Uncertainties, Values, and Climate Targets

Mathias Frisch*

Using climate policy debates as a case study, I argue that a certain response to the argument from inductive risk, the hedging defense, runs afoul of a reasonable ethical principle: the no-passing-the-buck principle according to which scientists ought to offer their scientific judgment when such judgment is both possible and needed in public policy debates.

1. Introduction. Can nonepistemic values play a legitimate role in science? The influential argument from inductive risk (Rudner 1953; Douglas 2000) gives an affirmative answer to this question. According to this argument, decisions whether to accept or reject a scientific hypothesis under conditions of uncertainty are influenced by nonepistemic values, since such decisions require value-laden judgments on acceptable standards of evidence. At least in the case of scientific hypotheses with implications for public policy, judgments on acceptable standards of evidence will reflect nonepistemic values.

Defenders of the value-free ideal have replied that scientists ought not to accept or reject hypotheses under conditions of uncertainty but instead should report to what extent the evidence supports the hypotheses under consideration (Jeffrey 1956; Betz 2013; Parker 2014). I want to call this reply the *hedging defense*. According to the hedging defense, it is always possible for scientists to hedge their commitments to hypotheses in ways that do not require value-laden judgments.

In this article I argue that the hedging defense fails as a general reply to the argument from inductive risk, since it is incompatible with what I take to be a reasonable demand on scientists' involvement in public policy debates: the no-buck-passing principle (Elliott 2011; Steel 2014) according to which

*To contact the author, please write to: Leibniz University Hannover, 30451 Hannover, Germany; e-mail: mathias.frisch@philos.uni-hannover.de.

Philosophy of Science, 87 (December 2020) pp. 979–990. 0031-8248/2020/8705-0017$10.00
Copyright 2020 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/710538 Published online by Cambridge University Press

scientists ought to offer their scientific judgment on hypotheses, when such judgments are needed to reach scientifically informed policy decisions and such judgments are possible.

In the next section, I briefly review the argument from inductive risk and the hedging defense. In section 3, I introduce the no-passing-the-buck principle and argue that this principle requires that it is possible to pair different types of putatively value-free hedged scientific hypotheses with decision-theoretic frameworks in a manner that allows policy makers to reach decisions even in the case of strongly hedged inputs by scientific experts. I then turn to climate policy as a case study. In light of the deep uncertainties governing precise, quantitative predictions of changes to the climate system and associated damages, many authors advocate a broadly precautionary approach to climate policy. Since precautionary reasoning appears to require as scientific input no more than a range of scientifically plausible futures, invoking a precautionary principle seems to allow us to satisfy both the value-free ideal and the no-passing-the-buck principle even in situations of deep uncertainty. Yet while precautionary reasoning arguably can underwrite a qualitative argument for the general need for climate action, it cannot support any specific climate policy target, at least not in a way compatible with the hedging defense. Deciding on a specific climate target requires scientists to stick out their necks and make precisely the kinds of value-laden judgments on what constitutes sufficient evidential support that lie at the heart of the argument from inductive risk. Thus, the hedging defense fails.

2. The Argument from Inductive Risk and the Hedging Defense. Rudner (1953) famously argued that "the scientist *qua* scientist" needs to make value judgments. Rudner points out that any decision to accept or reject a hypothesis involves a nonnegligible risk of error. Thus, scientists need to decide what constitutes sufficient evidence for accepting or rejecting a hypothesis, and these decisions can have ethically relevant consequences. Moreover, as Elliott (2011) and Steel (2016) have argued, when a scientific decision can have ethically relevant consequences, scientists have a moral responsibility to take these consequences into account (see also Douglas 2009). This, Steel argues, is just an instance of the general human responsibility to take foreseeable harmful consequences of our actions into account in decision making. The demand that scientists ought to take the ethically significant consequences of their scientific decisions into account results in a normative version of the argument from inductive risk: since decisions about sufficient evidence can have ethically significant consequences, there are cases in which ethical considerations should influence scientific standards on what counts as sufficient evidence to accept or reject a hypothesis. And, hence, there are circumstances in which scientists qua scientists ought to make value judgments.

https://doi.org/10.1086/710538 Published online by Cambridge University Press

Rudner's argument was criticized by Jeffrey (1956), who argued that scientists are not in the business of accepting or rejecting hypotheses but should instead report the probabilities of hypotheses given the available evidence. And it is these probabilities that scientists should then report to policy makers. I here want to focus on just one line of response to this criticism that was initiated by Rudner himself (and in fact was already discussed in Rudner's original paper before Jeffrey's criticism). Rudner points out that reports of probabilities can themselves be uncertain with a nonnegligible risk of error. And accepting or rejecting a certain probability distribution for a certain quantity can also have ethically significant consequences. For example, which probability distribution we accept for the equilibrium climate sensitivity, ECS (the amount of warming that results from a doubling of atmospheric CO_2 concentrations), has ethically significant effects on the social cost of carbon recommended by climate-economy integrated assessment models (Weitzman 2009; Frisch 2013). Thus, the argument from inductive risk reemerges at the level of second-order uncertainties (see also Steel 2016). This does not, however, settle the issue, since Jeffrey's reply can be generalized into a broader hedging defense: when there is a significant chance of error scientists ought to hedge their hypotheses sufficiently to report only claims about which no significant uncertainties remain.[1]

Hedging can occur along different dimensions and can take several different forms. First, hypotheses may be coarse-grained. Thus, instead of making precise predictions for future temperature increases, climate scientists might report temperature bands or intervals, within which future increases are expected to lie, or they might report only a sign for future temperature changes. Predictions may also be coarse-grained spatially in that scientists might only predict global average temperature changes rather than local temperatures.

Second, hypotheses can be modally weakened. Instead of accepting a specific value for a certain parameter, we can accept a probability distribution over different values of the parameter. When these probability distributions are themselves uncertain, scientists could report second-order probabilities by assigning them to different first-order probability distributions, or scientists could commit merely to probability bands instead of precise probabilities, or they might endorse merely "possibilistic" (see Betz 2013) claims and report only which predictions are scientifically possible.

Proponents of the argument from inductive risk state that the hedging defense fails, since nonnegligible risks of error can remain even for higher-order claims. Advocates of the hedging defense, by contrast, argue that any

1. Havstad and Brown (2017) call this defense the "deferred decision response." Their criticism focuses mainly on the fact that decisions carrying inductive risks occur at many different stages in the scientific process.

https://doi.org/10.1086/710538 Published online by Cambridge University Press

uncertainty can be made explicit until the claims scientists commit to "are actually as well-confirmed as [scientifically widely accepted] benchmark statements" (Betz 2013, 215). How do we adjudicate this disagreement? As Steel points out, not all weakening is acceptable as hedging, since in the limit we reach tautologies, which while not uncertain are not the kind of statements scientists are expected to make. Steel proposes that we limit the scope of discussion to *scientifically informative statements* and argues that advocates of the hedging defense have not been able to establish that such statements cannot be infected by nonnegligible uncertainties.

3. No Buck Passing and Climate Policy Decisions. How do we determine whether a statement is sufficiently informative? Steel maintains that both the possibilistically hedged statements discussed by Betz and the statements of probability bands discussed by Parker (2014) remain uncertain enough to provide a purchase for the argument from inductive risk. Yet his argument appeals to the assumption that even broadly accepted background theories can be less certain than scientific truisms such as 'coal burns', and one might question the normative import of the very general observation that science often falls short of (almost) complete certainty. Instead I here want to appeal to a richer, context-dependent criterion of informativeness: whether a claim is sufficiently informative partly depends on its normative implications in a given decision context.

In addition to the risks of falsely accepting and of falsely rejecting a hypothesis, there is a third risk scientists face: that of suspending judgment, which also can have significant ethical consequences. In the case of foundational research in physics, there may be no ethically significant consequences associated with suspending judgment on the truth of a hypothesis. The situation is different in the case of applied research. Suspending judgment on whether a drug is safe and effective may cost lives, as the controversy surrounding the approval of the AIDS medication AZT illustrates.

Thus, Steel proposes a generalized version of the argument from inductive risk that begins from the premises that scientists must make decisions whether to accept, reject, or suspend judgment about hypotheses and that such decisions involve a nonnegligible risk of ignorance (see also Wilholt 2009). According to Steel, a person is *ignorant* with respect to a hypothesis P if and only if either she is in error with respect to the truth of P or she suspends judgment with respect to P. Note that a person also is ignorant with respect to a hypothesis P if she accepts only a hedged claim V that is the result of coarse-graining or modally weakening P but suspends judgment with respect to P itself.

For scientists to suspend judgment can run afoul of what I take to be two reasonable normative principles proposed by Kevin Elliott: the *ethics* principle states that scientists have an ethical responsibility to consider major

https://doi.org/10.1086/710538 Published online by Cambridge University Press

societal consequences of their decisions, and the *no-passing-the-buck* principle states that it may be "socially harmful . . . for scientists to respond to uncertainties by withholding their judgment or providing uninterpreted data to decision makers" (2011, 55). One way in which scientists may violate the no-passing-the-buck principle is to commit only to a strongly hedged variant of a hypothesis *P* while withholding judgment on *P* itself.

Obviously, there are situations in which, considering the available evidence and the normative consequences at stake, a relatively strongly hedged variant *V* of *P* may be appropriate for scientists to pass on to decision makers. Thus, I am not arguing here that scientists always have to make definitive unhedged judgments or that such judgments are always needed for policy decisions. What an appropriate level of hedging is depends on the evidence available in support of a hypothesis but also on the context of the decision problem at issue. Depending on the context, suitably hedged claims may be sufficient. Passing on only a strongly hedged variant *V* of *P* to policy makers may be appropriate, when there exist significant uncertainties concerning the adequacy of *P* and there is a decision procedure available that uses *V* as its input and, given our preferences and normative commitments, provides a ranking of different actions as output. But if (i) there is no decision procedure available that can use *V* as its input, (ii) the stakes are high enough, and (iii) there is additional scientific evidence available concerning *P* or a variant of *P* less hedged than *V*, then scientists may need to 'stick out their necks' and provide decision makers with predictions that may meet only a lower standard of evidence than a more fully hedged prediction does.

What do I mean by a decision procedure being available? According to the value-free ideal, there exists a sharp division of labor between scientists, on the one hand, and policy makers or stakeholders, on the other: scientists provide policy makers with value-free scientific expertise on an issue, and policy makers combine this scientific input with normative and other evaluative considerations to reach policy decisions. Different types of scientific input need to be paired with different types of decision procedures. If the hedging defense is to be compatible with the no-passing-the-buck principle, there have to be decision procedures available that take suitably hedged predictions as their input. Conversely, if a given policy issue requires a certain type of decision procedure, then scientific advice, if possible, has to be as fine-grained and modally committed as that decision procedure requires.

What type of ethical or decision frameworks are appropriate in a given context is itself a difficult issue. For ease of exposition I here restrict myself to broadly consequentialist frameworks, yet I do not intend this restriction to imply that I do not take considerations of justice to be important as well. The most fine-grained and least hedged predictions are predictions that a decision-relevant quantity will definitely take on a specific value given specific policy decisions. If deterministic predictions are possible without

https://doi.org/10.1086/710538 Published online by Cambridge University Press

significant uncertainties, then decision makers can use these as input in a cost-benefit analysis and choose the action that, given their utility function, maximizes overall utility or welfare.

If all that scientists can commit to with certainty are probabilistic predictions, then expected utility theory provides an appropriate decision framework. In fact, expected utility theory plays a particularly prominent role in climate policy discussions. It provides the framework for economy-climate optimization integrated assessment models (IAMs; see, e.g., Nordhaus and Boyer 2003), which are used to predict optimal emission pathways. Different IAMs recommend dramatically different pathways, and there is a sizable literature criticizing several of the central assumptions made by prominent IAMs. One overarching lesson to be gleaned from these discussions is that various probabilistic inputs into IAMs are infected by often quite large uncertainties (Frisch 2017, 2018).

Now, the philosopher John Broome argues that these uncertainties are no reason to abandon the framework of expected utility theory: "The lack of firm probabilities is not a reason to give up expected value theory. . . . Stick with expected value theory, since it is very well founded, and do your best with probabilities and values" (2012, 129). The hedging defense, by contrast, would recommend that scientists should not commit to specific probability distributions and instead hedge their hypotheses even further, requiring us to adjust our decision framework accordingly. Indeed, in light of the deep uncertainties plaguing probabilistic climate predictions, one might argue that climate scientists should only commit to specifying ranges of scientifically plausible outcomes. There are a number of decision frameworks available that require as input only a set or range of possible outcomes, such as maximin or minimax regret, which requires only a matrix of possible outcomes over a set of decision options.

Arguably even further hedging is possible, if possible outcomes cannot be specified precisely numerically. Steel (2014) articulates and defends a qualitative precautionary approach to decision making under deep uncertainty, using climate change as one of his paradigm examples. One version of a precautionary principle states that "if a scientifically plausible mechanism exists whereby an activity can lead to a catastrophe, then that activity should be phased out or significantly restricted" (28). Like maximin or related rules, Steel's precautionary approach proceeds by evaluating a matrix of possible actions versus possible states of the world, but unlike the former, the precautionary principle requires no quantitatively precise inputs. In particular, the principle is insensitive to small differences between states and asks us to look only at possible worlds or states of affairs in which the precautionary measures would make a sizable difference. In the case of climate change, the principle requires us to evaluate decision matrix 1 (table 1; see Steel 2014, 60). In trying to decide whether to reduce emissions, we are

https://doi.org/10.1086/710538 Published online by Cambridge University Press

TABLE 1. Decision Matrix 1

	Good World	Bad World	Ugly World
No mitigation	Little costs	Catastrophe	Catastrophe
Mitigation	Modest costs	Modest costs	Catastrophe

asked to consider three kinds of possible worlds: "good world" in which climate change will be much less severe than feared with only minimal negative impacts, even if we do not undertake any meaningful mitigation measures; "bad world" in which doing nothing would result in catastrophic outcomes but mitigation would allow us to prevent most of the negative effects of climate change; and "ugly world" in which it is already too late to prevent catastrophic climate change. Since mitigation measures only make an appreciable difference in a bad world, we need only to consider this kind of world, and since mitigating would prevent a catastrophe in this world while doing nothing would not, the precautionary principle says that we should engage in mitigation measures.

What particular version of a precautionary principle may be defensible or plausible is not important for the argument I want to develop here. What is important for the dialectics of the argument is that there is some such decision strategy available that will allow scientists to retreat to reasonably certain strongly hedged claims without violating the no-passing-the-buck principle—that is, that there are decision procedures that will allow policy makers to reach decisions on issues of vital importance, even with the highly hedged inputs that scientists clinging to the value-free ideal can at best provide. In response to Steel and Elliott's worry that sometimes scientists cannot withhold judgment on an issue, even when significant uncertainties remain, because the social or moral risks of ignorance are too high, the defender of the value-free ideal will have to point to something akin to the precautionary strategy to maintain that even relatively weak, severely hedged claims can provide sufficient input for policy makers to act.

Yet, as I want to argue now, the scope of precautionary reasoning is somewhat limited and, pace Steel, even in the context of the climate problem needs to be supplemented by other decision strategies that require less hedged scientific inputs. Consider one of the central aims of climate policy discussions: policy makers are looking for guidance on not only whether we should reduce our greenhouse gas emissions but also what appropriate mitigation pathways are. How much should we reduce emissions and how quickly? Will the consequences of climate change be rather benign in the medium term, as one of the developers of climate-economy IAMs (Tol 2002) has argued, or are dramatic mitigation measures needed, as other economists and many climate scientists insist (Stern 2013; Schellnhuber et al. 2016)? While the precautionary argument we have considered so far may tell us that some mitigation

https://doi.org/10.1086/710538 Published online by Cambridge University Press

measures are needed, it does not tell us what ought to be to done. It does not provide us with a policy target—it does not, for example, tell us whether we should aim to keep the rise in global temperatures to 2°C or aim for a more ambitious 1.5°C goal.

We can try to model a decision process aimed at settling on a specific climate policy target as a series of pairwise comparisons of two possible climate targets, to which we apply precautionary reasoning. In this case we would have to evaluate decision matrix 2 (table 2), which allows us to compare three policies—do nothing and two different climate targets—in four different worlds: in addition to the good, bad, and ugly worlds, we now also need to consider a world ('worse world') in which the weaker emission target would be insufficient to prevent a catastrophe but the stronger target would be sufficient to prevent a catastrophe. If such a world is scientifically plausible, then precautionary reasoning will recommend that we adopt the stronger of the two emission targets.

The problem is that the process of comparing a given policy with a more stringent competitor can be repeated for ever more stringent mitigation policies. We can imagine ordering policies P according to the stringency of their target and in subsequent steps comparing ever more stringent neighboring pairs of policies. At each step n of comparing the policy $P(n-1)$ recommended at step $n-1$ with its more stringent neighboring policy, the precautionary approach will recommend that we adopt the more stringent alternative, until one of two limiting conditions is satisfied: (i) the costs of mitigation policies will become comparable to their benefits or (ii) we have reached a mitigation policy so stringent that there are no plausible scenarios that allow us to discriminate it from a more stringent rival policy.

If the costs of mitigation policies become comparable to their benefits, then one of the conditions of applicability of the precautionary principle is no longer satisfied. Precautionary reasoning, as Steel understands it, only applies when a consistency condition is satisfied, which "states that a precaution should not be precluded by the same version of [a precautionary principle] used to justify it" (Steel 2014, 28). But if the costs of mitigation policies became comparable to their benefits—that is, if dramatic, disruptive mitigation measures aimed at preventing environmental catastrophes were themselves to result in catastrophic consequences—then precautionary

TABLE 2. Decision Matrix 2

	Good World	Bad World	Worse World	Ugly World
No mitigation	Little costs	Catastrophe	Catastrophe	Catastrophe
Weaker target	Modest costs	Modest costs	Catastrophe	Catastrophe
Stronger target	Modest costs+	Modest costs+	Modest costs+	Catastrophe

https://doi.org/10.1086/710538 Published online by Cambridge University Press

reasoning would both recommend these policies and recommend against them (see also Gardiner 2006).

Both limiting conditions depend on an assessment of what scenarios are scientifically plausible. It may be possible, for any mitigation policy we investigate, to imagine a scenario in which this policy leads to catastrophe. But unless the scenario in question is scientifically plausible, it need not be considered in applications of precautionary reasoning. Similarly, it may be possible for each pair of mitigation policies to imagine a scenario in which the more lenient but not its more stringent rival results in catastrophe.

Now, assessing either when there exist no more scientifically plausible scenarios distinguishing between two ever more stringent mitigation policies or when mitigation policies become so severe that they themselves could plausibly result in catastrophic costs may require precisely the kind of more specific yet uncertain judgments that the hedging strategy aims to avoid: When we compare different mitigation pathways and climate projections, which projections are plausible enough that they need to be included in our assessment? Many of the probability distributions derived for the climate sensitivity ECS within different modeling frameworks are 'fat-tailed' distributions that go to zero more slowly than a normal distribution and assign very high values of ECS small but nonnegligible probabilities (Weitzman 2009). When are mitigation policies stringent enough that catastrophic scenarios resulting from high values of ECS become so implausible that they no longer need to be taken seriously scientifically? How do we distinguish intuitively improbable yet scientifically legitimate threats from those that are scientifically implausible? Several climate threats exhibit tipping point behavior, but at which temperatures various tipping points will be reached is deeply uncertain (Smith et al. 2009). At which point is a mitigation strategy stringent enough that the threat of the climate system's reaching catastrophic tipping points is no longer scientifically plausible? While it is clearly not the case that scientists are completely ignorant about answers to these questions, plausible answers are infected by deep uncertainties.

I am not here arguing that scientists always have to provide policy makers with definitive, unhedged judgments and that policy makers are helpless if they do not receive guidance in the form of completely unhedged claims. Rather, I am arguing that there are decision contexts, such as that of deciding on a specific climate policy target, in which policy makers require information that is less fully hedged than the hedging strategy would require and, hence, that there are decision contexts in which scientists need to 'stick out their necks' and make judgments that go beyond what they can be reasonably certain about.

As a matter of fact, climate scientists do not follow the recommendation of the hedging strategy. Consider discussions of the two climate targets—the 1.5°C and 2°C targets—of the Paris Agreement. On the one hand, some

https://doi.org/10.1086/710538 Published online by Cambridge University Press

climate scientists have argued for the need of a target more stringent than the 2°C target, on the basis of our best scientific understanding of the climate system. On the other hand, there are lively discussions of the uncertainties and knowledge gaps concerning the effects of different targets, as is evident, for example, from a recent special Intergovernmental Panel on Climate Change (IPCC) report (Hoegh-Guldberg et al. 2018) that catalogs the following gaps in our knowledge and corresponding scientific needs: "Robust probabilistic models which separate the relatively small signal between 1.5°C versus 2°C from background noise, and which handle the many uncertainties associated with non-linearities, . . . local scales, latent or lagging responses in climate" (sec. 3.7); a better understanding of ranges of uncertainties under different pathways; an understanding of critical thresholds for extreme events; and a better understanding of various tipping point behaviors.

In fact, IPCC reports go to great lengths to catalog various uncertainties, as good scientific methodology would recommend. Thus, one might take the IPCC report as providing an example of the hedging defense, as Betz does (2013, sec. 5). Yet while the IPCC report provides hedged judgments, the resulting claims are not sufficiently hedged not to fall prey to considerations of inductive risk. Many of the characterizations of uncertainties are based on expert judgments and, pace Betz, are paradigm examples of ethically responsible reactions to a situation of inductive risk. While trying to communicate to policy makers that significant uncertainties concerning predictions of the future evolution of the climate system remain, the IPCC report and other scientific publications try to give guidance that is detailed enough to be useful in policy decisions.

4. Conclusion. In this article I argued that the hedging defense in response to the argument from inductive risk fails and runs afoul of the no-passing-the-buck principle. I developed my argument by considering a precautionary approach to climate policy, arguing that a precautionary approach that requires only strongly hedged and reasonably certain inputs cannot yield a specific climate target. Climate scientists can only provide input on what specific climate targets to pursue if they are willing to "stick out their necks" and are willing to make judgments that go beyond what they can be reasonably certain about. Thus, climate policy decisions provide a case in which the hedging defense runs afoul of the no-passing-the-buck principle.

What an appropriate decision framework is in this context remains, at least as far as the current discussion is concerned, an open question. Defenders of the hedging defense are right, I take it, in arguing that scientific integrity demands that uncertainties not be swept under the rug and be made explicit as best we can. Thus, I would argue that Broome's position that we stick with expected utility theory even in the face of deep uncertainties is

https://doi.org/10.1086/710538 Published online by Cambridge University Press

inadequate. By the same token, retreating to fully hedged claims leaves policy makers with less guidance than scientists reasonably could—and, I argue, ought to—provide. The question how these demands that to some extent pull in different directions might be integrated into a decision framework goes beyond the scope of this article.

REFERENCES

Betz, Gregor. 2013. "In Defence of the Value Free Ideal." *European Journal for Philosophy of Science* 3 (2): 207–20.

Broome, John. 2012. *Climate Matters: Ethics in a Warming World*. Norton Global Ethics Series. New York: Norton.

Douglas, Heather. 2000. "Inductive Risk and Values in Science." *Philosophy of Science* 67 (4): 559–79.

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

Elliott, Kevin C. 2011. *Is a Little Pollution Good for You? Incorporating Societal Values In Environmental Research*. Oxford: Oxford University Press.

Frisch, Mathias. 2013. "Modeling Climate Policies: A Critical Look at Integrated Assessment Models." *Philosophy and Technology* 26 (2): 117–37.

———. 2017. "Climate Policy in the Age of Trump." *Kennedy Institute of Ethics Journal* 27, no. 2 (Suppl.): E87–E106.

———. 2018. "Modeling Climate Policies: The Social Cost of Carbon and Uncertainties in Climate Predictions." In *Climate Modelling: Philosophical and Conceptual Issues*, ed. Elisabeth A. Lloyd and Eric Winsberg, 413–48. Cham: Springer.

Gardiner, Stephen M. 2006. "A Core Precautionary Principle." *Journal of Political Philosophy* 14 (1): 33–60.

Havstad, Joyce C., and Matthew J. Brown. 2017. "Inductive Risk, Deferred Decisions, and Climate Science Advising." In *Exploring Inductive Risk: Case Studies of Values in Science*, ed. Kevin C. Elliott and Ted Richards. Oxford: Oxford University Press.

Hoegh-Guldberg, O., et al. 2018. "Impacts of 1.5°C Global Warming on Natural and Human Systems." In *Global Warming of 1.5°C: An IPCC Special Report*, ed. V. Masson-Delmotte, et al. New York: Intergovernmental Panel on Climate Change.

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

Nordhaus, William D., and Joseph Boyer. 2003. *Warming the World: Economic Models of Global Warming*. Cambridge, MA: MIT Press.

Parker, Wendy. 2014. "Values and Uncertainties in Climate Prediction, Revisited." *Studies in History and Philosophy of Science* A 46 (June): 24–30.

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

Schellnhuber, Hans Joachim, et al. 2016. "Why the Right Climate Target Was Agreed in Paris." *Nature Climate Change* 6 (7): 649–53.

Smith, Joel B., et al. 2009. "Assessing Dangerous Climate Change through an Update of the Intergovernmental Panel on Climate Change (IPCC) 'Reasons for Concern.'" *Proceedings of the National Academy of Sciences* 106 (11): 4133–37.

Steel, Daniel. 2014. *Philosophy and the Precautionary Principle: Science, Evidence, and Environmental Policy*. Cambridge: Cambridge University Press.

———. 2016. "Climate Change and Second-Order Uncertainty: Defending a Generalized, Normative, and Structural Argument from Inductive Risk." *Perspectives on Science* 24 (6): 696–721.

Stern, Nicholas. 2013. "The Structure of Economic Modeling of the Potential Impacts of Climate Change: Grafting Gross Underestimation of Risk onto Already Narrow Science Models." *Journal of Economic Literature* 51 (3): 838–59.

Tol, Richard. 2002. "Estimates of the Damage Costs of Climate Change." Pt. 1, "Benchmark Estimates." *Environmental and Resource Economics* 21 (1): 47–73.

https://doi.org/10.1086/710538 Published online by Cambridge University Press

Weitzman, Martin L. 2009. "Additive Damages, Fat-Tailed Climate Dynamics, and Uncertain Discounting." *Economics: The Open-Access, Open-Assessment E-Journal*, 3 (2009-39): 1–22. http://dx.doi.org/10.5018/economics-ejournal.ja.2009-39.
Wilholt, Torsten. 2009. "Bias and Values in Scientific Research." *Studies in History and Philosophy of Science* A 40 (1): 92–101.

https://doi.org/10.1086/710538 Published online by Cambridge University Press