# The Scientist qua Policy Advisor Makes Value Judgments

Katie Steele*†

Richard Rudner famously argues that the communication of scientific advice to policy makers involves ethical value judgments. His argument has, however, been rightly criticized. This article revives Rudner's conclusion, by strengthening both his lines of argument: we generalize his initial assumption regarding the form in which scientists must communicate their results and complete his 'backup' argument by appealing to the difference between private and public decisions. Our conclusion that science advisors must, for deep-seated pragmatic reasons, make value judgments is further bolstered by reflections on how the scientific contribution to policy is far less straightforward than the Rudner-style model suggests.

**1. Introduction.** The role of science advisors in policy making has expanded and become formally recognized, particularly in the area of environmental regulation. A significant worry is that scientists' specialist knowledge permits or rather forces them to infringe on the role of those democratically elected to decide what is good for society. This article argues that the worry is justified; for deep-seated pragmatic reasons, scientists must make value judgments when presenting their opinions on factual matters. This has implications for how scientists should be engaged in policy decisions.

At first blush the worry might seem justified only at the practical, and not the theoretical, level. While facts and values are closely entangled, in theory, at least, they can surely be separated out: scientists provide the data, and policy makers evaluate consequences and make final decisions accordingly. Richard Rudner, however, famously argued (1953) that the fact and value

*To contact the author, please write to: Department of Philosophy, Logic, and Scientific Method, London School of Economics and Political Science, Houghton Street, London WC2A 2AE; e-mail: K.Steele@lse.ac.uk.

†Many thanks to Jan Sprenger, Jack Justus, Helen Regan, Hasok Chang, Rosa Runhardt, and Kevin Elliott for valuable feedback and comments on earlier versions of this article.

Philosophy of Science, 79 (December 2012) pp. 893–904. 0031-8248/2012/7905-0004$10.00
Copyright 2012 by the Philosophy of Science Association. All rights reserved.



https://doi.org/10.1086/667842 Published online by Cambridge University Press

dimensions cannot be separated so easily and that scientists do make value judgments when presenting their results. Rudner's claims about values in science are in fact directed at the heart of scientific practice; Rudner argues that scientists qua scientists must make value judgments.[1] His points are arguably more compelling, however, when applied to the science-for-policy domain. In short, Rudner appeals to the fact that scientists never know things for sure but rather assess the weight of evidence at a particular point in time to determine which hypothesis to accept as true. The consequences of getting it wrong must drive this assessment. For instance, if the consequences of being alarmist about climate change were considered much worse than the consequences of false security, then much evidence would be required to support, say, an announcement that only 450 parts per million by volume greenhouse gases will lead to a 2°C rise in average global temperature. If, however, false security were considered the really bad outcome, then less evidence would suffice to support the same assertion.

One might immediately see problems with Rudner's argument. We need not look further than the previous example: climate scientists do not typically report on possible temperature-rise scenarios in the manner above—they rather give a probability distribution, or probability bounds, over the possible temperature increases (see, e.g., IPCC 2007). Indeed, section 2 of this article formalizes Rudner's argument in order to bolster the commonsense view that, as it stands, the argument is too crude. Moreover, Rudner's backup argument, while having more intuitive appeal, is not plausible as it stands. Sections 3 and 4 revisit the two strands of Rudner's argument and improve on them. The upshot is that we do in fact require scientists to make value judgments, for more and less deep-seated pragmatic reasons. Objections to this position are raised and answered in section 5, and section 6 concludes.

**2. Rudner's Argument Clarified.** While Rudner's (1953) paper concerns core practice in science, it can be construed as arguing that scientists make ethical or social-political value judgments when delivering advice to policy makers because they must decide how much evidence is sufficient to warrant stating, at some point in time, that a genetically modified (GM) crop is safe to cultivate, say, or that we should expect at least a 2°C rise in average global temperature. That is, scientists do not just report the facts; they rather make decisions, taking into account potential consequences, about what should be regarded as the facts at a particular time, given the available evidence. And where such decisions are being made, values must be playing a role.

1. Note that Jeffrey (1956) and Levi (1960) both respond to Rudner's argument that scientists qua scientists make value judgments.

https://doi.org/10.1086/667842 Published online by Cambridge University Press

TABLE 1. The Policy Maker's Decision

| | $S$ | $\neg S$ |
|---|---|---|
| $A_1$ | $O_1$ | $O_2$ |
| $A_2$ | $O_3$ | $O_4$ |

It is important to see clearly what are the ingredients of this claim. A number of assumptions are involved, some of which are idealizations that strengthen the conclusion, because it is more surprising in the ideal setting that scientists must make value judgments. At least one assumption, however, is problematic; we will get to it shortly. To begin with, let us cast Rudner's argument in terms of a policy-decision setting, as per table 1. To keep things simple, there are only two possible policy acts, $A_1$ and $A_2$, and two relevant states of the world that may arise, affecting the consequences of these acts. For instance, the acts may be 'permit planting GM crop' ($A_1$) and 'ban planting of GM crop' ($A_2$), and the states of the world that affect the consequences may be 'crop becomes invasive' ($S$) and 'crop does not become invasive' ($\neg S$). The scientist must deliver information about the states $S$ and $\neg S$, and the policy makers subsequently evaluate outcomes and decide on the acts.

Let us first consider the idealizing assumptions, over and above the simple structure of the decision problem. One is that the scientists giving advice to policy makers share the same opinions or else that there is just one scientist involved. In the real world, scientists disagree, and this considerably complicates the decision as to what facts to report to the policy maker (see Hansson 1993). Our question is whether a single scientist necessarily makes value judgments when communicating opinions. Note also that the single scientist is assumed to be an expert on the entire question at issue: whether $S$ or $\neg S$ is true. Another important idealization is that the scientist in question wants to respect the democratic process as far as possible and does not try to act strategically; he or she presents results as truthfully as possible. The Rudner-style argument is that the scientist must inevitably make value judgments that are properly in the domain of policy makers; the argument is not that it is possible for a scientist so inclined to have an undue influence on the outcome of a policy decision.

Now for the more worrisome assumptions. Rudner assumes that scientists must deliver a binary assessment of the relevant hypotheses—they either accept or reject hypotheses. In our example, this is to say that the scientist is required to state whether $S$ or $\neg S$ should be accepted as true. If $S$, then $A_2$ would be chosen (i.e., the ban) because presumably planting a GM crop that then becomes a weed is worse than not planting the GM crop at all ($O_1 < O_3$). If, however, $\neg S$ is reported true, then the policy makers would choose not to impose a ban because presumably $O_2 > O_4$.

https://doi.org/10.1086/667842 Published online by Cambridge University Press

TABLE 2. The Scientist's Decision

| | $S$ | $\neg S$ |
|---|---|---|
| $R_S$ | $O_a$ | $O_b$ |
| $R_{\neg S}$ | $O_c$ | $O_d$ |

A further assumption is that the scientist is unaware of precisely how the policy makers value the decision outcomes. The scientist nonetheless faces a decision of his or her own that has related outcomes—whether to report that $S$ is true (call this act $R_S$) or whether to report that $\neg S$ is true (act $R_{\neg S}$; see table 2). Let us assume the standard theory of decision: expected utility theory. Then the value of each act, $R_S$ and $R_{\neg S}$, is just the sum of the probability of each state multiplied by the value of the consequence the act yields in that state.[2] The scientist should make the report judged to have greater expected value/utility.

Note that the scientist-advisor may not actually go through the cognitive process of deciding, along the lines of table 2, which state to report as true at a particular time, but that is so much the worse for his or her advice being rationally justified. The claim is just that the scientist's report amounts, explicitly or implicitly, to a decision that involves value judgments. In table 2, the outcome $O_b$, for instance, is the consequence of reporting that $S$ is true (e.g., the GM plant will become invasive) when in fact it is not (rather, $\neg S$ is true). This has some relationship to the outcome $O_4$ in table 1, but we are assuming that the scientist does not know the value the policy makers attribute to $O_4$ relative to the other policy outcomes, or even the precise form of the policy makers' decision, and thus the scientist must simply estimate and evaluate the outcomes of the decision problem—$O_a$ to $O_d$—as best he or she can. This ignorance assumption is crucial to the argument, but it seems reasonable given the separation of the scientist and the policy maker, in spatial or temporal or institutional terms. Nevertheless, we will return to this point later.

The current concern is the first worrisome assumption regarding how scientists must communicate their results: true/false judgments about states. Anticipating criticism on this point, Rudner allows that the scientist may rather offer a probability function over the state space, but he claims that this just 'pushes the problem one step back' (e.g., Douglas 2009). Instead of having to decide what state to report as true at a particular time, the scientist must now decide what probability function over the states to report.

The push-the-problem-one-step-back argument is misguided or at least incomplete, however, because the probability function over decision states

2. So, e.g., the value of $R_S$, $U(R_S)$, is just $U(R_S) = \Pr(S) \times U(O_a) + \Pr(\neg S) \times U(O_b)$.

https://doi.org/10.1086/667842 Published online by Cambridge University Press

is typically understood as representing the scientist's beliefs.[3] One might object that the beliefs of scientists rarely live up to the (Bayesian) ideal of a precise probability function, but that amounts to a mere amendment to our objection to the push-the-problem-one-step-back argument. The core point is that scientists can simply report their beliefs to policy makers, using whatever representation that best 'captures' these beliefs, whether this be a probability function, a set of probability functions, nonadditive probabilities, or something else.[4]

Let us elaborate on the above. Scientists do not choose their beliefs in a manner that takes into account real-world consequences. They deliberate, to be sure, in order to arrive at well-reflected beliefs, but they should not subsequently revise these beliefs according to the type of world the beliefs imply. The criteria informing rational belief are arguably purely concerned with the truth. Various authors (e.g., Longino 1990; Douglas 2009) dispute this latter point, arguing that social-political values inform theory choice or inform a scientist's belief that a hypothesis is true. They claim that since evidence underdetermines hypotheses, when scientists assess the relative plausibility of these hypotheses, they must appeal to other nonempirical criteria, such as simplicity, explanatory power, and so on. These criteria are not obviously related to truth, and moreover, scientists' understanding of at least some of these criteria is affected by their social-political outlook. That may be so; there is no need to dispute the matter here. It suffices to say that there is a large difference between social-political background subtly affecting scientists' judgments of simplicity that in turn affect their beliefs about hypotheses and the type of value judgments that Rudner is concerned with that explicitly take account of real-world consequences. It is not helpful to mix these two issues together. Rudner's second line of argument fails—if scientists can simply communicate their beliefs, then their advice to policy makers is not a choice that necessitates weighing up various real-world consequences.

**3. Generalizing Rudner's First Argument.** This section and the next offer revised and improved versions of both Rudner's lines of argument and, thus, serve to bolster his conclusion. We begin with the first argument, which apparently rests on the assumption that scientists must convert their beliefs from probabilistic to binary form when they report to the policy maker. While this seems an unreasonable restriction on the science-policy commu-

3. Jeffrey (1956) and McMullin (1982), among others, criticize Rudner along these lines. They too argue that scientists can simply report probabilities for hypotheses and avoid value judgments.

4. Note that later in the article we may assume, at various points, that a scientist's beliefs can be represented by a precise probability function, but that is just for simplicity.

https://doi.org/10.1086/667842 Published online by Cambridge University Press

TABLE 3. Intergovernmental Panel on Climate Change Scientist's Decision

| | $S$ | $\neg S$ |
|---|---|---|
| Report high confidence | $O_a$ | $O_b$ |
| Report medium confidence | $O_c$ | $O_d$ |

nication process, it can be replaced by a more general assumption. Indeed, perhaps Rudner only intended the binary-form conversion to be illustrative of a more general phenomenon. The key point is simply that scientists must often convert their beliefs from probabilistic form to some other form. One example is the conversion from a probability function over states to a function that assigns 1 (true) to one state and 0 (false) to all others. This is Rudner's focus. But there are other examples of the general phenomenon: cases in which scientists must report their findings using some standard measure that is cruder than the measure required to accurately represent their beliefs.

For example, in the IPCC (Intergovernmental Panel on Climate Change) AR4 Synthesis Report (2007), scientists are asked to express their expert judgments on certain issues in terms of one of the following categories: very high confidence, high confidence, medium confidence, low confidence, and very low confidence. Guidelines are given regarding what quantitative levels these qualitative categories correspond to, but the levels are approximates only. For instance, high confidence is said to correspond to a chance of "about 8 out of 10," whereas medium confidence corresponds to a chance of "about 5 out of 10." Clearly there is room for discretion in determining the cutoff between these two categories. Consider a statement in the report (topic 3.3): "There is *medium confidence* that crop productivity will increase slightly at mid to high latitudes for local mean temperature increases of up to 1 to 3°C depending on the crop, and then decrease beyond that in some regions" (48). The 'scientists' have presumably translated their own probabilistic beliefs to the standardized judgment of confidence. For instance, their subjective probability may be 0.6, and they chose to report medium confidence.

Table 3 illustrates the decision that the IPCC scientist makes. For our example case, $S$ is 'crop productivity changes'. It may well be very difficult for scientists to determine what the outcomes $O_a$ to $O_d$ actually amount to. 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. (In this case, it is likely that the consequences of each report are similar, and thus the expected value of the reports is roughly equal. That may not always be the case, however.) The mere fact that the outcomes are difficult to specify and evaluate does not deny that a decision about what to report must be made. Scientists commit to an evalu-

https://doi.org/10.1086/667842 Published online by Cambridge University Press

ation of the desirability of outcomes, even if they make the reporting decision quickly and carelessly and the judgments are only implicit in their decision.

One might object that the IPCC confidence scale is crude because scientists are unlikely to have precise probabilistic beliefs on the relevant climate change matters. The translation problem remains, however, even if we relax the admittedly unrealistic assumption of precise probabilistic beliefs. For instance, perhaps the scientist's belief regarding crop productivity increase is best expressed by the interval [0.55, 0.65] (à la Levi 1985). Here too, the scientists must convert their belief to the high confidence, medium confidence, and so on, standard scale, and this involves, perhaps implicitly, evaluating consequences.

There will be many cases in which fact finders are asked to report their beliefs in a standardized manner so as to simplify/streamline the job of the public decision maker. Another example is the juror's role in reporting whether the accused is guilty according to the scale of preponderance of evidence, beyond reasonable doubt, and so on. The juror's decision about what to report in a particular case is analogous to those decisions modeled in tables 2 and 3. In short, while Rudner's assumption that scientists must report certainties to policy makers is not convincing, there is a more plausible general assumption that delivers Rudner's conclusions: scientists must often, for pragmatic reasons, deliver advice to policy makers in terms of a standardized evidence/plausibility scale that is cruder than the scale appropriate for representing their beliefs. In other words, the structure of the beliefs of scientists is often more complex than the predefined evidence scales, and there is no canonical projection. In these situations, scientists cannot avoid making value judgments, at least implicitly, when deciding how to match their beliefs to the required scale.

**4. Completing Rudner's Second Argument.** We said earlier that Rudner's push-the-problem-one-step-back argument runs counter to the standard interpretation of probability functions in decision models. There is nonetheless something appealing about this line of argument. This section completes Rudner's second argument by arguing that, in the context of public decision making as opposed to private decision making, there is not the tight connection between belief and choice, and thus the ideal measure over the state space in a public decision model may not be the scientist's own beliefs.

The right way to conceive of public decision models is a difficult open question. Here we merely suggest one possibility for the appropriate measure on the state space: not the scientist's exact belief function but rather some kind of confidence interval—or in Bayesian terms, credible interval—that is constructed from the scientist's belief function. This may be

https://doi.org/10.1086/667842 Published online by Cambridge University Press

an interval for the true or objective probability function over the state space. It may otherwise be an interval for a continuous state variable, like predicted temperature change. The latter in fact commonly features in science reporting, as the following statement in the IPCC AR4 Synthesis Report indicates: "Temperature change at 2090–2099 relative to 1980–1999, under the most optimistic emissions scenario considered, is *likely* to be in the range 1.1 to 2.9°C with a best estimate of 1.8°C" (IPCC 2007, table 3.1).

This statement includes a point estimate for the quantitative variable of interest, as well as a likely interval. The IPCC report advises that "likely" amounts to a probability greater than 0.66.[5] This is a fixed threshold, but it might just as well have been left open to the scientist's discretion. If so, the problem of working out the appropriate threshold for the credible interval is indeed a case of pushing the problem one step back.

In general, if public decision models require credible intervals for state variables/objective probabilities, the scientist may need to choose the appropriate size of the interval: 70%? 90%? 95%? Scientists must consider whether the pragmatic context requires a cautionary approach, which typically amounts to a wider credible interval, or whether a narrow credible interval will suffice. This choice involves value judgments—implicit in the decision is an assessment of the pragmatic consequences of error.

**5. Objections: Reconceiving the Scientist's Role.** We will not attempt to settle the issue of how public decision models should be constructed. Our general argument (unifying the claims in secs. 3 and 4) is that when communicating their findings, scientists are often asked to make the following translation:

subjective beliefs about states → some other measure on the state space.

In many cases, such a translation involves a decision on the part of scientists: given their beliefs about the world and their evaluation of the consequences that may arise, scientists must decide how to report, according to the specified scale, their assessment of the hypotheses.

We consider two objections to this conclusion. The first questions the role assigned to the scientist and essentially argues that if we were to demand more from policy makers, we could avoid the conclusion that scientists must make value judgments. In short, this is a fair but not terribly interesting objection. The second objection runs deeper because it permits us the assumption that scientists, in their role as policy advisors, must often convert their beliefs about state hypotheses to some purportedly more usable form. The

5. Section 3 discussed a different uncertainty scale used in the IPCC report, one expressed in terms of degrees of confidence. Arguably, it is not necessary to appeal to more than one scale of uncertainty in the report, but let us not dwell on that here.

https://doi.org/10.1086/667842 Published online by Cambridge University Press

argument, however, is that even when asked to play this role, scientists should not engage in decision-theoretic reasoning.

Let us start with the first objection. One might argue that the scientist can simply provide a sufficiently detailed probabilistic belief function on the state hypotheses, and the policy maker can then convert the function to a form suitable for the decision problem. Indeed, one might infer from our analysis that this is what should happen because scientists are not responsible for the evaluative dimension of a public decision.

Our conclusion can indeed be avoided in this manner. We are departing from an idealized setting in stipulating that it is the role of scientists as policy advisors to convert their probabilistic beliefs to whatever is the required measure on the state space. There are a couple of different pragmatic reasons, however, for making this assumption. To begin with, the cruder measure on the state space may simply be easier for the scientist to express or else easier for the policy makers to interpret. When it comes to credible intervals, for instance, the scientist may better appreciate the significance and proper interpretation of the interval and so be in a better position to determine its appropriate size. Otherwise, quite apart from these issues of human capabilities, the scientist may be required, for cost-effectiveness reasons, to do some of the evaluative work when contributing to a public decision. It may be too costly to have policy makers participate in every decision of a certain kind, so they rather specify broad values that apply to a whole class of cases. The burden would then rest with the scientist to deliver the results (e.g., determine the threshold for a beyond-reasonable-doubt judgment or the exact size of a credible interval) in a way that attends to the finer details of the particular case at hand.

One might insist that our appeal to pragmatics is misplaced here or that pragmatics in fact dictates the opposite division of labor: it is easier for policy makers to translate scientists' beliefs to the form suitable for policy making. Either way, it is important to recognize that someone must traverse the fact-value divide. Either the scientists must venture into social-political value territory, or at least adopt the policy makers' values in all their subtlety, or the policy makers have to reckon with probabilities and the interpretation of credible intervals.

Let us move on to the second objection, which is potentially more damaging to our argument. The objection is that, even if scientists, for pragmatic reasons, must translate their beliefs to some other required form, this should not be conceived as a decision, of the kind described in table 3. The antagonist here owes us an account, however, of exactly how, and according to what principles, the translation task should be done. It is no use appealing to epistemic principles. Assessments of evidential weight, simplicity, and so on, inform the scientists' beliefs—or, in other words, their all-things-considered assessment of the relevant hypotheses—but we cannot appeal to

https://doi.org/10.1086/667842 Published online by Cambridge University Press

epistemic principles a second time over to inform how these all-things-considered assessments should be converted to a cruder measure. As emphasized throughout the article, this final step is not an epistemic but rather a pragmatic matter. As such, surely it should be governed by decision theory. But that is precisely the objection. The only remaining option seems to be a translation rule that lacks any normative justification whatsoever. But is that really a plausible alternative?

The best way to defend the antagonist's proposal is by appeal to convention. In this way, a procedure may lack normative justification and yet not be ad hoc. It may be better in the long run for scientists to simply follow a convention when translating their beliefs to some other measure, rather than try to tailor their reports to the idiosyncrasies of the particular situation at hand. For example, classical statistical tests are often conducted with blind adherence to conventional significance levels, typically 0.01 or 0.05. This has been the target of much criticism, even within classical statistics, but arguably, communication is better and inferences more reliable when all scientists, at least within a subfield, report on statistical hypotheses according to the same level of stringency.

The convention proposal rests on a rule-based as opposed to an act-based justification for scientific reporting. (Think rule as opposed to act consequentialism.) The story would go like this: decision-theoretic reasoning would be employed, not in any particular instance but rather to settle reporting conventions for broadly characterized cases. The use of a rule would itself have virtues—public accessibility and reliability—and this would plausibly outweigh any advantages associated with scientists refining their reporting choices to the specifics of any particular scenario.

The above is a compelling objection to our argument that the scientist, qua policy advisor, must make value judgments. Note, however, that it appeals to a more complex account of the situation. Our claims concerned a scientist's reporting task at a snapshot in time. Admittedly, however, it is more reasonable to conceive of the scientist's predicament in dynamic or temporally extended terms: the scientist is making one reporting choice among many. In this more complex setting another sort of justification becomes available—we can appeal to reporting conventions that are advantageous in the long run. Perhaps the scientist need not engage in social-political value judgments after all, if the reporting conventions for broad classes of cases can be settled by policy makers. Let us not rush to this revised conclusion, however, as consideration of the dynamics of science reporting raises further issues.

**6. Final Remarks on the Dynamics of Science Advising.** If we were to adopt a more dynamic view of the science advising process, there would be other sources of value judgment to contend with. To begin with, the Rudner-style model assumes a fixed decision problem in which the sole input of the

https://doi.org/10.1086/667842 Published online by Cambridge University Press

scientist concerns the truth of the states. In reality, however, scientists would be involved in the initial setup and refinement of a policy decision problem. They would play a large role in initially identifying and characterizing the possible outcomes of any proposed policy act, before evaluating the likelihood of these outcomes/states. As such, the scientist needs to anticipate what sort of outcomes matter to policy makers. In many cases, this may well be obvious: for example, if the difference between a 2°C and a 4°C rise in average global temperature amounts to a large difference in the frequency and intensity of severe weather events, then clearly policy makers would want to take stock of these varied effects. In other cases, however, scientists may overlook or not communicate details they deem unimportant, when these details do in fact matter to various community interest groups.

The other prominent issue is the value of collecting more evidence. There is always the option of deferring a policy decision until more evidence can be brought to bear on the potential act outcomes. This may be evidence 'from the field' or else more analytic work or tweaking of a predictive model. Whether to defer a decision should ultimately be the policy maker's call: deferring may have costs that need to be weighed against the benefits of extra evidence. (This is the standard value-of-evidence decision problem, the basics of which are outlined, e.g., in Resnik [1987].) The issue here is that policy makers, in order to make the deferral call, require very subtle input from the scientist regarding the evidence that may arise from further field work/experimentation/analysis and how this evidence will bear on the specification of the policy decision problem. In practice, scientists will simply make the judgment themselves about whether more evidence would be useful. In effect, this is another stronger, pragmatic argument for why scientists must make value judgments. It is one thing for policy makers to be informed of the subtleties of the scientist's beliefs regarding states; it is another thing for policy makers to also be informed of how the scientist's beliefs may change given the various types of evidence that may be collected in the future.

In short, while appeals to the dynamics of science reporting may, to some extent, challenge the conclusion that the scientist as policy advisor must make social-political value judgments, it does not really help our antagonist's cause in the end. Once we consider the scientist's role in initially specifying a policy decision, as well as when to actually make the decision versus collect more evidence, we see how entrenched the scientist is in social-political value territory. This is a problem for democratic procedures. It can be mitigated, although not solved, only via sustained communication between policy makers and the scientists who advise them.

## REFERENCES

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

https://doi.org/10.1086/667842 Published online by Cambridge University Press

Hansson, Sven Ove. 1993. "Entscheidungsfindung bei Uneinigkeit der Experten" [Decision-making when experts disagree]. In *Die Modernisierung der Demokratie, Internationale Ansätze*, ed. Horst Zillessen, Peter C. Dienel, and Wendelin Strubelt, 87–96. Opladen: Westdeutscher.

IPCC (Intergovernmental Panel on Climate Change). 2007. *Climate Change 2007: Synthesis Report; Contribution of Working Groups I, II and III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Ed. Core Writing Team, R. K. Pachauri, and A. Reisinger. Geneva: IPCC.

Jeffrey, Richard C. 1956. "Valuation and Acceptance of Scientific Hypotheses." *Philosophy of Science* 23 (3): 237–46.

Levi, Isaac. 1960. "Must the Scientist Make Value Judgments?" *Journal of Philosophy* 57 (11): 345–57.

———. 1985. "Imprecision and Indeterminacy in Probability Judgment." *Philosophy of Science* 52 (3): 390–409.

Longino, Helen E. 1990. *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry.* Princeton, NJ: Princeton University Press.

McMullin, Ernan. 1982. "Values in Science." In *PSA 1982: Proceedings of the 1982 Biennial Meeting of the Philosophy of Science Association*, vol. 2, ed. Peter D. Asquith and Thomas Nickles, 3–28. East Lansing, MI: Philosophy of Science Association.

Resnik, Michael. 1987. *Choices: An Introduction to Decision Theory.* Minneapolis: University of Minnesota Press.

Rudner, Richard. 1953. "The Scientist qua Scientist Makes Value Judgments." *Philosophy of Science* 20 (1): 1–6.

https://doi.org/10.1086/667842 Published online by Cambridge University Press