation in which denying influence to some non-elected person’s moral judgments would be very costly. Policy-makers, in other words, should have the viable option of not letting their policy decisions be influenced by any non-elected person’s moral judgments. Only if that condition is satisfied can we say that by freely choosing to let one’s policy decisions be influenced by some person’s moral judgment, the policy-maker legitimized this influence. If the policy-maker merely resigns to having their decision be influenced by some person’s moral judgment because there is no viable alternative that does not involve doing so, then the policy-makers ‘decision’ to do so does not make the influence unproblematic. In those cases, the influence is undemocratic and procedurally unjust.

Requiring that such causal influence must always be the result of a choice of the policy-maker in a situation in which they could have viably avoided any such influence would be an implausibly strong condition on procedural justice. After all, policy-makers will inevitably engage in moral debates with other people and this will often result in those people’s moral judgments influencing policy-makers’ decisions. But we only need to claim that the condition is necessary for procedural justice in certain, *prima facie* suspicious cases: cases in which the causal influence is not mediated through a change of the policy-maker’s moral judgments. It is one thing if a policy-maker’s decision is influenced by another person’s moral judgments because that person makes moral arguments which the policy-maker finds convincing and changes their moral views in response to. Indeed, it is a feature of a healthy democracy that policy-makers

¹⁶ The thought that there is something anti-democratic about allowing scientists to use moral judgments when they communicate has been articulated before (see, for instance, Pielke (2007), Betz (2013), and de Melo-Martín and Intemann (2016); on the tension between democratic ideals and governance by scientists, see Sartori (1987, pp. 434–439). But often, this objection is raised without carefully explaining what is anti-democratic about scientists using moral judgments when they decide what to say about empirical matters. We hope that our discussion will help to sharpen the argument.

engage with other people's moral views and are open to changing their own views based on good arguments. However, it is quite another thing if a non-elected person's moral judgments causally influence a policy-maker's decisions, but that influence is not mediated by changing the policy-maker's moral judgments. If the non-elected person judges policy P to be more morally desirable than policy Q , then the policy-maker is more likely to choose P than if the person judges policy Q to be more morally desirable—even though the policy-maker's moral judgments are the same in both cases. To sum up, we claim that procedural justice requires that the policy-maker must be able to be free from such influence.¹⁷

This condition is satisfied in the case of the NGO. On the most natural way of imagining this case, the policy-maker will simply change their moral views in response to arguments they hear and therefore decide differently. Even if they do not change their moral views, they had the viable option to refrain from changing their decisions based on the arguments given by the member of the NGO. They freely decided to give the NGO member's moral judgments more weight and thereby legitimized this influence. The condition also explains why other cases are problematic. Suppose that a blackmailer credibly threatens to harm the policy-maker if they do not make policy decisions according to the blackmailer's moral judgments. The policy-maker decides to comply with the blackmailer's demands. Clearly, this influence of the blackmailer's moral judgments is in tension with democratic ideals. The condition we propose generates this verdict. The condition is violated because it is not a viable option for the policy-maker to deny the blackmailer's moral judgments influence on policy decisions: the costs of doing so would be prohibitively high. The policy-maker's decision to give weight to the blackmailer's moral judgments does not legitimize the influence from the point of view of procedural justice.

If scientists followed (Policy), scientists' moral judgments would influence policy decisions without that influence being mediated by changes in the policy-maker's moral views. After all, the scientist's moral judgments would influence what the scientist says about *empirical* questions, such as whether some pollutant increases the risk of cancer. This will influence the policy-maker's credences in empirical propositions but not their moral views. The condition we proposed for such kinds of influence would be violated: policy-makers would often lack a viable alternative to letting scientists' moral judgments influence their decisions. Thus, the influence of the scientists' moral judgments would be procedurally unjust. Scientist would lack a viable alternative because, first, the only alternative would be to not let their decisions be influenced by what scientists says, and, second, that is not a viable alternative.

First, if policy-makers changed what they do based on the scientist's empirical claims, then they would give the scientist's moral judgments weight. For example, if they let their decisions be influenced by the statement 'the pollutant is toxic', they could not help but let their decision be influenced by the scientist's moral judgments, in the sense that a scientist with different moral judgments might have made different empirical statements, which would have led the policy-maker to do something else.

¹⁷ We propose this condition as a necessary condition on procedural justice. For the purpose of this section—objecting against (Policy) and some of its prominent variants—our condition appears sufficiently strong. But there may be plausible, stronger versions of this condition. These would allow objecting to an even larger family of norms than our necessary condition does.

Thus, to deny the scientist's moral judgments influence on their policy decisions, the policy-maker would have to avoid changing what they do based on the scientist's empirical claims. If (Policy) was adopted, policy-makers would be put in a situation in which they have to decide between ignoring scientists altogether or having their policy decisions be influenced by the scientist's moral judgments.

But, second, ignoring scientists is not a viable option. The policy-maker knows that they would impose severe risks on others if they made decisions about policies that crucially depend on empirical matters without getting informed by scientists. In such a situation, many policy-makers would rightly accept that scientists' moral judgments influence their decisions, but they might only accept it because the alternative is out of the question. Thus, the moral judgments of some non-elected person would have a greater influence on a policy-maker's decisions but only because the policy-maker cannot viably deny them that influence. That, we contend, is problematic from the point of view of a broadly democratic conception of procedural justice.

In contrast, if scientists followed (Accuracy), then their moral judgments about policy consequences would not get a special influence on policy decisions. Scientists would decide what to say based on empirical judgments about what makes the policy-maker's credences accurate. They would not use moral judgments about policy consequences to decide what to say. Thus, if the policy-maker changes their policy decisions based on the empirical claims a scientist makes, they do not also have to accept that their decisions will thereby be sensitive to the scientist's moral judgments. Hence, (Accuracy) does not undermine procedural justice in the way (Policy) does.

As a response to this objection, a proponent of (Policy) might say that an advising scientist should explain to the policy-maker how their moral judgments shape what empirical claims they make.¹⁸ Then, the policy-maker could correct for the influence of the scientist's moral judgments on the scientist's empirical claims. In effect, the policy-maker could infer which empirical judgments and moral judgments the scientist combined in order to arrive at the statements they made. The policy-maker would have the viable option to permit the scientist's empirical judgments but not their moral judgments to influence policy decisions.¹⁹

The view that advising scientists should let their moral judgments influence what empirical claims they make but communicate those moral judgments so that scientists can undo the influence is unappealing for at least two reasons. First, it relies on the doubtful claim that policy-makers can subtract the influence of the scientists' moral judgments from their assertions. It assumes that they could update their credences as though they had perceived assertions scientists would have made had they not used their moral judgments to decide which empirical claims to make. Moral judgments can lead one to bias one's empirical claims in a myriad of ways and policy-makers cannot be expected to accurately assess the ways in which a particular piece of science

¹⁸ See Douglas (2008).

¹⁹ See Elliott and McKaughan (2014). It is important to distinguish the suggestion we discuss here from the different view that scientists should inform about empirical questions without consulting their moral judgments about policy consequences and then, in a second step, illustrate how the empirical information could be combined with moral values to arrive at policy decisions. We have no objections against that view, as long as the part that informs about empirical questions is conducted according to (Accuracy) rather than some other norm.

advice was biased, especially because they are not subject matter experts on the topics they are advised on. Second, it seems hard to motivate the claim that scientists should let their moral views influence what empirical claims they make but do it in a way that allows policy-maker to undo this influence. Presumably, a central motivation for (Policy) is that scientists have an obligation to bring about good policy consequences. But from that point of view, communicating in a way that allows policy-makers to undo the communication decisions that were supposed to bring about better policy consequences is hard to justify.²⁰

Alternatively, proponents of (Policy) might say that while being transparent about one's values does not enable policy-makers to rely on the empirical advice without also having their decision be influenced by the scientist's moral views, policy-makers could at least choose to listen to scientists whose moral views coincide with their own moral views. After all, while policy-makers cannot viably ignore all scientists, they can viably ignore some scientists.

But, first, given the diversity of moral views that policy-makers in pluralistic societies endorse, it is unrealistic that for any representative and any policy-relevant empirical question, there will be a scientist with the required combination of moral views and empirical expertise.²¹ And even if there was such 'matching' science advice available, the condition of procedural justice we proposed above would still be violated. Policy-makers should be able to not give any non-elected person's moral judgment causal influence on their policy decision, at least the kind of influence that is not mediated by a change in the policy-maker's moral judgments. It is hard to see why there should be an exception for scientists with moral judgments that coincide with the policy-maker's. Importantly, the influence exerted by an advising scientist with the same moral judgments will not always push the policy-maker towards choosing the option that they would have chosen anyways if they were acquainted with all the evidence that the scientist has. After all, the policy-maker is most likely aware of some consideration that bear on the policy decision and which the scientist is unaware of. Such considerations might well lead the policy-maker to choose a different option than the scientist judges to be morally best. Thus, the scientist's influence does not seem unproblematic just because their moral judgments match those of the policy-maker.

As a different response, the proponent of (Policy) might defend an 'institutional' version of (Policy). According to this view, rather than using their own moral judgments to decide what to say, there should be norms which capture particular moral judgments about policy consequences. Much as there are norms about which strengths of evidence you need to convict someone of a crime or of a civil offense, there could be norms about which strength of evidence you need to declare a substance safe that would kill people if it was unsafe or to declare a substance safe that would cause mild headaches if it was unsafe. If scientists followed such norms when they decided what to say, they would not use their own moral judgments about the moral disvalue of people being killed or experiencing headaches. Rather, by following such norms, they would be guided by the moral judgments that underlie the norms. Furthermore, the norms

²⁰ Also see de Melo-Martín and Intemann (2016, p. 512) and Elliott and Richards (2017b, p. 272) for worries about this response.

²¹ See de Melo-Martín and Intemann (2016, p. 511) for a similar objection in the context of the idea that diverse advisory boards should scrutinize how values are used in scientific studies.

and the underlying moral judgments could be democratically approved. Maybe they would not be directly voted on by elected representatives, but they could plausibly be defined by appointed officials who are accountable to the legislature.²²

The institutional version of (Policy) faces a dilemma. Either the norms are sufficiently precise to single out a narrow set of possible statements for any question a scientist is asked to give advice on and any evidential situation they might find themselves in. In that case, precise value trade-offs must underlie these norms. While these are democratically approved, many elected representatives will necessarily disagree, just as many elected representatives disagree with other democratic decisions. It seems highly undesirable that officials are presented with the choice of either not to listen to scientist or to accept moral views that they reject. Representatives should be able to decide in an informed way based on their own moral judgments or those of the constituency they represent, even if those moral judgments are not shared by the majority. Moreover, one might worry that spelling out norms to a degree of precision that leaves scientists with little freedom about what to say in any epistemic situation is impractical. The other option is that the norms leave scientists much freedom about what to say. But in that case, the proponent of the institutional version of (Policy) would presumably say that scientists should then fall back on their own moral judgments to decide what to say. But if scientists' moral judgments largely drive what empirical claims they make and norms only give rough guidance, the original objection applies. The institutional version of (Policy), whether it assumes precise or imprecise norms, fails to offer an appealing response to the objection.

As a final attempt to argue that (Policy) is not in tension with procedural justice, one might note that policy-makers can and do seek advice from many different scientists. While each advising relationship will lead to an influence of a scientist's moral judgment on a policy-maker's decisions, different scientists will make different moral judgments. The influence of all these moral judgments will cancel out in a system in which many scientists advise many policy-makers.

But there is strong evidence that scientists' moral views are systematically different from the moral views in the general population. For example, 81% of scientists in a survey in the U.S. identify as Democrats or lean towards the Democratic Party, compared to 52% in the general population.²³ This suggests that if scientists' moral judgments exert influence on policy decisions through science advice, the net effect would point in a clear moral direction rather than being morally neutral.

Instead of trying to argue that (Policy) is consistent with the alleged demands of democratic ideals, one might suggest that (Accuracy) also requires moral judgments about policy consequences. A proponent of (Policy) might say that it is simply

²² This institution-centered way of thinking about values in science has been proposed by Wilholt (2013) and Steel (2016), for instance. Steel (2016) argues that there are already such norms in place, and that scientists have no choice but to stick to them because, otherwise, they face the danger of having their work ignored or not receiving funding. We agree that concerning some decisions that scientists make—such as how to run an experiment—scientists do not have the freedom to pick whatever alternative they judge to have the best consequences. But when it comes to evidential thresholds in science advice, the norms we actually have clearly leave scientists with some freedom to make their own decisions (see also Steele 2012, p. 899). Even if there were such precise norms, it is unclear how reliably they would be enforced given that some science advice happens behind closed doors.

²³ See Kohut et al. (2009).

inevitable that scientists explicitly make or at least implicitly commit to moral judgments when they decide what empirical claims to assert.²⁴ Scientists must decide how to map their complex empirical judgments to less nuanced statements for policy-makers, and there is no appropriate basis on which to make this decision other than moral judgments about the consequences of saying various things. Maybe this is unfortunate from the point of view of procedural justice, and we should therefore try to reduce the discretion that scientists have when they make these decisions. But as long as we do not want to abandon the practice of science advice altogether, we will just have to accept that scientist's moral judgments have an increased influence on policy decisions.

However, (Accuracy) shows that there is a different possible basis for deciding what to say, namely, a concern for the accuracy of the policy-maker's credences. Moreover, that scientists explicitly make or at least implicitly commit to a moral judgment when they decide to assert some empirical claim *C* is not in tension with (Accuracy). Proponents of (Accuracy) can agree that scientists who decide to assert some empirical claim *C* explicitly make or implicitly commit to a moral judgment, namely, that *C* is the morally right thing to say in this situation. Proponents of (Accuracy) disagree with proponents of (Policy) merely in what other commitments this entails. Proponents of (Policy) claim that the moral choice-worthiness of saying something in a situation of science advice is determined by the moral value of its potential policy consequences, and therefore advising scientists commit to moral assessments of policy consequences. Proponents of (Accuracy) claim that the moral choice-worthiness of saying something in a situation of science advice is determined by the epistemic benefit it bestows upon the policy-maker, and therefore advising scientists do not commit to moral assessments of policy consequences. Conversely, to figure out what they should say, scientists do not need to confront ethical questions about how good or bad it would be if this or that policy was implemented. Instead, they need to consider empirical questions about how saying various things will affect the credences of the policy-maker.²⁵

A different objection would contend that telling actual scientists to follow (Accuracy) would lead many of them to decide what to say at least partially based on their moral judgments about policy outcomes. After all, it will often be unclear exactly what (Accuracy) entails in particular situations. For example, it might be unclear

²⁴ See Steele (2012) for an argument of this flavor.

²⁵ Here is a different way of making the same point. On (Accuracy), a scientist's idealized reasoning about what to say in a situation of science advice proceeds as follows:

1. Asserting *C* makes my addressees' credences most accurate in expectation.
2. (Accuracy)
3. Therefore, I should assert *C*.

The scientist clearly engages in moral reasoning. After all, the conclusion of the reasoning is a moral claim. But the only moral premise that the scientist needs is (Accuracy). They do not need to cast further moral judgments to decide what they ought to say. Rather, they need to cast empirical judgments about which assertions would be most accuracy-conducive.

In contrast, a scientist's idealized reasoning if (Policy) was followed would look as follows:

1. Asserting *C* has the morally most desirable policy consequences.
2. (Policy)
3. Therefore, I should assert *C*.

exactly how relevant different propositions are in a conversational context. In such situations, scientists could say many things while pretending—to others and possibly also to themselves—to be following (Accuracy). In other words, (Accuracy) leaves plenty of room for moral judgments, self-interest, and other factors to consciously or unconsciously bias scientists' communication decisions.

This is a fair worry. But it seems plausible that although actual scientists would not be able to completely shield their communication decisions from their moral views about policy outcomes, they would still do so much more if they followed (Accuracy) than if they followed (Policy). After all, (Accuracy) encourages them to bracket their moral judgments about policy outcomes whereas (Policy) does not. As an aside, it should also be noted that structurally similar worries apply to (Policy) and (Precision). These norms also leave plenty of room for conscious and unconscious bias to influence what scientists say. It will often be unclear what assertions have the best policy consequences, morally speaking. Thus, (Policy) leaves room for advising scientists being consciously or unconsciously guided by self-interest. And (Precision) would need to include some notion of relevance to specify which propositions should be asserted among the many that a scientist has high confidence in. As it will often be unclear exactly which propositions are relevant, scientists might be biased by their moral views or self-interest in what they assert even if they attempted to just say relevant things that they are confident of. Thus, the imperfect nature of actual scientists would, to some extent, undermine the claimed advantages of (Policy) and (Precision) as well.

A different objection of this flavor remarks that even if we assumed perfect compliance with (Accuracy), scientists' moral judgments might still influence policy decisions via science advice. Scientists must make decisions not just about what to say when they advise policy-makers, but also, for example, when they choose their research methodologies.²⁶ The norm (Accuracy) does not say anything about how scientists should make those other decisions. Thus, it is consistent with (Accuracy) that all of them are made based on the scientist's moral judgments. But these decisions influence what the scientist will tell policy-makers in situations of science advice. For example, if a scientist decides not to conduct a certain experiment, then this will influence what they can and will tell the policy-maker. The norm (Accuracy) allows that the decision whether to conduct the experiment was based on the scientist's moral judgments about policy consequences. Thus, (Accuracy) on its own does not rule out that scientists' moral judgments influence the policy-makers' decisions through science advice.

But, again, the reason to prefer (Accuracy) over (Policy) still applies. Even if scientists made all other decisions based on their moral values, scientists' moral judgments would have considerably less influence on policy decisions if (Accuracy) was fol-

Footnote 26 continued

Thus, even though a proponent of (Accuracy) need not deny that reasoning about what one ought to say involves moral claims, they can still point to an important difference between (Accuracy) and (Policy): (Policy) but not (Accuracy) requires scientists to supplement the norm of science advice with further moral judgments. Of course, we concede that scientists should be guided by further moral judgments about policy consequences in pathological situations—which we explicitly set aside in this paper—in which moral atrocities would be committed if scientists made a particular statement. Even the most staunch defender of value-free science should agree with this.

²⁶ See Douglas (2000, p. 565) and the contributions in Elliott and Richards (2017a, part 3).

lowed than if (Policy) was followed. The norm (Accuracy) would at least remove the influence of scientists' moral judgments from the decision how to communicate the results of the research they decided to conduct. If we are right that this influence is problematic from the perspective of procedural justice, then the fact that (Accuracy) diminishes it relative to (Policy) would still count as a reason in its favor.

3.2 Communication style should be context-sensitive

It is often claimed that scientists should make their uncertainty explicit when they communicate their findings. For example, Gregor Betz proposes that “[a]llegedly [...] value-laden decisions can be systematically avoided [...] by making uncertainties explicit and articulating findings carefully”.²⁷ The imperative to make uncertainty explicit would be supported by a norm such as (Precision).

The view that scientists should always make uncertainty explicit is in tension with research on how people change their credences when confronted with language describing uncertainty. Qualitative specifications of uncertainty, such as ‘it is unlikely that’, are interpreted as corresponding to different probabilities depending on the utility of the event whose likelihood is at issue.²⁸ This makes it hard to anticipate how the addressee will interpret what one says. Quantitative specifications are less ambiguous but might have other undesirable effects. First, they might simply be ignored by the policy-maker. For example, in a recent study, Eva Vivald and Aidan Coville find that policy-makers tend to ignore specifications of the variance of estimates.²⁹ Second, and more problematically, quantitative specifications might cause the policy-maker to misunderstand what the scientist says.³⁰ In such cases, it seems implausible that a scientist should make their uncertainty explicit, given that it leads to the policy-maker misunderstanding or ignoring what they say.³¹

²⁷ Betz (2013, p. 209). Betz’s view builds on Jeffrey’s (1956) well-known response to Rudner (1953).

²⁸ See Weber and Hilton (1990).

²⁹ See Vivald and Coville (ms).

³⁰ For example, Johnson and Slovic (1995) observed that “people seem to see lower risk estimates (10^{-6} , as opposed to 10^{-3}) as less credible”. Fischhoff (1995) suggests that if experts tried to just give their addressees ‘the numbers’, then

[c]onfused recipients of such raw materials may add some of their own uncertainty to that expressed by the analysts. Suspicious recipients may adjust risk estimates upward or downward to accommodate (what they see as) likely biases.

³¹ Frank (2017) and John (2018) also point out that making uncertainty explicit can be confusing and attack views such as (Precision) on these grounds. Relatedly, Elliott (2010) argues that advising scientists have a duty to promote their addressee’s ability to make informed decisions, and to fulfill this duty, they must make sure that their addressees understand what they say.

A document by the OECD on science advice emphasizes that “[c]ommunicating scientific advice in ways that maximise shared understanding (and minimise misunderstandings) is of key importance for scientific advisory bodies [...] [a] common mistake is for advice to be written in long, very technical reports [...] There is a balance to be achieved between oversimplification and the incomprehensibility of science language” (OECD 2015, pp. 21–22).

We contend that such cases favor (Accuracy) over (Precision). The norm (Precision) has the implausible implication that scientists should always make uncertainty explicit, independently of its effect on the listener. In contrast, (Accuracy) entails that if descriptions of uncertainty would cause confusion, scientists should not make their uncertainty explicit. Uncertainty should be made explicit only insofar as this is accuracy-conducive. Thus, the degree to which uncertainty ought to be made explicit depends on how receptive the addressee is to language describing uncertainty. This strikes us as an appealing implication of (Accuracy).

In addition to sometimes confusing the addressee, only saying things one has high confidence in—and thus always making uncertainty explicit—might also rule out other accuracy-conducive communication strategies. For example, studies suggest that strategically framing the facts one tries to communicate can positively influence one's addressee's ability to recall them later.³² A policy-maker might not remember much if the scientist carefully lays out the evidence about the effects of climate change on ecosystems. But if the scientist told a story about why the policy-maker's hayfever lasts longer each year, the policy-maker might be able to recall much more of the scientific insights that the scientist intended to communicate. Framing scientific information in compelling ways might in some cases be inconsistent with (Precision), if a compelling narrative requires making claims one has less confidence in—such as 'climate change makes your hayfever last longer'—rather than their high-confidence alternatives—such as 'there is strong evidence that...'. In such cases, (Precision) would not recommend employing such framing devices, even if they increased recall. In contrast, (Accuracy) would recommend doing so since they are conducive to making the policy-maker's credences accurate. Thus, for those who think that scientists should strategically use framing to communicate more effectively even if that requires making claims they are somewhat uncertain about, this speaks in favor of (Accuracy) over (Precision).³³ Analogous points could be made for other ways in which only saying things one has a very high confidence in stands in the way of inducing the most accurate credences.³⁴

While it strikes us as an appealing feature of (Accuracy) that it recommends scientists to tailor their communication style to suit their addressees, we also think that it would be desirable to have an explanation why, at least in some cases, scientists may permissibly simplify, omit mentioning uncertainties, and thus make assertions they think are not very likely to be true. After all, it generally seems morally dubious to assert things one does not have a high confidence in.

³² See Valkenburg et al. (1999) and Jaspert et al. (2011). We thank an anonymous referee for prompting us to discuss framing.

³³ See Nisbet (2009), Bolsen et al. (2014), and Yang and Hobbs (2020) for further discussion of framing in science communication.

³⁴ For example, Vivalt and Coville (ms) find that policy-makers tend to change their credences more in response to good news—such as that a program they want to implement is more effective than they thought—than in response to bad news. Hence, it might sometimes be more accuracy-conducive for scientists to emphasize uncertainties if the result is good news from the perspective of the policy-maker. After all, this might prevent the policy-maker from changing their credences too much. In contrast, if the result is bad news, it might be accuracy-conducive not to mention uncertainty. In general, any of the cognitive biases that policy-makers have is associated with a communication strategy that aim to mitigate its effect and thus is more accuracy-conducive (see also Akin and Landrum 2017).

One possible explanation is that there is a mutual understanding between the scientist and the policy-maker that the scientist will resort to asserting simplified claims in which they do not have a high confidence when that is accuracy-conducive.³⁵ Hence, a policy-maker cannot reasonably expect scientists to only say things in which they have a high confidence when simplifying or glossing over uncertainties is required for making credences in relevant propositions accurate. Rather, the policy-maker should expect that the scientist will communicate in this way. Thus, the scientist does not deceive the policy-maker. They do not represent themselves as having high confidence in the simplified and unqualified claims they make. This explains why it is morally unproblematic for scientists to say things they are not very confident in, even though in other contexts—such as on academic conferences—it would be deceptive and morally wrong to do so.

An argument in favor of this explanation is that the same phenomenon of partially relaxed norms of communication occurs in contexts other than science advice. Suppose you attend a public lecture about black holes on the open day of an astrophysics department. For the sake of conveying a few basic facts about black holes to the audience, the lecturing astrophysicist will simplify claims and gloss over uncertainties. They will assert claims in which they have low confidence and even claims they believe to be flat-out false, as asserted, if asserting more complicated, qualified claims would confuse the audience and distract from more basic facts they intend to convey. In this situation, you could not reasonably expect that the scientist will refrain from simplifications and only say things they are very confident in, even if that would confuse or in other ways hinder the communication of the central facts. You could not reasonably expect this because you know that the point of a public lecture on black holes is to help laypeople understand some basics about black holes—and maybe get some teenagers excited about studying astrophysics—and this requires partially relaxing norms which normally require people to only assert things they are confident in. According to the explanation we suggest, the same happens in situations of science advice. The point of science advice is to give policy-makers more accurate credences. Thus, norms that would otherwise require scientists not to assert claims they have low confidence in are partially relaxed, and scientists are therefore morally permitted to assert such claims.

4 The case against (Accuracy)

In the previous subsection, we claimed that in *most* cases, it seems right that scientists ought to simplify and gloss over uncertainties if that is accuracy-conducive. But suppose a scientist advises a policy-maker who has much too little confidence in a pollutant being toxic and who seems to be resistant to changing their credences. Suppose further that the vast majority of the evidence suggests that the pollutant is toxic, although there exists some weak counterevidence. After the scientist laid out the evidence for the toxicity of the pollutant, the policy-maker might ask in a skeptical tone of voice whether there is any counterevidence. Then, it might be most conducive to the accu-

³⁵ This explanation is inspired by Shiffrin (2016, ch. 1).

racy of the policy-maker's credence in the pollutant being toxic to falsely claim that there is no counterevidence. Granted, this would make the policy-maker's credence in the proposition that the scientist has counterevidence inaccurate. As this proposition is highly relevant in this context, the accuracy loss would make the utterance seem less appealing from the perspective of (Accuracy). But if the accuracy gain in the also highly relevant proposition that the pollutant is toxic is large enough, (Accuracy) might entail that the scientist ought to deny that they are aware of counterevidence.

Such an utterance would constitute a lie. In particular, by explicitly asking for any counterevidence, the policy-maker has brought about a context in which they have a reasonable expectation that the scientist will make any uncertainty explicit. The norms are no longer relaxed in a way that would permit not making uncertainty explicit. Falsely denying that one has counterevidence, even if this is the only way to get the policy-maker to have accurate credences, seems morally wrong.

To press this objection further, one might argue that (Accuracy) would end up recommending lying quite often. Policy-makers' credences are exposed to a variety of forces, and some of those forces do not push credences towards higher accuracy. In particular, interest groups might try to make policy-makers adopt certain credences not because they are rational given the evidence but rather because they speak in favor of policies that interest groups want to see enacted. For example, tobacco companies tried to make policy-makers believe that environmental tobacco smoke is not harmful, even though a balanced review of the scientific literature would have suggested otherwise.³⁶ Scientists must counteract such opposing forces when they attempt to make policy-makers' credences accurate.³⁷ Counteracting opposing forces requires exerting a stronger influence on the policy-makers' credences. One might suspect that, at least in some contexts, exerting such a stronger influence will often require telling lies.³⁸

There are different ways to derive an objection from the observation that (Accuracy) will sometimes recommend lying. First, one might object that (Accuracy) undermines the very concern that motivated it in the first place: a concern for the accuracy of policy-makers' credences. By telling scientists to greedily optimize for accuracy in

³⁶ Among other things, the tobacco industry set up the *Center for Indoor Air Research* which funded not only peer-reviewed research, but also research selected by tobacco industry executives. This research was more likely to confirm the tobacco industry's positions than peer-reviewed research projects (see Barnes and Bero 1996).

³⁷ Of course, a scientist will only know in some cases that the net influence of other people's influence on the addressee's credences is such that it would pull them, say, below the optimally accurate credence. In many other cases, as far as the scientist knows, a policy-maker might encounter both epistemic influences that lead them to increase their credence above what the scientist takes to be the most accurate credence and influences that lead them to decrease their credence below what the scientist takes to be the most accurate credence. If these two possibilities are about equally likely, then the influence is zero in expectation. In such cases, the scientist can maximize expected accuracy as if there were no other influences on the addressee's credences.

³⁸ In fact, one might worry that if scientists followed (Accuracy), then an unfortunate dynamic would unfold between scientists and their epistemic antagonists. Once scientists start lying, epistemic antagonists might respond by lying even more. In response, (Accuracy) requires scientists to engage in more lying to compensate for the increased misleading influence. It seems that if scientists followed (Accuracy), then the only equilibrium would be one in which both scientists and their epistemic antagonists vastly overstate their case in an attempt to compensate for unwanted influence from the other side. Thus, in the long run, (Accuracy) might end up recommending lying quite often. If scientists followed (Precision), this problem would not occur.

each particular case of science advice, it ends up undermining scientists' ability to influence policy-makers' credences. This is because, as explained above, optimizing accuracy in a particular case sometimes requires telling lies. But some of those lies will be exposed, and policy-makers might then lose trust in scientists.³⁹ Loss of trust in science makes scientists unable to influence policy-makers' credences, leaving them inaccurate. Hence, (Accuracy) undermines the concern for the accuracy of policy-makers' credences.

This objection is misguided. After all, as we explained in Sect. 2, (Accuracy) should be understood so that it gives some weight to the accuracy of credences in the future. Thus, (Accuracy) almost always prescribes not to tell lies because to maximize the accuracy of policy-makers' credences in the long run, scientists should avoid undermining trust in science.⁴⁰ The recommendations of (Accuracy) therefore often coincide with the recommendations of (Precision). But (Accuracy) would still be importantly different from (Precision). It would entail that (Precision) fails to latch onto the normative consideration that underlies the imperative not to lie in many circumstances, which is a concern about the long-term accuracy of addressees' credences. Moreover, (Accuracy) would still extensionally differ from (Precision) in at least two kinds of cases. First, simplification and omission of uncertainty in the cases discussed in the previous section would still be endorsed by this modified version of (Accuracy) but not by (Precision). Long-term accuracy is less likely to be jeopardized in such cases because, as explained above, there is a mutual understanding that scientists will simplify what they say, so no loss of trust has to be feared from it 'coming to light' that a scientist simplified their statements and omitted mentioning uncertainty.⁴¹ Second, even though (Accuracy) will rarely prescribe telling lies, it sometimes will. If a scientist can be fairly sure that the lie will never be exposed, and they think that it will make the policy-maker's credences more accurate, then (Accuracy) entails that they morally ought to tell the lie.

As a second objection, one might say that lying is always at least somewhat morally bad, and that this moral badness can surely sometimes outweigh the moral goodness of making the policy-maker's credences more accurate, in particular if the accuracy gains obtained by lying are small.⁴² In those cases, the policy-maker should not lie, but (Accuracy) would entail that they should. Thus, (Accuracy) is false.

One response to this objection is to reject its premise: that lying is morally wrong in these cases. This would require claiming that the imperative to make policy-makers' credences accurate always outweighs the moral wrong of lying. If falsely denying that one has any counterevidence makes the policy-maker's credences more accurate and will not undermine trust in science in the long run—e.g., because it is certain never

³⁹ See Nisbet (2009, p. 59).

⁴⁰ If one interprets (Accuracy) so that it gives little weight to credences in the future and to credences held by other persons' than the immediate addressees, and it is unlikely that others will quickly expose the lie, then (Accuracy) will tend to recommend lying more often.

⁴¹ A more complete 'extensional collapse' of (Accuracy) into (Precision) might happen in a society in which policy-makers are extremely scientifically literate. In such a society, scientists might communicate in the most accuracy-conducive manner if they carefully describe the scientific evidence, specify uncertainty, and so on.

⁴² See Kupfer (1982) for some discussion of the inherent wrongness of lying.

to be revealed that one actually had counterevidence—one should go ahead and tell a lie. Science advice, on this picture, is not for everyone. If you accept the invitation to advise policy-makers, you might be brought into situations in which you morally ought to lie for the sake of making the policy-maker's credences about the relevant questions accurate.

To make the implication that scientists should sometimes lie more palatable, note that the lies that the norm would require scientists to tell seem importantly different from paradigmatic cases of lying. While the lies that the norm prescribes are intended to give the hearer inaccurate credences in some propositions, these inaccurate credences are only caused in order to give the listener accurate credences in other propositions. Put differently, the inaccurate credences are not the point of the lie; they are mere epistemic collateral damage.⁴³

This response tries to soften the blow of the objection without modifying (Accuracy). Alternatively, one could concede that telling a lie is not what one should morally do in those cases and modify (Accuracy) to accommodate this verdict. One could add a hard side-constraint to (Accuracy) which excludes questionable means to maximize accuracy, such as lies. For example, Heather Douglas imposes such a restriction on her norm to avoid the conclusion that scientists ought to deceive their addressees in cases in which the unmodified norm would tell them to do so.⁴⁴ The same restriction could be added to (Accuracy).

As another alternative, one could endorse a balancing view. When we presented the three norms, we noted that all of them have counterexamples: when following them would foreseeably lead to the scientist being killed, the scientist should—at least in some such cases—not follow them. We could take this thought a step further and say that even in standard cases, there are multiple relevant considerations that bear on what the scientist ought to say. Three potentially relevant considerations are that one should say what one is certain in, bring about good policy consequences, and make the policy-makers' credences accurate.⁴⁵ What the scientist should say is determined by a balance of these considerations. On this 'balancing view', our paper argues that the first consideration often does not apply—making statements that the scientists are not confident in is often unproblematic—and that the second consideration is usually suspended for the sake of procedural justice. This is compatible with saying that in cases in which those considerations would be severely set back for small gains of accuracy, they outweigh the accuracy-related consideration, and the scientist should not maximize accuracy.

In sum, while the implication that scientists sometimes ought to lie is a reasonable objection against (Accuracy), there are ways to either soften the blow of this objection or to accommodate it while retaining the core idea underlying the accuracy-focused view.

⁴³ See also John (2018, p. 83).

⁴⁴ See Douglas (2009, p. 81). Relatedly, see Havstad and Brown (2017, footnote 14). As Douglas notes, imposing a hard side-constraint that rules out deception still leaves plenty of options for scientists to choose from. For example, scientists can "choose whether or not to emphasize the importance of [...] uncertainties".

⁴⁵ Other considerations might include, for instance, whether the scientist has promised not to deceive the policy-maker. For example, in most U.S. Congressional committees, witnesses can be sworn in.

5 (Accuracy) in practice

5.1 Assessing actual science advice

To get a better feel for what (Accuracy) would entail in practice, we will now assess the gap between actual and ideal science advice according to (Accuracy). This is meant to illustrate the implications of (Accuracy); it is not meant as an argument for or against it. After all, whether reality conforms to a suggested norm is, on its own, irrelevant for the plausibility of the norm.

Let us first look at written reports by scientists for policy-makers. When one surveys such reports, one quickly finds that many of them carefully indicate uncertainty. The norm (Accuracy) is more likely to recommend careful descriptions of uncertainty in contexts of written rather than verbal science advice, assuming that descriptions of uncertainty are more likely to be understood when presented in written form. That said, lengthy reports in which many claims are specified as ‘low confidence’ or ‘moderate confidence’ might not make policy-makers’ credences accurate. One might suspect that most policy-makers who are potentially interested in being advised by scientists on the topic will not take the time to work through such reports, yet alone pay close attention to specifications of uncertainty. They might simply ignore claims that are tagged with ‘low confidence’ or ‘moderate confidence’. If this empirical assumption is correct, (Accuracy) would recommend making reports more accuracy-conducive by not mentioning low-confidence claims and flat-out asserting high-confidence claims.

While it seems that (Accuracy) is critical of current practices of carefully reporting uncertainty, two qualifications must be made. First, reports which carefully make uncertainty explicit often begin with a one-page summary of the key findings. In those summaries, one encounters language that sacrifices careful descriptions of uncertainty for greater effect on the reader’s beliefs. For example, a summary in one report for policy-makers includes the claim that “[r]educing CO₂ emissions is the only way to minimise long-term, large-scale risks”.⁴⁶ After the summary, the full report follows, in which uncertainty is made explicit. There, one finds sentences such as “[i]f CO₂ emissions continue on the current trajectory, coral reef erosion is likely to outpace reef building sometime this century [High Confidence]”.⁴⁷ This communicative strategy resolves the accuracy-based worries against making uncertainty explicit. If the report makes uncertainty explicit only in the full report, but provides legislators with more digestible, less perspicuously qualified claims that are likely to effectively change their credences, then the report is perfectly consistent with (Accuracy).

Second, in cases in which the stakes are very high, policy-makers are more likely to put in the time and effort necessary to understand the nuanced epistemic situation of current science. One example of such a case is climate change. Consider the *Summary for policy-makers* of the *Fifth Assessment Report of the Intergovernmental Panel on Climate Change* (hereafter referred to as the IPCC Report). The report presents uncertainty by giving explicit definitions for terms such as ‘very likely’ in terms of numerical probabilities and then using these terms in its claims:

⁴⁶ IGBP, SCOR, & IOC (2013, p. 1).

⁴⁷ IGBP, SCOR, & IOC (2013, p. 16).

Anthropogenic influences have *very likely* contributed to Arctic sea ice loss since 1979.⁴⁸

The norm (Accuracy) vindicates that uncertainty is more carefully made explicit in the case of the IPCC report than in many other cases. The exceptional attention that was devoted to the report supports the assumption that addressees would muster the time and effort to take qualifiers such as 'very likely' into account as they updated their credences. Even if a majority of policy-makers ignored the qualifiers and only a small minority considered them, their inclusion would still have been beneficial in expectation. Thus, an argument for making uncertainty explicit based on (Accuracy) seems much more promising in the case of the IPCC report than in many cases of science advice, even though it is hard to know how much more accurate such qualifiers made policy-makers' credences.⁴⁹

In sum, (Accuracy), together with the empirical claim that explicit statements of uncertainty can lead to science advice being misunderstood or ignored, cautions against carefully stating uncertainty in all contexts. However, reports which include simplified summaries or are expected to receive a lot of attention by policy-makers can plausibly do so without violating (Accuracy).

It is worth mentioning that, beyond the decision about whether to indicate uncertainty or not, (Accuracy) can also guide detailed decisions about exactly how to communicate uncertainty. For example, a natural worry about the IPCC's current communication strategy is that some readers may—possibly unintentionally—revert to using their natural language understanding of terms such as 'very likely'. Ordinary meanings of these terms might come apart from the IPCC's explicit definitions of them. As a result, policy-makers might have inaccurate credences after reading the report.⁵⁰ Thus, a more accuracy-conducive strategy might have been to reiterate the corresponding probability range each time a term such as 'very likely' is used in the main text: 'anthropogenic influences have contributed to Arctic sea ice loss since 1979. [very likely, 90%-100%]'.⁵¹ If this is so, and (Accuracy) is correct, then the IPCC ought to present their uncertainty in this way. This illustrates that (Accuracy) provides a clear, applicable criterion that can guide even fine-grained decisions about how to present scientific findings.

Scientists also advise policy-makers in verbal conversations. In those contexts, (Accuracy) is likely to more often recommend not making uncertainty explicit. After all, it seems likely that descriptions of uncertainty lead to confusion more often in spoken rather than written language.

It will not surprise the reader that there are many examples of spoken science advice which are indeed less careful than written science advice. For example, in congressional hearings, one can find statements such as:

⁴⁸ IPCC (2014, p. 19 their emphasis).

⁴⁹ Let us stress that the question of whether making uncertainty explicit was accuracy-conducive in this case is an empirical question. A priori, doing so could also have been harmful, for instance by making the report harder to read.

⁵⁰ Patt and Schrag (2003) note this problem.

⁵¹ The findings by Budesu et al. (2009) suggest that this indeed makes readers more likely to correctly understand the authors' intended probabilities.

Many strands of mutually supporting evidence are woven into the confident knowledge that loss of land ice and warming of the ocean are driving sea-level rise.⁵²

While the speaker was involved in the compilation of the IPCC report, and thus certainly knew how to articulate uncertainty more precisely, it is consistent with (Accuracy) to here use the less careful phrase of ‘confident knowledge that p ’. Such a phrase is arguably more likely to make the policy-maker adopt a high credence in p than careful, probabilistic statements, such as ‘we are 95% confident that p ’. Hence, if the speaker also believed that the policy-maker’s credence should be high in order to be accurate, (Accuracy) would endorse the choice of a less careful phrase.⁵³

5.2 Deciding What to Say

The norm (Accuracy) specifies what scientists ought to say. It states that a relation holds between certain facts about the effects of utterances on credences and facts about which utterances advising scientists should make. It is a separate question how scientists should deliberate about what to say so that their advice will end up conforming with (Accuracy). The norm (Accuracy) does not answer that question. It is a criterion of rightness, not a decision procedure.⁵⁴

In particular, (Accuracy) does not entail that scientists should make explicit predictions about how accurate various credences of their addressees would be conditional on various statements they might make. In fact, making explicit predictions might sometimes not be a decision procedure that tends to produce science advice that conforms with (Accuracy). Just as trying to do explicit welfare calculations might be a bad way to maximize welfare, so trying to do explicit accuracy calculations might be a bad way to maximize accuracy.

What would be a better way for scientists to increase the chances that they say accuracy-conducive things? Scientists would need to get a good sense of how various things they could say would influence their addressees’ credences. To do so, gathering information about one’s addressees would probably be helpful. For instance, one might gauge the scientific expertise of one’s addressees to better anticipate how well they might understand various claims. Also, getting some acquaintance with widespread biases and misunderstandings of statistics might be helpful for communicating in a more accuracy-conducive manner. We take it that these are independently plausible prescriptions. The fact that they are also ways to make one’s advice more likely to conform with (Accuracy)—but not necessarily with other norms such as (Precision)—should be welcome news to those attracted to (Accuracy).

⁵² This quote is taken from Richard Alley’s statement in the congressional hearing titled *Earth’s Thermometers: Glacial and Ice Sheet Melt in a Changing Climate* on July 11, 2019.

⁵³ To be clear, we do not claim that the reason why actual scientists follow the norm is that scientists are concerned with the accuracy of policy-makers’ credences. It might simply be more difficult to carefully communicate uncertainty when one speaks than when one writes. All we claim is that simplification is, for whatever reason, more common in verbal science advice, as required by (Accuracy).

⁵⁴ See also Railton (1984).

Much more could be said about what good strategies for finding accuracy-conducive utterances are. But we lack expertise in the relevant fields to come up with much insightful on this question. The point that we want to stress here is merely that what scientists ought to say and what strategies they should employ to make it more likely that they say what they ought to say are related but distinct questions.⁵⁵

6 Conclusions

This paper developed the view that advising scientists should maximize the expected accuracy of policy-makers' credences. The view is appealing because it does not undermine the value of procedural justice in scientifically informed policy-making. Moreover, it yields the plausible verdict that in many cases, there is nothing wrong with simplifying and glossing over uncertainties if that is required to let policy-makers benefit epistemically from the conversation. While it overcomes some crucial issues of other proposed norms, it faces the objection that it recommends lying in situations in which lying seems hard to justify. We offered some responses to this objection and possible modifications of the norm which avoid it. In our discussion of how the norm would apply to real-world cases of science advice, we emphasized that it is able to guide fine-grained decisions about how to communicate uncertainty.

The view we explored is driven by the general idea that scientists should communicate in a way that maximizes the epistemic benefit of the addressee. In order to keep our discussion focused, we used a very specific notion of epistemic benefit: making credences accurate. While some have argued for accuracy as the sole fundamental epistemic value by which to evaluate how good or bad a given set of credences is, others might find the focus on accuracy or on credences overly narrow.⁵⁶ As we indicated in the introduction, we expect that much of our discussion would carry over to other norms which propose that science advice should be guided by maximizing some epistemic benefit to policy-makers. Therefore, we hope that even those who reject the focus on accurate credences will find our arguments worth considering when developing their own views.

Acknowledgements The second author is thankful for the financial support of the Studienstiftung des deutschen Volkes and the Claussen-Simon-Stiftung. The authors also received helpful comments by, among others, Jacob Nebel, Jonathan Quong, Jeffrey Russell, Ralph Wedgwood, participants of a poster session at the 27th Biennial Meeting of the Philosophy of Science Association, as well as participants of colloquia organized by Ulrich Gähde, Thomas Krödel, and Stephan Schmid. The views stated here are not necessarily the views of the supporting organizations and people mentioned in this acknowledgement. We are also thankful for the covering of the open access fees by the DEAL agreement.

⁵⁵ That said, we sometimes discussed decision procedures in this paper, too. For instance, we discussed the worry that scientist would unconsciously bias what they say based on their moral judgments if they tried to follow (Accuracy). That is a worry against scientists adopting a certain decision procedure, namely, thinking explicitly about what makes credences accurate. Strictly speaking, it is not a worry about (Accuracy), which is only a claim about what scientists ought to say.

⁵⁶ See Pettigrew (2016, e.g. pp. 5–7) for the view that all epistemic virtues that credences have reduce to accuracy.

Funding Open Access funding enabled and organized by Projekt DEAL. Studienstiftung des deutschen Volkes and Claussen-Simon-Stiftung.

Code availability (software application or custom code) Not applicable.

Declarations

Conflicts of interest/Competing interests (include appropriate disclosures) None.

Ethics approval (include appropriate approvals or waivers) Not applicable.

Consent to participate (include appropriate statements) Not applicable.

Consent for publication (include appropriate statements) Not applicable.

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/>.

References

- Akin, H., & Landrum, A. (2017). A recap: Heuristics, biases, values, and other challenges to communicating science. In K. H. Jamieson, D. Kahan, & D. A. Scheufele (Eds.), *The Oxford handbook of the science of science communication* (pp. 455–460). Oxford University Press.
- Barnes, D., & Bero, L. (1996). Industry-funded research and conflict of interest: an analysis of research sponsored by the tobacco industry through the center for indoor air research. *Journal of Health Politics, Policy and Law*, 21(3), 515–542.
- Betz, G. (2013). In defence of the value free ideal. *European Journal for Philosophy of Science*, 3, 207–220.
- Bolsen, T., Druckman, J. N., & Cook, F. L. (2014). How frames can undermine support for scientific adaptations: Politicization and the status-quo bias. *Public Opinion Quarterly*, 78(1), 1–26.
- Boulcault, M., & Schroeder, S. A. (2021). Public trust in science: Exploring the idiosyncrasy-free ideal. In *Social Trust* (pp. 102–121). Routledge.
- Budescu, D., Broomell, S., & Por, H.-H. (2009). Improving communication of uncertainty in the reports of the intergovernmental panel on climate change. *Psychological Science*, 20(3), 299–308.
- Christiano, T. (2008). *The constitution of equality: Democratic authority and its limits*. Oxford University Press.
- Christiano, T. (2012). Rational deliberation among experts and citizens. In J. Parkinson & J. Mansbridge (Eds.), *Deliberative Systems* (pp. 27–51). Cambridge University Press.
- Churchman, C. W. (1948). Statistics, pragmatics, induction. *Philosophy of Science*, 15(3), 249–268.
- de Melo-Martín, I., & Intemann, K. (2016). The risk of using inductive risk to challenge the value-free ideal. *Philosophy of Science*, 83(4), 500–520.
- Douglas, H. (2000). Inductive risk and values in science. *Philosophy of Science*, 67(4), 559–579.
- Douglas, H. (2008). The role of values in expert reasoning. *Public Affairs Quarterly*, 22(1), 1–18.
- Douglas, H. (2009). *Science, policy, and the value-free ideal*. University of Pittsburgh Press.
- Elliott, K. C. (2006). An ethics of expertise based on informed consent. *Science and Engineering Ethics*, 12(4), 637–661.
- Elliott, K. C. (2010). Hydrogen fuel-cell vehicles, energy policy, and the ethics of expertise. *Journal of Applied Philosophy*, 27(4), 376–393.
- Elliott, K. C., & McKaughan, D. J. (2014). Nonepistemic values and the multiple goals of science. *Philosophy of Science*, 81(1), 1–21.

- Elliott, K. C., & Richards, T. (2017). *Exploring inductive risk: Case studies of values in science*. Oxford University Press.
- Elliott, K. C., & Richards, T. (2017). Exploring inductive risk: Future questions. In K. C. Elliott & T. Richards (Eds.), *Exploring inductive risk: Case studies of values in science* (pp. 261–277). Oxford: Oxford University Press.
- Estlund, D. (2009). *Democratic authority: A philosophical framework*. Princeton University Press.
- Fischhoff, B. (1995). Risk perception and communication unplugged: Twenty years of process. *Risk Analysis*, 15(2), 137–145.
- Franco, P. L. (2017). Assertion, nonepistemic values, and scientific practice. *Philosophy of Science*, 84(1), 160–180.
- Frank, D. M. (2017). Making uncertainties explicit. In K. C. Elliott & T. Richards (Eds.), *Exploring inductive risk: Case studies of values in science* (pp. 79–100). Oxford University Press.
- Grice, P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and Semantics: Vol. 3: Speech Acts* (pp. 41–58). Academic Press.
- Havstad, J. C., & Brown, M. J. (2017). Inductive risk, deferred decisions, and climate science advising. In K. C. Elliott & T. Richards (Eds.), *Exploring inductive risk: Case studies of values in science* (pp. 101–123). Oxford University Press.
- IGBP, SCOR, & IOC. (2013). *Ocean Acidification Summary for Policymakers—Third Symposium on the Ocean in a High-CO2 World*. International Geosphere-Biosphere Programme.
- IPCC. (2014). *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. IPCC.
- Irzik, G., & Kurtulmus, F. (2019). What is epistemic public trust in science? *British Journal for the Philosophy of Science*, 70(4), 1145–1166.
- Jaspert, K., Van de Velde, F., Brône, G., Feyaerts, K., & Geeraerts, D. (2011). Does framing work? an empirical study of simplifying models for sustainable food production. *Cognitive Linguistics*, 22(3), 459–490.
- Jeffrey, R. (1956). Valuation and acceptance of scientific hypotheses. *Philosophy of Science*, 23(3), 237–246.
- John, S. (2018). Epistemic trust and the ethics of science communication: Against transparency, openness, sincerity and honesty. *Social Epistemology*, 32(2), 75–87.
- Stephen, J. (2019). Science, truth and dictatorship: Wishful thinking or wishful speaking? *Studies in History and Philosophy of Science Part A*, 78, 64–72.
- Johnson, B. B., & Slovic, P. (1995). Presenting uncertainty in health risk assessment: Initial studies of its effects on risk perception and trust. *Risk Analysis*, 15(4), 485–494.
- Kohut, A., Keeter, S., Doherty, C., & Dimmock, M. (2009). *Scientific achievements less prominent than a decade ago*. Pew Research Centre for the People & the Press.
- Kolodny, N. (2014). Rule over none ii: Social equality and the justification of democracy. *Philosophy & Public Affairs*, 42(4), 287–336.
- Kupfer, J. (1982). The moral presumption against lying. *The Review of Metaphysics*, 103–126.
- Nisbet, M. C. (2009). The ethics of framing science. In B. Nerlich, R. Elliott, & B. Larson (Eds.), *Communicating biological sciences: Ethical and metaphorical dimensions* (pp. 51–73). Routledge.
- OECD. (2015). Scientific advice for policy making: The role and responsibility of expert bodies and individual scientists. *OECD Science, Technology and Industry Policy Papers*.
- Parker, W., & Lusk, G. (2019). Incorporating user values into climate services. *Bulletin of the American Meteorological Society*, 100(9), 1643–1650.
- Patt, A., & Schrag, D. (2003). Using specific language to describe risk and probability. *Climatic Change*, 61(1–2), 17–30.
- Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford University Press.
- Pielke, R. A. (2007). *The honest broker: making sense of science in policy and politics*. Cambridge University Press.
- Railton, P. (1984). Alienation, consequentialism, and the demands of morality. *Philosophy & Public Affairs*, 134–171.
- Resnik, D. B. (2001). Ethical dilemmas in communicating medical information to the public. *Health Policy*, 55(2), 129–149.
- Rudner, R. (1953). The scientist qua scientist makes value judgments. *Philosophy of Science*, 20(1), 1–6.
- Sartori, G. (1987). *The Theory of Democracy Revisited*. Chatham House.
- Schroeder, S. A. (2020). Democratic values: A better foundation for public trust in science. *The British Journal for the Philosophy of Science*.

- Shiffrin, S. V. (2016). *Speech matters: On lying, morality, and the law*. Princeton University Press.
- Shrader-Frechette, K. S. (1994). *Ethics of scientific research*. Rowman & Littlefield.
- 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.
- Steele, K. (2012). The scientist qua policy advisor makes value judgments. *Philosophy of Science*, 79(5), 893–904.
- Valkenburg, P. M., Semetko, H. A., & De Vreese, C. H. (1999). The effects of news frames on readers' thoughts and recall. *Communication Research*, 26(5), 550–569.
- Vivald, E., & Coville, A. How do policymakers update their beliefs?.
- Weber, E., & Hilton, D. (1990). Contextual effects in the interpretations of probability words: Perceived base rate and severity of events. *Journal of Experimental Psychology: Human Perception and Performance*, 16(4), 781–789.
- Weingart, P., & Lentsch, J. (2009). *Scientific advice to policy making: International Comparisons*. Opladen/Farmington Hills.
- Wilholt, T. (2013). Epistemic trust in science. *The British Journal for the Philosophy of Science*, 64(2), 233–253.
- Wilson, D., & Sperber, D. (2012). *Meaning and Relevance*. Cambridge University Press.
- Yang, Y., & Hobbs, J. E. (2020). The power of stories: Narratives and information framing effects in science communication. *American Journal of Agricultural Economics*, 102(4), 1271–1296.

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