To appear in: Gonzalez, W. J. (ed.), 'Climate Change and Studies of the Future'

Roles for scientists in policymaking

Joe Roussos[1]

April 2023

Abstract. What is the proper role for scientists in policymaking? This paper explores various roles that scientists can play, with an eye to questions that these roles raise about value-neutrality and technocracy. Where much philosophical literature is concerned with the conduct of research or the transmission of research results to policymakers, I am interested in various non-research roles that scientists take on in policymaking. These include raising the alarm on issues, framing and conceptualising problems, formulating potential policies, assessing policy options for expected efficacy, and more. I consider examples from climate change and Covid-19 policymaking. My intention is to encourage philosophers to expand their interest in values in science out from the conduct of research to the wide array of roles that scientists play in policymaking. The paper is therefore an overview of the landscape of potential research questions, rather than a presentation of a single argument.

1 Introduction

What is the proper role for scientists in policymaking? During the Covid-19 pandemic, many governments proudly reported that they were "following the science", implying a central and guiding role for scientists in at least this policy arena. In this essay I explore various roles scientists can play in policymaking, with an eye to two philosophical issues: values in science and technocracy. I am interested in a broad range of science-based policy areas, but I draw my examples from climate change and the Covid-19 pandemic. What these two policy arenas have in common is the great uncertainty facing decision-makers as well as the high stakes of the decisions. Covid-19 was of course a higher urgency problem, in the sense that its harmful consequences were felt on a scale of months whereas the harmful consequences of climate change will be felt over decades.

I am motivated in part by a particular issue gaining currency in the literature on climate change adaptation. "Off the shelf" academic information is often inaccessible to decision-makers and unsuitable for direct use. This has led to calls for increased "usability", with the suggestion that "co-production of climate information" might secure this (Kirchhoff, Lemos, and Dessai 2013; Jebeile and Roussos 2023). Broadly speaking, this is often interpreted as a call for scientists

[1] Joe.roussos@iffs.se, Institute for Futures Studies, Stockholm

to engage more closely with policy. There is significant discussion of what this means in terms of research—including various models of stakeholder engagement, citizen science, and so forth. My interest, however, is in other research-adjacent or non-research activities that scientists engage in, such as assessment of evidence, participation in expert panels, submission to regulatory bodies, and direct work with policymakers.

The starting point for this paper is two linked ideas about the role of expert advisors, and specifically scientists, in policy decision-making.

The first is value-free ideal, which says roughly that the core of science should be free from non-epistemic values. This is a contrast class with "epistemic values", a set of values the pursuit of which is thought to support key mission of science: finding the truth. Examples of epistemic virtues are simplicity, breadth, explanatoriness. By contrast non-epistemic values are things like justice, equity, welfare, respecting the dignity of people, and so on. The "core" of science here refers to evidence gathering and processing, and the acceptance of hypotheses.

The second background idea is one of division of labour. Two famous, and perhaps apocryphal, quotations by British politicians serve to illustrate the presumed division between scientists and political officials. According to Margaret Thatcher "advisors advise, ministers decide." Or, in the words of Winston Churchill, "experts should be on tap, but not on top." For broadly democracy-promoting reasons, experts should not make decisions, though they do have a crucial role in informing them.

Together, these two ideas form what I will call *the neutral advisor ideal for scientists in policymaking*. The ideal isn't always clearly articulated, but it is often taken to imply that scientists ought to be restricted to providing (certain kinds of) value-free information to policymakers, rather than playing a larger role in decision-making or a more value-laden one. The neutral advisor ideal is, I think, prevalent among research scientists, science advisors, science scholars and philosophers. Discussions of the ideal often lump together different things that scientists are meant to ensure don't influence their research work and advice: financial interests, personal biases, political ideologies, moral commitments, democratically-informed values.

Is this vision of a neutral advisor possible? Is it desirable? My survey of the roles scientists play will be conducted with an eye to these questions. These aren't questions of the scientist qua scientist but, in Katie Steele's (2012) language, of the scientist qua policy advisor.[2]

Much work has been done in philosophy of science to show that science qua science is not and cannot be value-free (or, in some cases, ought not be). This is valuable work. But it has two restrictions, from my perspective. First, it focuses on research and often fails to consider the many non-research roles scientists play. Second, it is overly focused on a particular site of values in science: the inductive risk involved in making scientific inferences and choosing methodologies. Inductive risk has come to dominate discussions of values in science, and is now the default analytical frame for many philosophers even when thinking about science-for-

[2] But where Steele (2012) focuses on how communicating scientific results to policymakers may necessitate value judgements, I consider a wider set of actions the scientist may take as policy advisor.

policy—cf. Douglas's (2009) landmark book and, more recently, Parker and Lusk's (2019) work on values in climate services.

The two limitations interact: the focus on inductive risk is supported by a common but narrow view of the role of scientists in policymaking. The "scientific information" that advisors are meant to provide under the neutral advisor ideal is often assumed to be information about the world is or might be. This is certainly a common and useful form of scientific information—science is, of course, centrally engaged in making predictions. We can frame this view in the language of decision theory: scientists should supply policymakers with information about which states of the world are possible and should provide probabilities for them. Thus, we discuss whether probabilities in climate science are objective or subjective, and how they should inform decision-making (Winsberg 2018). We worry about whether they are free from the values of the individual scientists. Implicitly, such discussions assume that the remainder of the decision process—or, put another way, the other elements of a decision called for by decision theory—is the domain of the policymaker and therefore that this remainder does not raise questions about the value-neutrality of science.

But, in reality, scientists play a much wider role in policymaking. They help policymakers develop their understanding of which questions are relevant and which decisions are important. They help conceptualise options, by proposing actions (using their knowledge of the mechanisms in the system) or commenting on which putative actions are likely to be efficacious. They provide various tools to help non-experts achieve (partial) understanding, including heuristics, simplified scenarios, analogies, and narratives. Scientists from different disciplines may get involved with different parts of the decision process—some proposing what is likely to happen in the relevant system; others answering questions about how the system will respond to changes; still others relating proposed policy measures (like introducing a tax) to variables of interest (like emissions). In order to "join up" this process, each scientist may be asked to give input on other parts, despite the work falling outside of their core area of expertise.

These roles can be similarly housed within a decision theory framework, and we can use tools from decision science more broadly to think about where and how these roles raise questions of value-neutrality. I will therefore structure the core part of my discussion, in §§4–**Error! Reference source not found.**, around the "decision table", an analytical tool from decision theory. My goal in doing this is twofold. First, I think it is a helpful expository tool for organising the different roles scientists can play. Second, I think that some of the focus on inductive risk and on probabilities comes from the fact that philosophers find technical analyses and formal tools attractive. By showing that these further roles for scientists can be embedded in a decision-theoretic framework I hope to bask in some of that glow. Before this, I begin with two brief discussions of the literature. In §2 I discuss the philosophical literature on values in science, and in §3 I discuss the science-for-policy literature, which is mostly written by non-philosophers. Neither is intended to be a complete review, and the science-for-policy literature in particular is very large.

2 Values in science: epistemic and democratic arguments

In this section I will provide a potted history of some of the discussion about the role of values in science that has taken place in philosophy, with an eye to lessons that can be used when discussing science-for-policy. It isn't intended to be comprehensive, so much as to draw out some trends that I see as significant for my purpose here.

To start, a caveat: there is a significant discussion of semantic fact-value entanglement, which is important for understanding how definitions of key scientific terms can build in value-commitments. I do not discuss this here, focussing instead on the discussion of value-freedom in the methods of science.

Discussions of values in science often conceive of science as a relatively isolated practice of producing research, which only interacts with the rest of the world at either end of the process. For example, Max Weber distinguished four stages at which values might influence science (Weber (1949b) discussed by Reiss and Sprenger (2017)). The stages are:

i. the choice of a scientific research problem,
ii. the gathering of evidence in relation to the problem,
iii. the acceptance of a scientific hypothesis or theory as an adequate answer to the problem on the basis of the evidence,
iv. the proliferation and application of scientific research results.

Social and political values can legitimately play a role in stages (i) and (iv), the thinking goes. Indeed, much attention in the values in science literature focuses on stage (iii). Heather Douglas (2000) has shown that this process is too simple and highlighted stages between (ii) and (iii), involving choice of methods, classification of evidence, and choice of inference rules, where values may play a role. But nevertheless, the focus is often on the "core" of science as it was here that philosophers perceived the greatest threat from value-ladenness.

One can understand the literature on this question in terms of two questions. The first is whether it is *possible* for science (implicitly: in the core stages (ii) and (iii)) to conducted without the use of a class of values often called non-epistemic values. (The thought behind this terminology is that there are epistemic values which can and should guide science. But a value for, e.g., human welfare, is non-epistemic and should not be a part of science.) The second is whether it is *desirable* for science to be free from these values. The view that it is possible and desirable for science to be value-free is called "the value-free ideal for science." Recent philosophical work on the value-free ideal has focused on arguments against it, either by arguing against the possibility of value-free science or by arguing that value-free science is undesirable. I will briefly outline various positions that philosophers have taken against the value-free ideal.

A first group argue against the possibility of value-free science, on the grounds that it is not possible to perform inferences that are a necessary part of science without making non-epistemic value judgements. Richard Rudner's classic (1953) paper argues that the acceptance of a hypothesis requires a judgement about how strong evidence would need to be, and that this in turn requires evaluating "importance, in the typically ethical sense, of making a mistake in accepting or rejecting the hypothesis" (Rudner 1953, 2). Helen Longino (1996) makes a more general argument about the underdetermination of theory by evidence, arguing that in choosing a theory scientists must appeal to what I am here calling non-epistemic values. Similar arguments about the necessity of value judgements can be found made by Carl Hempel (1965). There is a long tradition of arguing that the social sciences in particular cannot be value-free because of the crucial role that the scientist's point of view plays in structuring social scientific investigation (Weber 1949a; Reiss and Sprenger 2017).

A second group argue that scientists ought to make value judgements. Heather Douglas argues that scientists have the same moral obligations as non-scientists, and that they ought to consider the downstream implications of their decisions (Douglas 2000; 2009). Thus, when evaluating the inductive risk in choosing a method, classifying evidence, selecting an inference rule, or accepting a hypothesis, scientists should make ethical considerations. Longino (1996) sometimes also argues this way.

A third group argue that the value-free ideal is not a good way of pursuing the truth. Rather society will do better at achieving this goal of science if it establishes institutions of science which allow for and exploit the foibles and values of individual scientists. I place Philip Kitcher (1990), Sandra Harding (1995), and Longino (2004), in this group.

A fourth group argue that the value-free ideal is not a good way of achieving other important goals. Kitcher (2001; 2011) sees science as part of democracy, and "well-ordered science" will be directed towards societal goals, infused with democratic values and a concern for human welfare. While much of his focus is on the direction of research and allocation of funding (Weber's stage (i)) but includes helpful discussions of the use of science by society. On this view, values play a necessary and important role in directing science to fulfil its purpose. Neutrality is replaced by transparency.

Much of this work (though notably not Douglas) focuses on academic research science. Many philosophers and philosophical works don't fall neatly into one of these four groups. For example, Matt Brown's recent (2019) argument against value-free science crosses several of these boundaries.

So much for arguments against the value-free ideal. But presumably in order to attract such attention the ideal must be powerfully attractive. What are the arguments for it? Unsurprisingly, there are different motivations for it. One category focuses on the harm values might to do science's aim of learning about the way the world is. The worry is that scientists might look for answers which accord with their values, effectively predetermining the conclusions of their research. Science ought to accept conclusions on the basis of evidence, and "values are not evidence; wishing does not make it so" (Douglas 2009, 87). The problem is that allowing values to determine the conclusions we accept would be tantamount to wishful thinking. As Douglas herself notes, this is a compelling argument for a *restriction* on the roles values ought to play in science: they should not replace evidence. But this argument leaves open the possibility that values can reasonably play some smaller role, such as Douglas's "indirect role" in inference: acting only once all evidence has been accounted, and only if there is a gap left over by that evidence.

Another set of motivations, perhaps more important for my purposes, are what I will call democratic defences of the value-free ideal. I will consider two such defences here: one due to W. E. B. Du Bois and the other emerging from the political theory literature on technocracy.

Liam Kofi Bright (2018) argues that Du Bois (1898) held that value-free science was necessary in order to ensure public trust in science and facilitate its use for its intended purpose. Du Bois held that "the mediate aim of science [i.e., the purpose for which society undertakes it] … was to provide information that can fruitfully be used to guide policy in democratic states." (Bright 2018, 2231). I read these arguments as being concerned with the proper functioning of science. (In this way, Du Bois is most naturally in conversation with modern authors like Kitcher.) In order to fulfil this aim, science must be trusted by the public—so that they allow the research to occur and so that they are willing to enact its results. Securing public trust, in turn, requires that scientists be motivated only by the search for truth—i.e., that science be free from non-epistemic values.

The anti-technocracy motivation is implicit in much of the political theory literature on the role of experts in democracy. Policy decision-making is part of the democratic process and must include (democratically selected) values. What makes a policy decision the right one *democratically* is that it responds appropriately to the values that the policymaker represents, in virtue of representing the people who democratically selected them (or, if they are civil servant, the political authorities who direct their efforts). We want our policies to be evidence-based, but we want the values that drive those decisions to come from the democratic process. If science is tinged with the values of the individual scientists who conduct it, then our policy decisions will be influenced by the (non-epistemic) values of those scientists. But they were not democratically selected (nor should they be!), and their values are unlikely to be representative of the demos. (Or simply *are not representative*, depending on one's favoured analysis of democratic representation.) Thus, the thinking goes, science must be value-free lest it undermine the proper functioning of democracy. (Note that this is argument is about undue influence on a process in which official decision-making authority rests with policymakers. It is not about scientists directly making decisions, e.g., in regulatory settings and the technocratic threat that this poses (cf. Salter 1988; Jasanoff 1990).)

My focus is on these defences of the value-free ideal because the goods they attempt to secure are those which seem most directly threatened by the potentially problematic engagement of scientists in policymaking. These democratic defences are not particularly concerned with "pure" science or with scientific research as opposed to other scientific activities. Du Bois is partly interested in the perception of science and scientists, which is surely a function of all of their perceptible behaviours. Political theorists, on the other hand, are concerned with control of democratic decision-making. Research is probably the least effective avenue by which scientists could exert unsavoury political control.

3 Values in science-for-policy and regulatory science

An important role that science plays is in advising policy decisions. One way this happens is through the direct involvement of scientists in a policymaking or regulatory process, by making submissions, participating in expert panels, or acting as advisors. Although this is a very different domain than research, the value-free ideal is once again prevalent—often as part of the more complex neutral advisor ideal. The difficulties around maintaining neutrality inform the decisions that many make not to contribute to policy discussions beyond publishing their research and hoping that it has downstream effect (Lach et al. 2003). Former deputy director of the EPA's national laboratory, Robert Lackey, has gone so far as to call value-laden science in a policy advising context a kind of corruption (Lackey 2004).

In this section I summarise some literature relevant to the neutral advisor ideal, from a handful of related disciplines: policy, STS, and sociology of science. To begin, it will be helpful to introduce a simple but common framework for thinking about policymaking called "the policy cycle". Many versions of it exist, but there is significant overlap between them. The framework should be understood as a simplified and idealised description, which neglects considerable messiness and complexity present in real policymaking. Here are some common stages included in the policy cycle.

1. Problem emergence: the identification of an issue, growth of awareness of and discussion about it.
2. Agenda setting: the entry of an issue into the decision agenda of political actors. Often grouped with (1).
3. Formulation of policy options: the problem is formally characterised and analysed, and solutions are proposed.
4. Decision-making: a governmental decision is taken on a policy or policies to address the problem.
5. Implementation: actors, typically within the government, implement the policy.
6. Evaluation: the state of the problem is assessed, and the policy's implementation is examined for efficacy.

An immediate observation we can make is that scientists get involved at many different stages of the policy cycle.

1. Problem emergence: scientists can uncover a problem, e.g., Climate Change, or they can be asked to investigate a potential problem, e.g., acid rain.
2. Agenda setting: scientists can lobby for a problem to be taken seriously, e.g., climate change in the era before the IPCC.
3. Formulation of policy options: scientists can use scientific knowledge to suggest and develop policy options, e.g., carbon capture and storage.
4. Decision-making: scientists can support a decision process by working as advisors, either directly with policymakers or on expert panels, supplying evidence on the nature of the problem, its expected evolution, potential policy efficacy, etc.

5. Implementation: scientists can work with implementation teams, providing more detailed and granular information and advice as the policy is fleshed out and as issues arise in implementation.
6. Evaluation: this is often a social scientific exercise.

As the above list indicates, scientists get involved through different channels: expert advisory committees, individual "science advisor" positions, in-house science teams in government departments, and scientific associations—not to mention their "standard" role as researchers whose work might be used by others doing policy work. Each of the above stages has distinct entry points for potentially problematic value-ladenness, related to the different tasks that scientists might get involved in. I cannot address all these details here, and will instead try to draw out significant themes.

There are several attempts to think about the objectivity of scientists in policy which cut across the policy cycle. One popular approach is due to Roger Pielke Jr (2007), who introduced four "stances" that scientists can take when engaging with policy. (These stances could be taken by individual scientists or e.g., interdisciplinary advisory councils.) He uses a toy example of choosing a restaurant to dine at for illustration. The first stance is the Pure Scientist, a disinterested approach that will make available whatever off-the-shelf scientific research seems fitting. To help you choose a restaurant, the Pure Scientist might offer research on the fundamentals of nutrition. The second stance is the Science Arbiter, in which the scientist positions themself as a somewhat more helpful resource. You can ask the Science Arbiter specific questions, like "What is the closest vegan restaurant?" and they will supply you with the best scientific answer they can. The Science Arbiter does not tell the decision-maker what they ought to prefer, which differentiates them from the third stance: Issue Advocate. The Issue Advocate will try to convince you to eat at a particular restaurant, or in a less-extreme case to eat at a category of restaurant, e.g., Thai. The scientist who engages in Issue Advocacy may think that they understand the decision-makers interests particularly well, or have personal or political motivations for supporting a particular option. The Issue Advocate tries to limit the scope of the decision-maker's choice to one or a few options. By contrast, the Honest Broker of Policy Alternatives—the final stance—tries to expand the scope of choice. The Honest Broker provides information on all relevant restaurants, along with whatever information is helpful in deciding between them. Choice is left to the decision-maker; the Honest Broker tries to empower and inform them. The Broker is more active than the Arbiter; the former tries to understand the nature of the decision-maker's problem and engage in creating answers which fit, while the latter is more like a human search engine.

Pielke regards all stances as valuable and necessary in certain contexts.[3] Broadly speaking, he thinks that less engaged stances, such as the Pure Scientist and Science Arbiter, are suited to decision contexts where there is agreement on values and low uncertainty (Pielke 2007, 76–77).

[3] That said, Pielke also argues that the Pure Scientist is rarely seen in reality and is often a rhetorical device masking stealth Issue Advocacy.

These stances are viewed by many scientists as safer than more engaged stances (Lach et al. 2003), regardless of the context. They are the natural stances for scientists under what it called the "linear model" of the relation between science and society, in which science first does its work and achieves scientific knowledge and this knowledge is then transferred to the policy arena. Under this model it is often presumed that consensus on the science is necessary and sufficient for political consensus (Pielke 2007, 77–79).

But they are also seen as insufficient for difficult areas of policymaking. Issue Advocacy and Honest Broker stances are required when values are contested, and uncertainty is high. "Cooperation between scientists and politicians is of particular importance when the knowledge of experts is contested or uncertain and when political party lines are ill-defined with respect to an issue", writes Wolfgang Krohn (2005) in a review of a type of European body called an Enquete Commission, which involves both experts and political representatives. In much of the literature that followed Pielke, the Honest Broker is taken as the ideal stance for scientists, with many citing the "opening up" function as crucial to democratic decision making (e.g., Moore 2017, 46).

Issue Advocacy is controversial. Pielke (2007, 94) opposes stealth advocacy but thinks that open advocacy is right when scientists feel strongly about policy issues. Scientists who work in policy often regard any advocacy as unethical (e.g., Lackey 2004; 2007 and references therein). However, as Kitcher has noted, the discussion about values in policymaking tends to blur together two distinct issues related to values in science-for-policy: (1) scientists "judging that the data are good enough, given the more or less precisely envisaged consequences for human welfare, and [2] accepting a conclusion because it will make you money or advance your favoured political cause" (Kitcher 2011, 164). The latter seems straightforwardly problematic, while the former is subject to much debate.

Liora Salter makes a similar point, using three categories: "We believe it is essential to distinguish between science and values, science and interests, and science and bias, even though all these relationships are often described in terms of science and values" (Salter 1988, 194). In her usage, "values" refers roughly to moral considerations, and become problematic when used by scientists in advocating for a policy, based on their research but where the moral considerations are a necessary bridge between the research findings and the advocated policy. In the kind of regulatory science she studied, "the problem in the relationship between science and advocacy [arises] not in the conduct of the research but in its presentation to a non-scientific community" (Salter 1988, 193). "Interests" refers to the interests and desires of various stakeholder groups. Salter's study is about regulatory standards setting and so the interest groups include companies, workers in factories, consumers, and so forth. Science can be conducted with certain interests in mind, for example the safety of workers in a factory which uses a potentially dangerous chemical. Scientists can safely express conclusions in terms of the interests of various parties without undermining their science. "Bias" refers to straightforwardly unethical practices: scientists who are "bought", data which is falsified, or key facts which are strategically omitted in a regulatory submission.

Writers on science-for-policy have spent significant time on the problem of technocracy, or government by unelected expert elites. A classic study on US science advising is Sheila Jasanoff's *Fifth Branch*. Jasanoff's study is social scientific, and part of her aim is to understand and illuminate the roles scientists play and the influence they wield. But she has normative concerns, particularly related to democracy. The title of her book signals that experts have become an unofficial "fifth branch" of the government and play a significant, and potentially worrying, role in decision-making. The work that scientists do on advisory committees, says Jasanoff, is "not 'science' in any ordinary sense, but a hybrid activity that combines elements of scientific evidence and reasoning with large doses of social and political judgement" (Jasanoff 1990, 229).

This is despite the careful specification of roles that the government agencies commissioning these committees have sometimes attempted. For example, Jasanoff describes an episode in which scientists working for the US Environmental Protection Agency (EPA) are instructed to assess the safety of a particular pesticide but end up acting as broader policy analysts. They are embedded within an EPA process, they understand its broader aims, and they naturally come to consider the downstream implications of their judgement. If this pesticide is banned, what alternatives remain? Will a judgement to ban in this case lead to similar judgements down the line? What impact will this have on this whole area of regulation and industrial activity? These judgements are reminiscent of those Douglas (2009) discusses, but are not necessarily restricted to what Douglas calls an "indirect" role, nor are they framed in terms of inductive risk.

At other times, the agencies themselves place scientists in positions of administrative authority. Jasanoff discusses the US Food and Drug Administration's use of expert panels in place of judges in the 1980's to evaluate the technical claims of regulatory adversaries and pass judgement on their claims (1990, 222–26).

Regardless of the administration's openness to technocratic delegation,

> Scientists in the aggregate wield influence, and they do so, moreover, through proceedings that lack many of the safeguards of classic administrative decision-making. Participation of lay interests is limited and often one-sided, cross-examination is almost unknown, and committee recommendations, however much weight they carry, are seldom accompanied by detailed explanations or considerations of alternatives. (Jasanoff 1990, 229)

Jasanoff's view of value-freedom in regulatory science is pessimistic: it is not possible to restrict scientist's to technical issues, nor is it plausible that their subjective values are irrelevant (1990, 230).

Pielke, too, is negative about the possibility of value-free science. In his view, "objectivity is more possible in cases where the decision context is highly specified or constrained … in circumstances where the scope of choice is fixed and the decision-maker has a clearly defined technical question" (2007, 6). Under such conditions, scientists are able to act as Science

Arbiters, which Pielke takes to be the most objective stance. In the absence of these conditions, and in particular when there is no consensus on the values to be served "there is very little room for arbitrating science in the process of decision-making, and even good faith efforts to provide such a perspective can easily turn into a political battleground where political debate is couched in the guise of a debate over science (and the expert may not even be aware of his/her arguing politics through science)" (Pielke 2007, 6).

This completes my extract of salient bits of the literature. It highlights the ways in which the discussion of science advisors in policy is more complex than that of values in scientific research, and less focused on the epistemic goals of science (though these are of course central to the authority science wields). We can now turn to the core questions for the neutral science advisor ideal: (1) Is it possible? (2) Is it desirable? It should be clear now that answering these questions is quite distinct from answering the analogous questions about the value-free ideal for research science. We will need to pay proper attention to the roles that scientists actually play in policymaking, work out what the relevant challenges are with respect to neutrality, and then attempt to answer them. This paper does not attempt to do all this, however. I aim merely to identify various roles scientists play and make some preliminary comments with respect to the role of values.

4 The decision table as framing device

This is a paper about the role of scientists in policy decision-making, but it is also a paper about what decision theory can add to analyses this topic. I will structure my discussion using one of the most basic tools of decision theory: the decision table. This is a representation of a decision and, in particular, I will use it to represent the decision facing the policymaker. At a very high level I will assume that policymakers face some problem or challenge and that they are engaged in formulating a policy response to it (though I will also consider how issues come to be on the policymaker's agenda). I will represent this as a decision in which the policymaker selects one act (policy) from a menu of options. Taking an option (enacting a policy) is associated with a number of possible consequences or outcomes. These outcomes are more or less desirable, from the perspective of the policymaker's aims and values, and we represent this with a partial ranking of the outcomes. Which outcomes actually results from taking an option depends on the way the world is—what else happens, the way the policy interacts with the problem, people's responses to the policy's enactment, and so forth. We use the term "state of the world" to refer to a description of all the features that are relevant to determining which outcome results from taking an option. A decision problem can be represented by a table showing, for each available option, the consequence that follows from its exercise in each state of the world.

Here is a simple example. You have an appointment that you don't want to miss. If you walk you will arrive a little late. If you take the bus and the traffic is light, you should arrive a little ahead of time. On the other hand, if the traffic is heavy then you will arrive very late, perhaps so late that the appointment will be lost. Is it worth risking it? This decision is represented in Table 1.

Table 1. Example decision: take a bus

	LIGHT TRAFFIC	HEAVY TRAFFIC
Take a bus	Arrive on time, cool and calm	Arrive significantly late
Run	Arrive slightly late and sweaty	Arrive slightly late and sweaty

There is, in principle, no restriction on the number of options a policymaker might consider, nor the number of states that might be relevant to their decision. So a general decision table looks like Table 2.

Table 2. General decision table

	STATE 1	…	STATE N
Option 1	Consequence 11	…	Consequence 1n
…	…	…	…
Option m	Consequence m1	…	Consequence mn

In orthodox decision theories, the option which ought to be chosen is the one with the highest value (or any one of those, if there is a tie). The value of an option is a function of its consequences and the states in which they come about. In Bayesian decision theory the value of an option is its subjective expected utility. Each state of the world is assigned a probability, representing the decision-maker's subjective assessment of its likelihood. (In some variations, the probabilities are assigned to what I am calling consequences and are conditional probabilities: the probability of the state conditional on each option being chosen.) The consequences are assigned utilities, representing the desirability of each consequence according to the decision-maker. This is displayed mathematically in Table 1.

Table 3. General decision table with probabilities, utilities, and option values

	Pr(S1)	…	Pr(Sn)	
O1	U(C11)	…	U(C1n)	$V(O1)=E_{Pr}[U(C1j)]$
…	…	…	…	…
Om	U(Cm1)	…	U(Cmn)	$V(Om)=E_{Pr}[U(Cmj)]$

When analysing a policy decision, the "personal" elements of this table will be replaced by something more collective. Instead of having utilities represent the preferences of an individual, the evaluative elements might represent collectively determined values, as interpreted by democratic representatives. Instead of using probabilities which represent subjective credal states, we might want the probabilities to reflect scientific or scientifically informed judgements of likelihood. (These might be generated from the subjective probabilities of multiple experts, but they aren't the decision maker's own.)

I should point out that actual policy analyses will rarely, if ever, utilise a table of this form. I make no claim that this tool is useful in the practice of policymaking, or that it captures everything that goes into a real decision. My purpose in using it is to structure the coming discussion of inputs to a policy decision, and the roles scientists can play therein.

In philosophical decision theory, these tables are often specified at the outset of the discussion, and they are, for the most part, tools for facilitating a discussion of some principle of rationality or decision. But in decision analysis and support, formulating and filling in these tables is a difficult task. When I teach decision theory, I ask my students to represent a simple decision using a table like this. They inevitably find the task difficult, requiring several iterations and refinements. It is often unclear what the relevant partition of states is until the options are outlined. But what the relevant options are can depend on what the decision-maker is trying to achieve, and the states are part of specifying that. Consequences are occasionally described as simple conjunctions of options and states, which is helpful for determining the level of grain required when describing the states, but is *unhelpful* when first formulating the table because desired consequences are the part of the decision which are most accessible from the outset.

Here is a stylised example of the process of formulating a decision table. Suppose that you are responsible for London's Thames flood defences. Currently, a system called the Thames Barrier protects London from flooding. It consists of 10 movable metal gates, which can be raised to hold back water. Roughly speaking, the Barrier was designed to protect the city from a hundred-year flood. (This is a measure of severity in terms of probability: increasingly severe floods are decreasingly likely. The return period is the expected time between events of this severity. Another way of stating it is that the annual probability of a flood worse than this is 1/100 or 1%.) As sea levels rise due to climate change, the Barrier will become less protective against floods. A simple way of putting this is that the level of flooding it can protect against will decrease, so that instead of protecting against a hundred-year flood it might come to protect against only a fifty-year flood.

Suppose that you are deciding on a new system to replace the Barrier. To simplify I will suppose that you are replacing it with a similar system, but with taller, stronger gates. How much taller the gates should be will depend on how high the sea level rises. You want to protect London against a hundred-year flood relative to new sea levels—this is the desired consequence which we can use to begin constructing our decision table. We will call a barrier "safe" if it protects against a hundred-year flood.

	SEA LEVEL RISE BELOW XMM	**SEA LEVEL RISE ABOVE XMM**
Do nothing	Current barrier safe	Current barrier unsafe
?	?	?
?	?	?

Perhaps a first question is whether the sea level rise will be significant enough to require a new barrier at all. Some level of sea level rise will push the Barrier below the current safety threshold of a hundred-year flood. Exactly what that level of rise is, denoted X in the table, is a matter for scientific investigation. Let's suppose that X=30cm, for specificity. Now supposing that a new barrier is required, there are likely to be different estimates of how high the level will go. You might engage scientists to provide estimates along with probabilities. These estimates can come in different forms. Ideally, you would get a continuous distribution of levels of sea rise and associated probabilities. Much more likely, you will receive interval estimates of expected rise, e.g., 70-90cm of rise. In a good case, this will come with a level of confidence: e.g., the expected level of sea rise is 66% likely to fall within 70-90cm.

	R<30	**30<R<70**	**70<R<90**	**90<R**
Do nothing	Current barrier safe	Current barrier unsafe	Current barrier unsafe	Current barrier unsafe
?	?	?	?	?
?	?	?	?	?

This gives you some sense of the possibilities you face. With these estimated levels of sea rise in hand, you can begin to devise policy options: different replacement barriers which are suitable to different levels of sea rise. Which options you consider will depend on the states your scientists considered possible and on the probabilities they assign to those states. Let's suppose that things are such that 70cm is an important threshold—there is a replacement barrier which is effective up to 70cm, and another more expensive option which is tolerant above 70cm.

	R<30	**30<R<70**	**70<R**
Do nothing	CB sufficient	CB insufficient, damaging floods	CB insufficient, catastrophic floods
Build 70-tolerant barrier	Barrier sufficient but unnecessary	Sufficient and cost-effective	Insufficient, money wasted and damaging floods
Build >70-tolerant barrier	Barrier sufficient but unnecessary	Sufficient but excessive	Sufficient and necessary

Now, depending on the costs and the evaluations of different consequences, we may have a table with the right resolution to facilitate a successful decision.

This somewhat lengthy example is meant to highlight a few things. First, that the decision maker's conception of the decision changes as we proceed with the analysis. Second, that the menu of options depends in part on the expected states of the world, but once the options are on the table the relevant description of those states might change. Decision-making is often iterative: one identified rough states of the world, formulates options, refines the states, considers consequences, reconsiders options, etc. I noted at the outset that the decision table is an idealised tool, so what I intend here is not a reflection on the decision theorist's exercise of constructing such a table. Rather, these seem to me to be important features of decision making itself, brought to light by this somewhat artificial exercise.

4.1 Scientists at every stage

One of the aims of this paper is to highlight that scientists can and do helpfully participate throughout this iterative process, and not just in providing states of the world or probabilities for them. Scientists can help fill in each of the three major areas of the decision table. Starting with structure:

- Framing: giving input on what the decision is that needs to be made
- Columns: highlighting important possibilities, i.e., identify states of the world
- Rows: suggesting policy options
- Cells: working out consequences for policy options

They can also play a role in both of the attitudes central to decision-making.

- Beliefs: assigning probabilities to states
- Desires: assigning utilities to consequences

It may seem surprising that scientists are involved in the assignment of utilities, which is an evaluative task. Indeed, I do not mean that scientists are or should be involved in assessing how good various outcomes would be. What I mean is that scientists are often intimately involved in answering questions which are required in order for an evaluation to take place. Each of these roles will be discussed in the following sections.

5 Problem emergence and decision framing

Scientists serve a crucial function in making policymakers aware of potential issues and helping to formulate a well-described problem which can then be the target of policymaking.

We begin with identifying issues. Examples of scientists raising potential issues are numerous and obvious, climate change, for instance. This awareness raising is prior to policymaking, and presumably after (at least some) research, so it might seem to be unproblematic in terms of scientists involving themselves in an essentially political process. Nevertheless, scientists who attempt to conform with a severe value-free ideal often eschew this kind of engagement. This is possibly because it threatens to undermine their image as disinterested researchers, which is crucial to maintaining their credibility in other spheres of scientific activity, for example when providing scientific evidence *during* policymaking. This potential tension between scientists' engagement in problem emergence and their status as sources of reliable and objective information motivates for a further discussion of the roles scientists can play in problem emergence and decision framing.

There are several different stances scientists could take towards "raising the alarm" about a potential issue. The most disinterested stance might say that scientists should play no role in policy or politics, and therefore take no steps beyond publishing results in scientific journals. Borrowing Pielke's (2007) terminology, we can call this stance the Pure Scientist. Pure Scientists do not raise issues, but instead rely on an (often vaguely conceptualised) pipeline of dissemination, which is meant to ensure that the important results that they publish in academic venues filter through to public and political awareness.

Another of Pielke's roles has a useful analogy here: the Issue Advocate. This scientist would independently determine that there are potential problems and attempt to bring them to the attention of policymakers. This might involve writing and speaking in public venues, making submissions to legislative or regulatory bodies, or getting involved directly in science advising. Note the important difference between an Issue Advocate at the problem emergence stage and an Issue Advocate at the decision stage. As I am imagining it, the IA advocates for an issue to receive attention at the problem emergence stage. "The world is warming—take action!" At the decision stage, the IA advocates for a specific policy solution. "Implement a carbon tax!" An individual scientist may take on both of these roles, but they are easy to distinguish and nothing about the former implies the latter. These are very different and raise different questions about the role of scientists in democratic decision-making. In particular, the IA during problem emergence does not seem to raise the issues which concern Pielke.

The other two stances that Pielke highlights don't seem to have useful parallels during problem emergence. Regarding the Google-like Science Arbiter: how would the policymaker know what to ask about? With an issue on the table, questions become obvious. (Once you are choosing where to dine, you know to ask about restaurants.) But asking blindly about potential problems ("Do I need to worry about sea level rise?) seems too random to be a good strategy. If the policymaker were to ask more open-ended questions, such as "are there any relevant issues?", this would force the Science Arbiter to move into a more active stance.

The same goes for the Honest Broker, a stance which focuses on opening up the decision space. Prior to policymaking beginning, the analogous goals seem best served by most effectively getting issues onto the policy agenda. We might try to distinguish between a *specific* Issue Advocate and a more general Honest Broker of Potential Policy Problems. Where the IA is focused on garnering attention for a single problem or class of problems—like climate change—perhaps the Honest Broker is more of an equal opportunity issue raiser—"Pay attention to global warming *and* global poverty *and*…". But given that we are discussing scientists raising issues which emerge from their disciplinary expertise, this stance would dilute the scientist's authority and undermine their efficacy. Climate scientists can speak with authority when raising issues related to the climate, but they have no particular authority with respect to potential economic policy issues.

Presumably all issue raising involves value judgements. The scientist must do more than note that increasing greenhouse gas concentrations will lead to global warming, they must judge that global warming will have harmful effects on people, animals, the environment, etc. In a sense, these value judgements are external to both the scientific research and to the policy cycle that the alarm-raiser hopes to initiate. But they are a necessary component of issue raising. Scientists with different values will therefore raise different problems, and what one takes to be an issue another might think is not worthy of attention. If scientists play a crucial role in raising awareness of potential policy problems in democracies then we may wish to worry about which values are informing their issue raising behaviour.

We now turn to framing a problem for decision-making. So far, I have spoken of "problem emergence" or "issue advocacy" as though it is clear what the problem or issue is. But in these early days, part of the work of scientists, activists and policymakers will be defining the issue. This involves the process outlined above with my Thames Barrier example but goes further, since that example did not involve much focus on refining our understanding of the consequences of each option. This process can involve significant work and iteration. As James

M. Joyce says, "Choosing is really a two-stage process in which the agent first refines her view of the decision situation by thinking more carefully about her options and the world's state until she settles on the 'right' problem to solve and then endeavors to select the best available course of action by reflecting on her beliefs and desires in the context of this problem" (Joyce 1999, 72).

Decision theorists refer to this as the "framing" of a decision problem. The same decision, intuitively speaking, can be framed in different ways: by picking out different locations of benefit (people, places, times, etc.) and by resolving the possibilities at different levels of detail. Ideally this should not change the decision that gets made—decision theories should be tolerant to different representations of the same problem. In practice, however, this is not the case. For agents with resource constraints and only partial understanding, certain representations will be favoured because they clearly highlight the important factors on which the decision depends. Descriptions which are too sparse risk omitting important features, and thus failing to prompt the necessary considerations with the attitudinal portions of the table are filled in. Descriptions which are too detailed risk confusing or overwhelming decision makers (Bradley 2014, 4–7).

As we have seen above, arriving at a framing is an interactive and iterative process, where each aspect of the table is refined as other aspects become clearer. Which scientists are involved, and at which stages, can significantly influence the framing of a decision problem. Consider a flood risk, which is brought to a governmental authority's attention by water system scientists who present the problem in terms of the risk to a village, in terms of harms to people and property. The initial framing analysis involves scoping out possibilities, identifying categories of potential harms, and brainstorming initial policy options. One such option in our flood scenario might involve deliberately flooding a so-called "diversion area"—land which is away from the village where water can be directed so that it does not flood areas of concern. If the problem is initially framed in terms of harms to people and property in the village, this may neglect or downplay other stakeholders and other categories of harm, such as harms to agricultural land and livestock in the diversion area.

These concerns are not specific to science-led policy, nor to the science advising portion of policymaking in general. They are central to good democratic policymaking. However, there are two reasons to pay attention to the roles of scientists and to consider what constraints ought to guide their involvement. First, scientific-technical advice is often sought even when stakeholders such as impacted residents and farmers are not involved. So, scientists can act as a proxy, if a partial one, for stakeholders. A water management specialist might be well acquainted with the concerns of agricultural landholders and be in a position to raise their concerns in a discussion of flooding a diversion area. This could be done well either through the conscious deployment of representative values (e.g., as elicited via a public consultation) or through a diverse body of scientists whose values naturally represent those of the stakeholders. Second, scientists may get involved with a narrow scientific-technical purpose but end up participating in much wider discussions. This is in part because the iterative nature of exploring how to frame a decision requires close cooperation between science advisors and decision-makers.

Scientists can improve a policy process by adding important details in the framing process, such as additional aspects or locations of harms and benefits. There is some implicit value-ladenness to any such work because the scientist much judge that something is at least a potentially valuable consequence: a scientist who does not value agriculture may not think to suggest that those consequences be considered. But similar logic to the Honest Broker of Policy Alternatives can guide us here: scientists should attempt to open up avenues for consideration, rather than narrow the scope of the problem. Additional considerations and finer grained specification of possibilities is often helpful, and when it is not helpful this is usually due to the limitations of the policymaker. So adopting an "opening up" rather than "closing down" stance should once again be the default.

6 States, consequences, and uncertainty

The most well acknowledged role of scientists in policy advising is to offer information about the potential future states of the world. In outline, this involves two forms of information which often require different expertises. The first is information about the states of the world which might come about, absent any policy information. The second is information about the states of the world which might come about given the implementation of specific policies. The first kind is involved in problem emergence and policy formulation, and acts as a point of comparison in policy evaluation. The second kind is involved in policy formulation, decision-making, implementation, and evaluation.

Certain forms of decision theory presume act-state independence, in which the state of the world is given, and the decision-maker reacts to it as best they can. But policymakers often think of the state of the world as malleable, and changing it is the target of their policymaking. Decision theories which insist on formulating decision problems so that acts and states are independent can therefore be unnatural and difficult to use when thinking about this kind of policy. At other times there is no clash, as policies are reacting to a state of the world. In climate policy, we see the two forms of policy operating in tandem. In mitigation policy, we attempt to alter the degree of climatic change in the future, by reducing emissions or sequestering greenhouse gas. In adaptation policy, we assume that some level of climate change will occur and attempt to respond to it.

These complexities inform my decision to discuss states and consequences together. The line between them can be blurry, depending on one's theoretical commitments. But there is an important difference in the scientific work required to provide the two kinds of information I mentioned at the outset of this section. Analysing a policy which aims to change the state of the world requires understanding the policy, which in turn requires significant collaboration between different scientific domains and working closely with the policymakers themselves. We can see this by considering a Covid-19 example. Policymakers who wanted to control the spread of the virus, bringing transmission to a low-enough level that the disease would be eliminated. One important proxy variable for this was the disease's reproduction number, denoted r. Policies to control the spread of the virus were described as aiming at bringing r below 1. Analysing such policies required a combination of scientific expertises. Epidemiologists and virologists were required to estimate the basic reproduction number, r_0, which describes the expected spread of the disease in a completely susceptible population, without any interventions. An intervention might then be considered, such as closing schools. Various social scientists might contribute to the discussion about the effects of this policy. Its primary aim is to reduce contact between households and thus reduce the spread of infection. How big an impact is desired? Understanding this requires assessing the expected impact of other current policies, as well as the natural, self-driven reduction in contacts between people who fear falling ill. This effect might fall within the domain of behavioural scientists or psychologists. Estimating the efficacy of school closures might involve expertise in human mobility to determine the impact this will have on the number of contacts people have. It also requires understanding the nature of the virus and the disease: how is the disease spread, and are children an important source of transmission? The expected effects of the policy on human movement and household contact then need to be combined with the understanding of the disease to estimate the effect on the reproduction number r.

6.1 Which scientists, which possibilities

A major problem in science advising during the first two years of the Covid-19 pandemic was that the advice was insufficiently integrated. Responding to a viral pandemic involves decisions whose considerations and consequences cut across numerous fields of expertise: from epidemiology and public health to economics and education. Policymakers combatting the pandemic seek to minimise the negative impact of the pandemic on people's wellbeing, which is a function of their mental and physical health, but also their income, lifespan, happiness, opportunities and capacities for self-creation, and more.

Which scientists are involved in policy advising determines in part which elements of wellbeing are considered as targets for policy, modelled in the decision-making process, and measured during policy evaluation. For example, the UK's Scientific Advisory Group for Emergencies (SAGE) was composed almost entirely of epidemiologists, medical professionals, and biomedical scientists.[4] The discussions they held and advice they provided therefore focussed largely (and indeed exclusively, to begin with) on the physical health detriments of the coronavirus disease (Bradley and Roussos 2021).

A further issue, beyond the kinds of scientists involved in policymaking, is the integration of their scientific assessments of consequences. The different components of wellbeing are interrelated, and policies targeting one component will inevitably impact others. In some cases, these impacts are positive, and a full accounting for the policy's positive impacts will be incomplete without considering these interactions. A mask mandate might reduce the burden of other infectious diseases in addition to Covid-19 and thus reduce overall harm and demand for health resources more than expected. In other cases, the impacts will be negative and so a

[4] I base this claim on an analysis of the publicly listed professions of SAGE members at all SAGE meetings during 2020 and 2021. Thanks to Izzy Gurbuz for performing the analysis.

policy may seem better than it is in the absence of an accounting of interactions. School closures will have negative effects on children's education, which in turn harms their wellbeing, possibly over a long period of time. So, what is ideally required is an integrated assessment of the impact a policy will have on wellbeing, once the dependencies and interactions are taken into account.

A common but insufficient way of accounting for a policy problem's impact on different components of wellbeing is to engage scientists from different disciplines, but then to pit their advice against one another in an adversarial environment. This risks a series of decisions in which the levels of power are handed back and forth between competing interests. Arguably this is what we saw in the UK in mid-2020, when the first Covid-19 lockdown was followed immediately by a policy called "Eat Out to Help Out", in which Britons were encouraged to eat in restaurants to help the flailing hospitality sector. This policy was the brainchild of the Treasury (the UK's finance ministry) whose concern is the health of the economy. Unsurprisingly, Eat Out to Help Out was short-lived and preceded another wave of Covid-19 infections and ultimately a series of further lockdowns.

There is a delicate balancing act required here, however. Policymakers who try to consider all effects on wellbeing risk becoming overwhelmed by the myriad side-effects of and interactions between policies. This is related to what is known in decision theory as a "grand-world" decision problem. This is a decision problem which includes all relevant consequences, all feasible options, and is specified at the finest level of detail. By contrast a "small-world" decision problem is coarser, leaving out options and details to create a more manageable problem (Savage 1954). In principle, whenever we decide we really face a grand-world decision problem; just as, in principle, every policy aims to maximise societal wellbeing across all of its constituents. There are various decision theoretic complexities here that need not detain us (see Joyce 1999, 70–77, 121–22), the point of interest is that resolving a decision into finer detail, and integrating various scientific analyses, is another area where the "more is better" attitude of the Honest Broker stance does not quite work. Policy decision-makers must, like individual decision-makers, ultimately settle on a small-world decision to confront, knowing all the while that it is an inadequate representation of the true problem they confront.

6.2 Uncertainty and uncertainty attitudes

There has been significant discussion in the philosophy of science literature about the role that scientific uncertainty plays in policymaking, and the role that value judgements play in managing scientific uncertainty. Rather than recapitulate the lessons of this literature, I want to highlight a less-discussed issue.

This concerns role of uncertainty attitudes on the part of science advisors. An uncertainty attitude is a liking of or aversion to uncertainty, which manifests in behaviour as a willingness to trade material consequences in exchange for a reduction of the uncertainty associated with making a decision. I intend to use the term "uncertainty attitude" to cover two kinds of attitude discussed in the decision science literature: risk aversion and ambiguity aversion (as well as their neutral and positive attitude alternatives). Speaking broadly, an aversion to uncertainty is a preference for decisions which involve less uncertainty over those which involve more uncertainty. For example, a risk averse person prefers a guaranteed $10 to a gamble with 50% chance of winning nothing and 50% chance of winning $20.

These attitudes encode evaluations—they reflect values rather than beliefs. It is a matter of some discussion amongst ethicists what the proper uncertainty attitude is for a policymaker who makes decisions for a population (e.g., Rowe and Voorhoeve 2018; Stefánsson 2021). So, one potential way that scientists' values could play an illegitimate role in policymaking is if their own uncertainty attitudes influenced their science advising or the eventual policy decision, rather than the morally correct uncertainty attitude for policymaking (be that the attitude of the policymaker, an attitude which represents those of the population, or an attitude which is objectively best for policymaking no matter the attitudes of the policymaker or population.) As it is presently unclear both what the ethical issues are surrounding uncertainty attitudes in policymaking, and how the attitudes of advisors may play influence policymaking, it is difficult to outline a view of what scientists ought to do (if anything) to remain neutral.

A second way that uncertainty attitudes may play a role is in motivating the scientist's engagement throughout the policy cycle. Two scientists might agree on all the facts about a potential problem, and yet one of them feel that it is a greater cause for concern than the other, because of their different attitudes to uncertainty. This in turn might motivate the concerned scientist to sound the alarm, suggest policies for addressing the problem, volunteer for work on expert panels, and so forth. As I have noted, this is all vitally important work. So, does it matter that there is going to be a selection effect in who gets involved in this work? And is it important that one factor influence this self-selection will be uncertainty attitudes?

7 Options

Scientists play a significant role in the formulation of policy options. This is especially obvious in cases where the options in question require new technologies, such as proposals for carbon dioxide removal and sequestration as a policy for mitigating climate change. Carbon dioxide removal involves "removing CO_2 from the atmosphere and durably storing it in geological, terrestrial, or ocean reservoirs, or in products" (IPCC 2018), as a form of geoengineering aimed at reducing net emissions. The formulation of a concrete policy based on carbon dioxide removal would likely involve climate scientists, other scientists, and engineers working in concert. For example, one such proposal involves capturing carbon dioxide by the carbonation of silicates. In laymen's terms: certain rocks absorb carbon dioxide, and so crushing these rocks and exposing a large surface area has the potential to absorb significant amounts of CO_2. This is clearly a proposal which required significant expertise simply to propose.

The role of scientists in policy formulation has received significant discussion in the policy literature. It is in this context that Pielke outlined his four stances for science advising and advocated the Honest Broker of Policy Alternatives. For this reason I won't have much new to add to the theoretical landscape here and will make only a few notes. As we saw in §3, Pielke argues that scientists ought to "open up" policy discussions, rather than narrowing them down by advocating for a specific policy or set of policies. It is only in rare cases that policymakers can truly "follow the science": in particular, when there are "shared values and low uncertainties about the relationship of alternative courses of action and those valued outcomes" (Pielke 2007, 36), in which case the role of science is simply to identity the policy which best achieves the agreed-upon goals. This situation, Pielke notes, is exceedingly rare.

However, Pielke's "opening up" strategy may be unfeasible in practice. There are a great many potential policy options which may apply to any given problem. True neutrality is thus practically impossible: scientists will inevitably make some decisions about which alternatives are more feasible or attractive (Havstad and Brown 2017). Feasibility judgements may encode evaluative assumptions, such as the judgement that the government will not spend the required money to develop a nascent technology or that the costs are prima facie too high.

Value judgements in option formulation can be unintended and unnoticed. For example, consider early discussions of policy responses to the Covid-19 pandemic in meetings of the UK's SAGE committee in February and March 2020.[5] In discussing international travel restrictions and domestic movement restrictions as potential policy responses to the spreading virus, the scientific advisory committee declared these policy options to be "draconian"—a clearly value-laden term, which motivated the exclusion of those options from early consideration. This judgement was bound up with a prediction: that UK residents would not tolerate such restrictions and that the policy would therefore be ineffective. This is an empirical hypothesis that could be supported by evidence, though it is unclear that the medical and epidemiological expertise of the committee members put them in a good position to have access to this evidence. But whether the policy is acceptable in a democracy, or respects liberty, or accords with national values, are clear not scientific or value-neutral questions.

This is an example of "closing down" a policy discussion rather than "opening it up". But I think it is likely that the scientists involved did not actively make the judgement, by reflecting on a set of values and judging consciously that the policy options were unacceptable. This was neither open Issue Advocacy nor stealth Issue Advocacy. Instead, it was subconscious Issue Advocacy. I suspect that the scientists looked at the government of China's response to the initial outbreak in Wuhan and had an emotional response, *feeling* that it was unacceptable and not how things ought to be done. The entanglement between the judgement that the policies are draconian and that they won't work in the UK shows the difficulty in identifying value judgements. The scientists may well have expressed the judgement primarily as being about feasibility, not permissibility. Options which clash with a scientist's subconscious moral commitments may simply not appear to them as feasible alternatives.

This is a particularly difficult form of value-ladenness for scientists to avoid. But the case of the UK's response to Covid-19 shows why it is important to combat: SAGE advisory meetings on Covid-19 began in late January, and in early February policies such as international travel

[5] In particular, the meetings called SAGE 3 and SAGE 7. Minutes, which record the term "draconian", are available here: https://www.gov.uk/government/publications/sage-minutes-coronavirus-covid-19-response-3-february-2020 and https://www.gov.uk/government/publications/sage-minutes-coronavirus-covid-19-response-13-february-2020

restrictions were discussed (and rejected). It was not until mid-March that the UK enacted Covid-19 policies, beginning with the closure of schools and soon moving to a “lockdown”—a set of policies which restricted movement outside of the home, closed shops and businesses, and extended the closure of schools and universities. Lockdowns of this sort involved the curtailment of liberties and exercise of state power in a way that may well have been unthinkable in the months before the pandemic. Enacting such policies required an evaluation of the costs and benefits, made by a fully informed and democratically authorised decision maker. Scientists lacked both the authority and the requisite information to make such a judgement and should not have removed options from the agenda because they pre-judged them to be unacceptable.

8 Conclusion

This paper has aimed to map out a landscape of different roles that scientists play in policymaking, and to identify different sites where value judgements might be made—knowingly or unknowingly—by those scientists. The philosophy of science literature has been overly focused on one activity—research—and one way that values can play a role—inductive risk judgements. I have tried to highlight areas where values may enter into science advising that have little or nothing to do with inductive risk, in part because they have little or nothing to do with conducting research. My conviction is that philosophers should pay attention to all the roles that scientists play in policymaking and extend our thinking about the desirability and nature of value-freedom to encompass this wider set of activities.

A secondary aim has been to provide this map in terms of the language of decision theory. This is because of a tendency that I have detected for philosophers to associate the formal tools of decision theory with a limited view of the role scientists play in policy and value-free view of science itself. Such an association, if it exists, is entirely contingent: decision theory is a flexible tool that can be used to facilitate discussions of many roles for scientists and of value-ladenness in science at many points in the decision process. I do not claim that decision theory is a blank slate—any decision theory involves idealising assumptions including some which are explitictly normative (cf. Roussos 2022)—but even within the relatively orthodox theory I have used here, we have found tools to discuss a range of roles for scientists in policy support.

Acknowledgements

My climate change-related research takes place within the project "Climate Ethics and Future Generations" at the Institute for Futures Studies. Financial support from Riksbankens Jubileumsfond (grant M17-0372:1) is gratefully acknowledged. The Covid-19 aspects of this research were supported in part by a grant from the LSE Research Support Fund (PI: Richard Bradley). Thanks to our research assistant, Izzy Gurbuz, who reviewed and annotated all the minutes from SAGE meetings on Covid-19 and whose work has been very valuable in formulating these thoughts.

References

Bradley, Richard. 2014. 'Decision Theory : A Formal Philosophical Introduction', 1–42.

Bradley, Richard, and Joe Roussos. 2021. 'Following the Science: Pandemic Policy Making and Reasonable Worst-Case Scenarios.' *LSE Public Policy Review* 1 (4:6): 1–7. https://doi.org/10.31389/lseppr.23.

Bright, Liam Kofi. 2018. 'Du Bois' Democratic Defence of the Value Free Ideal'. *Synthese* 195: 2227–45.

Brown, Matthew J. 2019. *Science and Moral Imagination*. Pittsburgh: University of Pittsburgh Press.

Douglas, Heather. 2000. 'Inductive Risk and Values in Science'. *Philosophy of Science* 67 (4): 559–79. https://doi.org/10.1086/392855.

———. 2009. *Science, Policy and the Value-Free Ideal*. Pittsburgh: University of Pittsburgh Press.

Du Bois, W. E. B. 1898. 'The Study of Negro Problems'. *The Annals of the American Academy of Political and Social Science* 11 (1): 1–23.

Harding, S. 1995. '"Strong Objectivity": A Response to the New Objectivity Question'. *Synthese* 104 (3): 331–49.

Havstad, J, and Matthew J. Brown. 2017. 'Inductive Risk, Deferred Decisions, and Climate Science Advising'. In *Exploring Inductive Risk: Case Studies of Values in Science*, edited by K Elliott and T Richards, 1–21. New York, NY: Oxford University Press.

Hempel, Carl G. 1965. 'Science and Human Values'. In *Aspects of Scientific Explanation*, 81–96. New York: The Free Press.

IPCC. 2018. 'Annex I: Glossary'. In *Global Warming of 1.5°C*, edited by J.B.R. Matthews. https://www.ipcc.ch/sr15/chapter/glossary/.

Jasanoff, Sheila. 1990. *The Fifth Branch*. Cambridge, MA: Harvard University Press.

Jebeile, Julie, and Joe Roussos. 2023. 'Usability of Climate Information: Towards a New Scientific Framework'. *Wiley Interdisciplinary Reviews. Climate Change*.

Joyce, James M. 1999. *The Foundations of Causal Decision Theory*. Online (20. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511498497.

Kirchhoff, Christine J, Maria Carmen Lemos, and Suraje Dessai. 2013. 'Actionable Knowledge for Environmental Decision Making: Broadening the Usability of Climate Science'. *Annual Review of Environment and Resources* 38: 393–414.

Kitcher, Philip. 1990. 'The Division of Cognitive Labor'. *Journal of Philosophy* 87 (1): 5–22. https://doi.org/10.2307/3655750.

———. 2001. *Science, Truth, and Democracy*. Oxford: Oxford University Press.
———. 2011. *Science in a Democratic Society*. New York: Prometheus.
Krohn, Wolfgang. 2005. 'Enquete Commissions'. In *Encyclopedia of Science, Technology, and Ethics*, edited by Carl Mitcham, 641–44. Farmington Hills, MI: Macmillan Reference.
Lach, Denise, Peter List, Brent Steel, and Bruce Shindler. 2003. 'Advocacy and Credibility of Ecological Scientists in Resource Decisionmaking: A Regional Study'. *BioScience* 53 (2): 170–78. https://doi.org/10.1641/0006-3568(2003)053[0170:AACOES]2.0.CO;2.
Lackey, Robert T. 2004. 'Normative Science'. *Fisheries* 29 (7): 38–39.
———. 2007. 'Science, Scientists, and Policy Advocacy'. *Conservation Biology* 21 (1): 12–17.
Longino, Helen E. 1996. 'Cognitive and Non-Cognitive Values in Science: Rethinking the Dichotomy'. In *Feminism, Science, and the Philosophy of Science*, edited by L. H. Nelson and J Nelson, 39–58. New York: Kluwer.
———. 2004. 'How Values Can Be Good for Science'. In *Science, Values, and Objectivity*, edited by Peter Machamer and Gereon Wolters, 127–42. Pittsburgh: University of Pittsburgh Press.
Moore, Alfred. 2017. *Critical Elitism: Deliberation, Democracy, and the Problem of Expertise*. Theories of Institutional Design. Cambridge: Cambridge University Press. https://doi.org/10.1017/9781108159906.
Parker, Wendy S, and Greg Lusk. 2019. 'INCORPORATING USER VALUES INTO CLIMATE SERVICES'. *American Meteorological Society* 100 (9): 1643–51.
Pielke, Roger A., Jr. 2007. *The Honest Broker*. Cambridge: Cambridge University Press.
Reiss, Julian, and Jan Sprenger. 2017. 'Scientific Objectivity'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2017. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/win2017/entries/scientific-objectivity/.
Roussos, Joe. 2022. 'Normative Formal Epistemology as Modelling'. *The British Journal for the Philosophy of Science* forthcoming.
Rowe, Thomas, and Alex Voorhoeve. 2018. 'Egalitarianism under Severe Uncertainty'. *Philosophy and Public Affairs* 46 (3): 239–68.
Rudner, Richard. 1953. 'The Scientist Qua Scientist Makes Value Judgments'. *Philosophy of Science* 20 (1): 1–6.
Salter, Liora. 1988. *Mandated Science*. Dordrecht: Kluwer.
Savage, L J. 1954. *The Foundations of Statistics*. 1972nd ed. Dover.
Steele, Katie. 2012. 'The Scientist qua Policy Advisor Makes Value Judgments'. *Philosophy of Science* 79 (5): 893–904. https://doi.org/10.1086/667842.
Stefánsson, H. Orri. 2021. 'Ambiguity Aversion behind the Veil of Ignorance'. *Synthese* 198 (7): 6159–82. https://doi.org/10.1007/s11229-019-02455-8.
Weber, Max. 1949a. '"Objectivity" in Social Science and Social Policy'. In *The Methodology of the Social Sciences*, translated by Edward A. Shils and Henry A. Finch, 50–112. New York: The Free Press.
———. 1949b. 'The Meaning of "Ethical Neutrality" in Sociology and Economics'. In *The Methodology of the Social Sciences*, translated by Edward A. Shils and Henry A. Finch, 1–49. New York: The Free Press.
Winsberg, Eric. 2018. *Philosophy and Climate Science*. Cambridge: Cambridge University Press.