
Climate Change and Second-Order Uncertainty: Defending a Generalized, Normative, and Structural Argument from Inductive Risk

Daniel Steel

University of British Columbia

Philosophers of science have recently debated whether second-order uncertainty in climate change forecasts demonstrates the applicability of the argument from inductive risk (AIR) to this case. This paper defends a generalized, normative, and structural interpretation of AIR to address challenges raised in this literature. The interpretation of AIR proposed is generalized by including the possibility that scientists may suspend judgment rather than accept or reject a hypothesis. In addition, it distinguishes between descriptive and normative versions of AIR, and provides reasons for preferring the latter. Finally, it emphasizes advantages of applying AIR at a structural rather than individual level.

1. Introduction

This article critically examines a recent philosophical debate on the role of values in climate change forecasts, such as those found in assessment reports of the Intergovernmental Panel on Climate Change (IPCC). On one side, several philosophers (Biddle and Winsberg 2010; Intemann 2015; Winsberg 2010; Winsberg 2012) insist that the argument from inductive risk (AIR), as developed by Rudner (1953) and Douglas (2009) among others, applies to this case. AIR aims to show that ethical value judgments should influence decisions about what is sufficient evidence for accepting scientific hypotheses that have implications for policy issues. Advocates of extending AIR to climate science claim that values are deeply implicated in attempts to probabilistically quantify uncertainty, for instance, about future rises in sea levels. This claim, if correct, would undermine a classic line of response to AIR made by Jeffrey (1956). According

to Jeffrey, “the scientist qua scientist” should avoid value judgments by not accepting hypotheses; instead, scientists should assess the probabilities of hypotheses and pass these probabilities along to policy makers. But if ethical value judgments are already in the probabilities, then Jeffrey’s proposal—at least when it comes to climate change—is a nonstarter. On the other side of the debate are several scholars who argue against the claim that decisions made in the course of developing climate models are inevitably influenced by ethical values (Betz 2013; Morrison 2014; Parker 2014).

To discuss this dispute in greater depth, it is helpful to introduce the concept of second-order uncertainty, that is, uncertainty about an assessment of uncertainty. To illustrate, consider the claim that under a “business as usual” emissions scenario, global mean sea level (GMSL) will rise by less than 1 meter by 2100. A first-order assessment of uncertainty might judge the probability of this claim to be 66%. Second-order uncertainty would mean uncertainty about that probability estimate itself, and instead of a single number an interval, say 33% to 90%, might be given. Winsberg (2012) argues that second-order uncertainty is ineliminable when it comes to climate change forecasts and furthermore that this entails an inevitable role of values in assessments of first-order uncertainty. His critics, meanwhile, contest this reasoning at one or both of these steps. They argue that it is possible to substantially reduce second-order uncertainty in scientific policy advice (Betz 2013; Parker 2014) or that second-order uncertainty does not necessarily entail that values have influenced how uncertainty has been quantified (Betz 2013; Morrison 2014; Parker 2014). In this essay, I develop a generalized, normative, and structural version of AIR designed to address these objections, and demonstrate its relevance to climate science.

To start, in section 2.1, I claim that talk of acceptance and rejection in AIR should be understood in relation to claims that are informative within some decision-making context, such as climate change policy. This clarification supports an argument that critics have failed to adequately explain how second-order uncertainty regarding climate change can be reduced to a negligible level. The claims suggested by these critics are either significantly uncertain in their own right or not sufficiently informative to be useful for decision-making. I regard this first point as a clarification of AIR, rather than a modification or extension. However, the next three points constitute developments that distinguish the version of AIR defended here from its classic formulations.

First, in section 2.2, I propose that AIR be generalized to allow that scientists may suspend judgment rather than accept or reject an informative hypothesis. In contrast, classic statements of AIR typically presuppose a situation in which a specific informative hypothesis must be accepted or

rejected by a deadline, and hence suspending judgment is not an option (cf. Rudner 1953). The generalization of AIR I propose leads to an argument against extremely high standards of evidence for accepting statements related to climate change. When scientific uncertainty is a serious concern, as in the case of climate change forecasts, extremely high standards of evidence naturally lead to suspension of judgment. That in turn risks transforming scientific uncertainty into a reason for inaction in the face of serious problems.

Secondly, I argue that AIR should be interpreted normatively rather than descriptively. I agree that the existence of second-order uncertainty alone is insufficient to establish that ethical or social values did in fact influence climate modeling. But this point only counts against descriptive versions of AIR, which aim to show that value judgments of an ethical or social nature inevitably influence scientists' decisions about which hypotheses to accept. While some classic statements of AIR do suggest such an interpretation (Rudner 1953), the argument can also be interpreted normatively (Douglas 2009; Elliott 2011). Construed in this way, AIR aims to show that, under certain circumstances, ethical values *should* influence judgments about what counts as sufficient evidence for accepting a claim. In section 3.1, I argue that there are strong reasons for preferring a normative interpretation of AIR.

Finally, in section 3.2, I consider what I call the *democratic objection*, according to which AIR threatens to circumvent democratic processes by encouraging a small group of scientists to impose their values upon society (Betz 2013; Steele 2012). I suggest that this objection implicitly presumes that AIR is primarily applied at the level of decisions made by individual scientists or research teams. In contrast, I propose that applications of AIR should emphasize decisions about evidential standards in science made at structural rather than individual levels. I explain how focusing on structural applications of AIR leads to a better answer to the democratic objection, and illustrate this point with a discussion of the structure of the IPCC process.

2. Scientific Uncertainty and AIR

In this section, I take a closer look at arguments by Betz (2013) and Parker (2014) that second-order uncertainty in IPCC reports can be effectively minimized. In section 2.1, I argue that neither has adequately explained how this is possible. In section 2.2, I show that extremely high standards of evidence, such as those recommended by Betz, do not avoid a generalized form of AIR, because such standards often incur costs due to extended suspension of judgment.

2.1. Can Science Advice Avoid Second-Order Uncertainty?

AIR is founded on the premise that accepting and rejecting hypotheses is one role of science and that such decisions come with a non-negligible risk of error. When these errors have ethically significant consequences, the argument concludes, ethical values are relevant to deciding what constitutes sufficient evidence for accepting or rejecting a claim (Braithwaite 1953; Churchman 1956; Douglas 2000, 2009; Elliott 2011; Rudner 1953; Steel 2010). The argument can be formulated as follows:

1. Scientists must make decisions about whether to accept or reject hypotheses.
2. Such decisions invariably involve a non-negligible risk of error.
3. Thus, scientists must make decisions about what constitutes sufficient evidence for accepting and rejecting hypotheses.
4. Decisions about what constitutes sufficient evidence can have ethically significant consequences.
5. If a scientific decision has ethically significant consequences, then scientists should take these into consideration.
6. Thus, there are circumstances in which ethical considerations should influence scientific standards of what counts as sufficient evidence.

The reasoning in the first three steps is that scientists are in a situation wherein evidence that fails to conclusively confirm (or refute) a hypothesis must be taken to be sufficient for accepting (or rejecting) it. Thus, decisions must be made about where to place the evidential threshold, that is, about what kind and how much evidence is enough. Premise 4 states that such decisions can have ethically significant consequences. For instance, environmental protection agencies often must issue rules about safe and unsafe exposure levels of various substances. Consequently, decisions must be made about which methods and data should be taken to be sufficient for establishing such claims, and these decisions can have implications for public health (cf. Steel 2011). Premise 5 in effect states that scientists cannot claim any special exemption from the ordinary human responsibility to consider foreseeable harmful consequences of their actions (cf. Douglas 2009, 2014). Premises 3 through 5, then, lead to the conclusion that, when scientific decisions about what constitutes sufficient evidence have ethically significant consequences, values of an ethical nature should influence these decisions.

One objection to AIR is to reject premise 2, according to which decisions about whether to accept or reject hypotheses involve a non-negligible risk of error (Levi 1980; Lacey 1999; cf. Steel 2013). Betz pursues this

line of objection, but with regard to second-order uncertainty (2013, pp. 211–214). He suggests that scientists should not accept “plain hypotheses” (e.g., GMSL increase by 2100 will be less than 1 meter) when they are uncertain but only “hedged hypotheses” that involve some qualification regarding probability, uncertainty, or degree of confirmation (e.g., there is a 66% probability that GMSL increase by 2100 will be less than 1 meter). However, defenders of AIR insist that assessments of uncertainty are inevitably uncertain themselves (Douglas 2009, p. 85; Winsberg 2012). Betz considers this response and replies as follows:

Does accepting hedged hypotheses, which are, thanks to epistemic qualification and conditionalization, weaker than plain ones, still involve substantial error probabilities? ... This much is clear: Sometimes a probability (or, generally, an uncertainty) statement cannot be inferred, based on the available evidence, without a substantial chance to err. But Douglas, or the methodological critique [i.e., AIR], needs more: Every hedged (e.g. epistemically qualified or suitably conditionalized) hypothesis involves substantial error probabilities.—And that seems to be plainly false. (Betz 2013, p. 214)

Betz then proceeds to suggest how descriptions of uncertainty could be qualified to such an extent that no substantial risk of error remains. Those suggestions will be examined presently. For now, consider the relationship between AIR and the above statement from Betz.

The issue at hand is whether premise 2 is true of decisions about whether to accept “hedged” hypotheses. Do such decisions invariably involve non-negligible uncertainty? Douglas answers yes, and Betz answers no. However, the dispute here is somewhat more complex than it might appear at first, because the scope of acceptance decisions at issue demands some clarification. Premise 2 cannot be meant to cover acceptance decisions regarding all possible statements. Otherwise, the observation that no substantial uncertainty is implicated in accepting tautologies or elementary truths of arithmetic would suffice to refute it. Consequently, premise 2 should be limited to the sorts of acceptance decisions that scientists are expected to make. In the case of climate change, these would include answers to such questions as the following. Is climate change primarily due to anthropogenic causes? Will GMSL increase by more than 1 meter by 2100? Will global mean temperatures increase by more than 2° Celsius? In short, science is expected to provide *informative*, if not necessarily conclusive, answers to questions. Properly understood, therefore, premise 2 makes a claim about the acceptance of informative scientific

claims.¹ Of course, a decision to accept an uninformative claim involves no serious risk of error. The harder task is to show that the same can also be true of scientific claims that are informative with respect to climate change policy. In the remainder of this section, I argue that Betz (2013) and Parker (2014) have not adequately explained how this can be achieved.

Betz argues that scientific claims given as advice to policy makers “can be justified beyond reasonable doubt, even under uncertainty” (2013, p. 215). In other words, although some residual uncertainty is unavoidable, Betz proposes that by “hedging” their claims scientists can reduce this uncertainty to a negligible level. Betz offers the following elaboration of this idea.

I take it that there is a vast corpus of empirical statements which decision makers—in private, commercial or public contexts rightly take for granted as plain facts. I’m thinking for instance of results to the effect that plutonium is toxic, the atmosphere comprises oxygen, coal burns, CO₂ is a greenhouse gas, Africa is larger than Australia, etc. ... The idea is that scientific policy advice comprises [only those] results that are equally well confirmed as those benchmark statements—so that policy makers can rely on policy advice in the same way as they are used to rely on other well-established facts. By making all of the policy-relevant uncertainties explicit, scientists can further and further hedge their scientific reported findings until the results are actually as well confirmed as the benchmark statements. (Betz 2013, p. 215)

According to Betz, then, scientific claims regarding climate change should be qualified by explicit discussion of uncertainties until they are as well confirmed as such statements as “coal burns” or “Africa is larger than Australia.” Betz cites the *Guidance Note for Lead Authors of the IPCC Fifth Assessment Report on Consistent Treatment of Uncertainties* (Mastrandrea et al. 2010) to support this recommendation (Betz 2013, pp. 216–18). In particular, Betz calls attention to the *Guidance Note’s* suggestion that a description of the “state of scientific understanding” be associated with claims (Betz 2013, p. 217). Examples of such descriptions include, “A variable is ambiguous, or the processes determining it are poorly known or not amenable to measurement,” and, “A probability distribution or a

1. Betz also appears to interpret AIR in this manner. He formulates AIR as resting on the premise, “To arrive at (adopt and communicate) policy-relevant results, scientists have to adopt or reject plain or hedged hypotheses under uncertainty” (2013, p. 213). Thus, in Betz’s formulation, it seems that AIR pertains to statements that are informative for policy decisions, as uninformative claims would presumably not be relevant.

set of distributions can be determined for the variable either through statistical analysis or through use of a formal quantitative survey of expert views” (Betz 2013, p. 217).

However, it is doubtful that the *Guidance Note* illustrates Betz’s proposal that scientists, when giving policy advice, should limit themselves to claims that are as well confirmed as CO₂ is greenhouse gas, etc. The difficulty is that assessments of the state of scientific understanding on topics relating to climate change are also likely to be uncertain (not to mention politically controversial!). For example, whether a probability distribution can be determined for a variable depends on a number of potentially uncertain factors. These include whether the assumptions of the statistical model are accurate and perhaps even philosophical debates about the underlying rationale of the statistical approach (e.g., frequentist versus Bayesian). Even negative assessments (e.g., that the value of a variable is ambiguous) can be uncertain. Perhaps the value is not ambiguous, but only appears such due to some error in data analysis. In sum, Betz has not adequately explained how enacting the *Guidance Note*’s recommendations would enable second-order uncertainty to be reduced to a negligible level.

Similar difficulties confront Betz’s other proposals for how to eliminate second-order uncertainty in scientific advice. For example, he proposes that general circulation models should be interpreted as generating “possibilistic,” rather than probabilistic, predictions about climate change (2007, 2009, 2010, 2015). A possibilistic prediction asserts that an outcome, such as a 1-meter rise in GMSL by 2100, is a serious possibility, where a “statement *P* is seriously possible if and only if it is consistent with the entire body of background knowledge *K*” (Betz 2015, p. 195). Betz (2010, p. 93) suggests that his possibilistic approach avoids the argument from inductive risk as championed by Rudner (1953). And Betz (2013, p. 214) gives, “It is possible (consistent with what we know), that ...”, as an example of a type of claim that could be asserted with negligible uncertainty. The thought, then, seems to be that second-order uncertainty could be reduced to a negligible level if a possibilistic approach to assessing uncertainty were adopted.

Although Betz’s discussion of possibilistic predictions is insightful, his approach does not escape second-order uncertainty because it can be uncertain whether a statement is seriously possible. A statement might be regarded as seriously possible so long as it has *not been refuted*, or only if its consistency with background knowledge has been *verified* (Betz 2015, p. 195). Yet uncertainty can arise in either case.² Betz primarily addresses

2. That it may be significantly uncertain whether a hypothesis has been refuted is, of course, a familiar theme in philosophy of science (Duhem 1991; Lakatos 1970; Kuhn 1970). Betz (2007, pp. 6–7) appears to acknowledge this.

the verification option. As Betz notes, it is not at all trivial to verify that statements about long-term consequences of climate change are seriously possible in the sense he defines (2015, p. 195). In particular, climate models, like scientific models generally, contain simplifying assumptions known not to be strictly true. Consequently, Betz (2015) considers several lines of argument that might be given to support the claim that a not-strictly-true model verifies that a climate change forecast is consistent with background knowledge. The model might be argued to be a good enough approximation of reality for intended purposes, or to have a good track record of predictive success. But such arguments about models, even when reasonable, are typically uncertain. Approximations thought to be harmless might surprisingly undermine predictions, and assessments of past predictive success might be uncertain or not good indicators of future accuracy. Furthermore, other sources of uncertainty are present in Betz's possibilistic approach. For example, judgments about serious possibility rest on prior decisions about what to accept as "background knowledge." From Betz's (2015) discussion it seems that current scientific theories are included in the body of background knowledge.³ But it is often thought that even well confirmed scientific theories fall far short of the certainty of truisms such as "coal burns" or "Africa is larger than Australia" (cf. Stanford 2006).

As a last resort, Betz suggests that second-order uncertainty can be eliminated for all practical purposes by a "fully comprehensive" description of relevant uncertainties (2013, p. 214). Of course, frank disclosure of uncertainties is laudable in climate science and elsewhere. However, it is questionable whether a "fully comprehensive" list of uncertainties is possible in relation to highly complex processes, such as climate change, wherein the potential for "surprises" is a familiar feature (Liu et al. 2007). In such circumstances, uncertainty about whether all sources of uncertainty have been adequately considered appears inevitable and non-negligible.

Taking a different approach, Parker suggests that the use of probability bands can be an effective means of reducing second-order uncertainty in climate change forecasts (2014, pp. 27–8; cf. Parker 2010a, p. 271; Parker 2010b, p. 996). Parker thus criticizes Winsberg for focusing "on methods for producing precise probabilistic estimates—which are known to be artificially precise, and she proposes to emphasize instead methods for producing coarser depictions of uncertainty" (Parker 2014, p. 27). To

3. If background knowledge did not include accepted scientific theories, then it would appear that nearly every climate forecast is seriously possible according to Betz's definition (since there would be nothing to link past observations with predictions). In this case, second-order uncertainty would be removed at the cost rendering claims about serious possibilities uninformative.

illustrate, consider the following statement taken from the IPCC's AR5, *The Physical Science Basis*, regarding GMSL rise by the end of this century.

GMSL rise for 2081–2100 (relative to 1986–2005) for the RCPs [Representative Concentration Pathways] will *likely* be in the 5 to 95% ranges derived from CMIP5 [Coupled Model Intercomparison Project Phase 5] climate projections in combination with process-based models of glacier and ice sheet surface mass balance, with possible ice sheet dynamical changes assessed from the published literature. (Stocker et al. 2013, p. 98)

The minimum of the range in question is 0.26 meters while the maximum is 0.98 meters (Stocker et al. 2013, pp. 98–99). The above claim, therefore, entails that if one of the five RCPs considered corresponds to the actual concentration pathway of greenhouse gasses in the atmosphere, then it is likely (i.e., probability $\geq 66\%$) that the GMSL rise by 2100 will be less than 1 meter.

The use of probability bands, rather than point estimates, is naturally understood as an attempt to reduce second-order uncertainty, as Parker suggests. But even if reduced to some extent, second-order uncertainty remains a serious concern for claims about probability intervals for several reasons. One is structural model error (SME), wherein the model incorrectly represents the dynamics of the target system. According to Roman Frigg and others, “if a nonlinear model has only the slightest SME, then its ability to generate decision-relevant probabilistic predictions is compromised” (Frigg et al. 2014, p. 32; cf. Winsberg 2012, pp. 116–22). In these circumstances, probabilistic forecasts may err not merely about precise probabilities but also as to which outcomes are relatively more or less probable, such as whether or not GMSL rise by 2100 will exceed 1 meter (Frigg et al. 2014, p. 39). Moreover, a case can be made that SME is a predicament of climate modeling (see Frigg, Smith, and Stainforth 2013; Stainforth et al. 2007). Similarly, Martin Weitzman (2009) argues that due to structural uncertainty probabilities of extreme future outcomes, such as a 3-meter rise in sea level by 2100, cannot be accurately estimated in a climate system that is undergoing rapid and largely unprecedented change. In addition, there is an empirical reason for suspecting second-order uncertainty in this case, namely, that past IPCC forecasts have tended to underestimate GMSL increases (Rahmstorf, Foster, and Cazenave 2012).

Parker and Betz deserve credit for examining ways in which second-order uncertainty might be managed or limited in the context of climate change science. But neither has seriously challenged the claim made by Rudner (1953) and Douglas (2009) that non-negligible uncertainty is

inevitable in the course of giving informative science advice to policy makers.⁴

2.2. The Risks of Ignorance

The arguments examined in section 2.1 were founded on the idea that extremely demanding evidential standards can enable one to avoid AIR. In this section, I argue that this is not true, because to adopt such exacting standards is to make a decision that may carry substantial ethical and practical consequences in its wake (cf. Cranor 1993; John 2015, pp. 7–8; Steel 2015). More specifically, I suggest that AIR can be generalized from risk of *error* to risk of *ignorance*. Strict evidential standards reduce the risk of error, but they also increase the risk that one will fail to accept or reject any informative claim whatever (i.e., suspend judgment), and hence remain ignorant. Moreover, ignorance can be harmful in the context of environmental and public health policy, as I explain.

To begin, note that the version of AIR presented at the head of section 2.1 implicitly presupposes a decision situation with an imposed time limit for accepting or rejecting an informative hypothesis. This point is illustrated by standard examples, such as industrial quality control and testing for toxic side effects of a pharmaceutical (Rudner 1953, p. 2). In both cases, a decision must be made within a restricted timeframe about whether to accept or reject a hypothesis with clear implications for a decision. And no matter which of those two options is chosen, some risk of error is present. Removing the presupposition that an informative claim must be accepted or rejected within a deadline leads to a more general version of AIR. Instead of a choice between two options, accept or reject, there are now three: accept, reject, or suspend judgment. Some simple logical relations hold among these three states: rejecting P is equivalent to accepting \neg P, and conversely, to accept P is to reject \neg P. Of greater interest for the present purposes, however, is that to suspend judgment regarding P is to neither accept P nor reject it. Consequently, one who suspends judgment runs no risk of error (i.e., of accepting P when it is false or of

4. An anonymous referee suggests the following as an example of a “practically certain, policy relevant” claim about climate change: “it is practically certain that no longer emitting GHGs from 2050 onwards would avert any catastrophe that might result from continuing to emit GHGs after 2050.” However, this claim is not informative. Naturally, if X does not occur, then X will not have harmful effects. But policy relevant claims say something about the effects of X vis-à-vis alternatives. For example, policy makers might ask: How costly would it be to reduce global net GHG emissions to zero by 2050? How effective would this be in mitigating adverse impacts of climate change in comparison to other proposals (e.g., a 90% reduction by 2050)?

rejecting P when it is true). However, error is not the only epistemic failure to which a person may be susceptible. Ignorance is another. If cognitive attitudes are limited to accept, reject, and suspend judgment, then a person can be said to be *ignorant* regarding P if and only if she is either in error or suspends judgment about P. For example, a person who accepts that Toronto is the capital of Canada is ignorant, as is one who suspends judgment about this statement.

The generalized version of AIR, then, can be formulated simply by reformulating the first two premises as follows:

- 1*. Scientists must make decisions about whether to accept, reject, or suspend judgment about hypotheses.
- 2*. Such decisions invariably involve a non-negligible risk of *ignorance*.

Since suspending judgment about P means neither accepting nor rejecting it, standards for what is sufficient for acceptance and what is sufficient for rejection jointly entail a standard for when to suspend judgment. That is, one suspends judgment when neither the evidential threshold for acceptance nor for rejection has been breached. Consequently, there is no need to explicitly mention suspension of judgment in any of the subsequent premises. The argument proceeds as before except that now an additional epistemic risk can be considered, namely, that extremely high evidential standards will result in a failure to accept or reject any informative claim. The generalized AIR shows that premises 1 and 2 in the original version are a stronger than necessary, since premise 3 follows from the weakened premises 1* and 2*. And because error (i.e., accepting a falsehood or rejecting a truth) is a species of ignorance, the generalized version of AIR retains the original as a special case.

The original AIR is driven by the simple observation that errors can be harmful. Accepting that tobacco is a health tonic when it is in fact a powerful carcinogen can plainly have negative consequences. But a generalized version of AIR requires that ignorance besides error can also be bad. It is, I think, straightforward to argue that this so. Suppose, for instance, that public health officials suspended judgment about whether tobacco smoke has harmful effects. Then they would be less likely to take appropriate actions to mitigate these harmful effects, such as placing warning labels on cigarette packages and prohibiting smoking in public places, especially as such measures would be met with vigorous political resistance. In general, the practical and ethical risk of suspending of judgment about P is that one will fail to take appropriate action that would ensue from correctly accepting or rejecting it. The generalized AIR, then, turns on the fact that extremely high evidential standards make prolonged suspension of judgment an overwhelmingly probable outcome. If scientific

advice must be limited to claims as well confirmed as “coal burns,” it is to be expected that scientists will have precious little informative advice to give. This in turn leads to serious difficulties if policy decisions are expected to be “science based” in the sense that the course of action chosen is justified by scientific knowledge. Justification in this context entails that scientific knowledge provides reasons for preferring the selected course of action to alternatives. Thus, extremely high evidential standards make it difficult to justify any course of action over others, and consequently are likely to result in prolonged delay in the response to serious problems confronting human health and the environment.

Moreover, this argument is far from hypothetical. A pair of volumes published by the European Environment Agency under the title *Late Lessons from Early Warnings* document over 25 cases environmental and public health cases illustrating this pattern (Harremoës, et al. 2002; European Environment Agency 2013; cf. Steel 2015, chapter 4). One of the lessons drawn from these cases pertains to the consequences of extremely strict evidential standards for accepting claims about harmful health or environmental effects. “The bias in science towards avoiding false positives [i.e., mistakenly accepting the existence of an effect] inevitably involves generating false negatives [i.e., failing to accept an effect when it exists], which, if they are human and or environmental disasters, as in most of these case studies, is not sound public policy” (Harremoës, et al. 2002, p. 208).

One might object to the above argument on the grounds that, if policy makers are very risk averse, then even relatively uninformative scientific advice may prompt action. For instance, scientists might report that they cannot conclusively refute the possibility of a 2-meter rise in GMSL by 2100 or the possibility of a 5-meter increase by 2200. If policy makers wish to avoid risk at all costs, this might prompt them to enact aggressive reductions of greenhouse gas emissions. However, this reasoning faces a straightforward difficulty that is commonly pointed out by critics of the precautionary principle (e.g., Manson 2002; Sunstein 2005). For scientists could be pressed to admit that they cannot conclusively refute the claim that even the most aggressive GHG reductions will fail to avert the threatened GMSL increases.⁵ Nor could they conclusively refute the possibility that such measures would have seriously harmful economic consequences.⁶ In general, that catastrophe is a possible consequence of an action is informative for decision-making only if similarly bad consequences are not

5. That significant increases in GMSL are “in the pipeline” appears to be regarded as a serious possibility by climate scientists (cf. Joughin, Smith, and Medley 2014; Rignot, Mouginot, and Scheuchl 2011).

6. For instance, see Stern (2007) and Nordhaus (2008).

possible given some alternative.⁷ Yet in the context of climate change and other complex phenomena, such impossibility claims will typically be uncertain.

Consequently, setting very high evidential bars for acceptance and rejection, as recommended by Betz, does not avoid the generalized AIR. It merely poses a distinct type of risk, namely, that scientific uncertainty may result in suspension of judgment and a consequent failure to take reasonable precautions. These concerns are also relevant to IPCC reports (cf. John 2015, pp. 7–8). For example, Hansen (2007) criticizes the IPCC for being given to “scientific reticence”—i.e., excessive epistemic caution—in regards to its projections of GMSL, and warns that such reticence could have disastrous consequences (cf. Brysse et al 2013; O’Reilly, Oreskes, Oppenheimer 2012).

3. What does AIR Aim to Show?

In this section, I consider the objection, made by Parker (2014) and Morrison (2014), that the presence of second-order uncertainty is not sufficient to show that ethical values did in fact influence modeling decisions. While I agree with Parker and Morrison on this point, I argue in section 3.1 that this observation only undermines a descriptive interpretation of AIR, and I argue that a normative interpretation is the preferable option. In section 3.2, I consider some of the implications of the normative version of AIR for climate change.

3.1. A Normative Interpretation

Even if substantial second-order uncertainty is inevitable, it might be resolved by some means besides an appeal to ethical values. Parker suggests that pragmatic considerations might guide such decisions.

For instance, the scientists might already have in hand some computer code for process P but not for processes Q, R, or S. Or they might judge that it will be much easier to incorporate P than to incorporate Q or R or S, given past choices in model building. Or they might be experts on P but have much less understanding

7. To illustrate, Betz (2009) defends a “modal falsificationist” approach (in which serious possibilities are those that have not been refuted) and ties this to the precautionary principle interpreted as the maximin rule. Yet if it is possible (i.e., not refuted) that all mitigation efforts will fail to avert catastrophe and possible that mitigation is costly, then maximin recommends *against* mitigation. That is because the worst case occurs if we mitigate and catastrophe happens anyway, since in that case we suffer the catastrophe plus the costs of mitigation (see Steel 2015, pp. 58–62).

of Q and R and S. Or it might be that a leading modeling group incorporated P for reasons like those just identified, and now it is seen as de rigeur for state-of-the-art climate models to include P. And so on. (2013, p. 27)

Margaret Morrison (2014) also argues that the presence of second-order uncertainty does not entail that values played a significant role in attempts to assess the first-order uncertainty of climate models. Subjective factors may be unavoidable, she argues, but these are not necessarily value judgments of an ethical or social nature (2014, p. 944). Perhaps scientists were unaware of the second-order uncertainty, used a particular model simply because it was the one they were familiar with, or were negligent in considering ways their model might be inadequate in a novel context (cf. Morrison 2014, pp. 256–57). To illustrate these claims, Morrison describes several cases of poor modeling decisions that resulted in catastrophes, such as the space shuttle Columbia disaster (Morrison 2014, pp. 257–58).

One might respond that even if scientists do not consciously consider ethical implications of their modeling decisions, value judgments may nevertheless be implicit in the system in which decisions are made. But turning attention to systemic aspects of scientific knowledge does not avoid the central difficulty, namely, that a thing's characteristics are not necessarily explained by its ethically significant consequences. This difficulty arises at both the level of social systems and individual decisions. Indeed, it is commonly regarded as a fallacy to assume that, if a social system S has a consequence X, then S originated, persists, or is prevalent because it has this effect (Kincaid 1996, pp. 122–36). The reasons for thinking this inference is fallacious are easy to appreciate: every social system has many consequences, but not all are relevant to explaining why it exists or has the features it does. Indeed, complex systems routinely generate surprising outcomes. Therefore, I agree with Parker and Morrison that the presence of ethically significant risks of error is not sufficient to prove an actual influence of ethical or social values.

However, what this entails for AIR depends on how that argument is interpreted. In particular, it depends on whether the argument is interpreted *normatively* (i.e., as attempting to show that values *should* influence scientific judgments about what constitutes sufficient evidence) or *descriptively* (i.e., as attempting to demonstrate an *actual* influence of values upon science).⁸ In sections 2.1 and 2.2 above, AIR is formulated in explicitly normative terms (note the presence of the word “should” in

8. This distinction is noted, for instance, in Betz (2013, p. 210).

line 6). Douglas (2009) also provides a good example of a normative interpretation of AIR. She argues that scientists have a moral responsibility to exercise reasonable foresight with respect to the consequences of error (2009, p. 67). Two things are worthy of note here. First, a claim about the moral responsibilities of scientists indicates a normative reading of AIR. Secondly, moral responsibilities worth talking about are ones that person might fail to live up to. This suggests that Douglas believes that scientists might not, and most likely sometimes do not, fulfill their moral responsibility to exercise reasonable foresight in relation to ethical concerns related to inductive risks. As a consequence, Douglas' position suggests that the presence of ethically significant consequences of error do not necessarily indicate that ethical values actually did influence decisions. Similarly, Kevin Elliott proposes the following three conditions for ethical values in "methodological and interpretative" decisions in science:

First, the "ethics" principle is that scientists have ethical responsibilities to consider the major societal consequences of their work and to take reasonable steps to mitigate harmful effects it might have. Second, the "uncertainty" principle is that those researching policy-relevant topics often face situations in which scientific information is uncertain and incomplete, and they have to decide what standard of proof to demand before drawing conclusions. Third, the "no-passing-the-buck" principle is that it is frequently socially harmful or impracticable for scientists to respond to uncertainty by withholding their judgment or providing uninterpreted data to decision makers. ... When these three principles apply to a case, there are ethical reasons for scientists to factor societal considerations into their responses to uncertainty. (2011, p. 55)

These three principles are intended to jointly characterize circumstances in which ethical values *should* influence scientific decisions.⁹ That is explicit in the final sentence in the passage: when the three conditions obtain, Elliott states that "there are ethical reasons" for values to influence decisions about what will constitute sufficient evidence.¹⁰ As before, a normative claim of

9. See Steel (2010, 2013) for additional explicitly normative renderings of AIR.

10. Winsberg also cites Elliott's three principles, but describes them as "three conditions under which scientists *should be expected to* incorporate social and ethical values" (2012, p. 130; italics added). The phrase "should be expected to" is ambiguous in a way that invites a descriptive interpretation, that is, under these conditions it is very likely that ethical values will in fact influence scientists' judgments.

this sort does not entail that values always do play this role under these conditions.

In contrast, Winsberg's statements about the inevitable role of values in climate modeling (cf. Winsberg 2012, p. 123) strongly suggest a descriptive interpretation of AIR. This reading is reinforced by Winsberg's formulation of the question he takes AIR to address: "the philosophically controversial question about social and ethical values is about the degree to which they *are* involved in the appraisal of hypotheses, or in reaching other conclusions that are internal to science" (2012, p. 114; italics added). In other words, the question is a descriptive one about the actual role of values in the assessment of hypotheses in science. Winsberg is not alone in proposing a descriptive construal of AIR, as some classic statements of the argument do the same. For instance, in Rudner's formulation, the conclusion of AIR states that, "scientists as scientists *do* make value judgments" (1953, p. 2; italics in original).¹¹ In addition, some more recent discussions of AIR also interpret it in a descriptive manner (cf. Steele 2012; Wilholt 2009).

However, the normative interpretation of AIR has two main advantages. First and most significantly, it avoids the fallacy of assuming that decisions and social systems are necessarily explained by their ethically significant consequences. To claim that, under certain circumstances, ethical considerations *should* influence evidential standards does not entail that they always do so. Indeed, the normative version of the argument is of interest precisely because ethical values sometimes fail to appropriately influence decisions about evidential standards. The second advantage of a normative interpretation of AIR is that it is more philosophically interesting insofar as it challenges the ideal of value-free science *as an ideal*. That is, it challenges the *desirability*, not just the *feasibility*, of value-free science (cf. Intemann 2015, p. 226). After all, defenders of value-free science are likely to agree that pristine value-freedom is impossible. Nevertheless, they would insist that the influence of ethical or social values upon scientific decisions about which claims to accept should be minimized as much as possible (Lacey 1999). Thus, even if an "inevitable influence" of values upon all aspects of science could be demonstrated, it would not thereby follow that the ideal of value-free science is mistaken. It is no doubt for

11. Some other classic formulations are ambiguous. For example, Hempel takes AIR to show that, "the justification of the rules of acceptance and rejection requires reference to value judgments" (1965, p. 92). Is Hempel making the descriptive claim that standards of acceptance and rejection used by scientists are in fact justified in this manner? Or is he instead making the normative point that, if one were to attempt to justify these standards, some appeal to values should be made?

such reasons that Douglas (2007) makes a point of arguing that value-free science is a bad ideal. According to Douglas, scientific communities have a responsibility to consider ethically significant foreseeable consequences of decisions concerning evidential standards, and failure to do so constitutes a form of negligence (2009, chapter 4). Moreover, a normative implication of AIR regarding climate change was already noted in section 2.2. That is, the generalized version of AIR highlights a danger associated with extremely high standards of evidence, such as those proposed by Betz, for accepting scientific claims. Namely, such standards are likely to lead to suspension of judgment, which in turn increases the likelihood of delay in the face of mounting evidence. In the case of climate change, inaction is likely to result in consequences that are progressively more severe and expensive to adapt to and mitigate. Thus, AIR entails that such factors should be considered when setting evidential standards, for instance, in IPCC assessment reports.

However, AIR itself says very little about just *how* values should influence scientific processes. The following three general types of criteria can be conceptually distinguished.

1. Procedure: Ethically significant consequences of evidential standards should be considered in ways that respect procedural norms of fairness.
2. Distribution: Evidential standards should distribute risks in ways that avoid unfairly benefiting some groups to the detriment of others.
3. Epistemic Integrity: Ethical or social values should influence standards of evidence in a manner that is compatible with the epistemic integrity of science.

All of these issues have been discussed in connection with AIR, especially the third (see Douglas 2009; Elliott 2011; Elliott and McKaughan 2014; Hicks 2014; Intemann 2015; Kitcher 2001, 2011; Shrader-Frechette 1991, 2011; Steel 2010, 2015; Steel and Whyte 2012). Nevertheless, much more work remains on clarifying all three of these and the relationships among them. In the next section, I discuss procedural fairness in relation to what I call the democratic objection to AIR.

3.2. Structure and the Democratic Objection

An explicitly normative interpretation of AIR invites objections that the role of values in science it recommends would be a bad thing. In this section, I consider one such challenge, which I call the *democratic objection*. The democratic objection asserts that ethical value judgments should not influence decisions about sufficient evidence in science because this

would undermine democracy by encouraging an elite cadre of scientists to impose their values on everyone else. A version of this argument is suggested by Katie Steele, who writes, “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” (2012, p. 893). In a similar vein, Betz writes, “As political decisions are informed by scientific findings, the value-free ideal ensures—in a democratic society—that collective goals are determined by democratically legitimized institutions, and not by a handful of experts” (2013, p. 207).

Advocates of AIR would likely respond that the democratic objection falsely assumes that scientists must make value judgments in an entirely non-transparent fashion and without consultation from non-scientists. To the contrary, defenders of AIR typically insist that the role of ethical values in decisions about what should count as sufficient evidence for accepting a claim should be as transparent as possible. That scientists should have some part in informing the value judgments does not to exclude others from participating, nor does require that values only operate behind closed doors. As Douglas writes, “Because of the need for democratic accountability in science for policy, the judgments made using these values should not be clandestine” (2009, p. 156). Moreover, AIR proponents often propose public or stakeholder participation as a means of incorporating values perspectives from outside the ranks of science (Douglas 2009; Elliott 2011; Intemann 2015; Kitcher 2011; cf. Dietz and Stern 2008). So, it might seem that friends of AIR have a quick and decisive answer to the democratic objection.

But there are several reasons why this response to the democratic objection is not sufficient. First, the majority of scientific research projects proceed without public or stakeholder involvement. Thus, if AIR recommends that individual scientists in general consider ethical values in selecting standards of evidence, then most of the time this will happen without the input of non-scientists. In addition, even when public/stakeholder participation is involved, the scientists are likely to have a predominant influence on the final product, that is, on the specific claims made and conclusions drawn in publications stemming from the research. Finally, some empirical research casts doubt on the extent to which transparency about non-scientific values or motives is an effective means for preventing problematic bias (Cain, Loewenstein, and Moore 2005; Loewenstein, Sah, and Cain 2012).

To better respond to the democratic objection, I suggest that it is necessary to explicitly consider several levels of social organization at which AIR might be relevant. The following are three organizational

levels at which ethical values might influence scientific decisions about which hypotheses to accept (cf. Steel 2015, pp. 215–16).

1. A structural framework of a system for generating and/or applying knowledge (e.g., a regulatory structure in which new pharmaceuticals must be tested found safe and effective before they can be marketed).
2. General methods or standards of evidence adopted within such a framework (e.g., randomized clinical trials as the “gold standard” for showing that a pharmaceutical is safe and effective).
3. Case specific methodological decisions made by individual scientists or research teams (e.g., choice of statistical model for analyzing a data set).

The distinction between levels 1 and 2 can be applied to the IPCC process. In this case, the structural level (i.e., level 1) would involve the overall process through which assessment reports are produced. This process begins with the IPCC Plenary, which consists of representatives of member governments, selecting a group of experts to serve on the IPCC Bureau. The IPCC Bureau then issues a call to governments and IPCC observer organizations for nominations of scientific experts for a Scoping Meeting. The Scoping Meeting produces an outline of the report, which then must be approved by the IPCC Plenary. At this stage, another round of nominations is sent out for authors for the chapters of the IPCC volumes, who will be selected by the IPCC Bureau. After the first draft of an IPCC report is written, it is submitted for expert review. One significant feature of this process is the distinct review procedures for the body of assessment reports versus the summary for policy maker (SPM). Both documents are reviewed by a large number of experts, where registration as an “expert” is very open. But unlike the main text of the assessment reports, the SPM is subject to a direct approval by representatives of governments (IPCC 2013, p. 1). However, aside from the SPM, government representatives are not permitted to directly write or modify the content of the IPCC reports. Within the IPCC process, the *Guidance Note* would be an example of level 2: it is a statement of evidential standards operating within the context of the broader structure just sketched. Such standards are conventional in that, while there are multiple sets of standards that might have been adopted, it is important that the same, or at least similar, ones be used by contributing authors (cf. Mastrandrea et al. 2010, p. 2; Willholt 2009).

Return, then, to the democratic objection. This objection maintains that the value-free ideal is necessary to prevent scientists from making ethical or social value judgments about inductive risks on behalf of the rest of society. One response to this objection is to observe that ethical values

can influence through democratic channels at levels 1 and 2. For example, decisions about reforms of structural features of a regulatory system can be openly debated in legislatures or parliaments, and input from stakeholders and scientific experts can be elicited in the process. Moreover, decisions made within an established structure (e.g., about standards of evidence for a classification of a chemical) can be subject to the approval of agency leaders appointed by elected officials. Thus, the democratic objection to AIR is much less acute in these contexts. Discussions of AIR sometimes note that values can interact with science at levels beyond that of individual scientists. For instance, Willholt (2009) discusses AIR in relation to conventions accepted within scientific communities, while Douglas (2014) emphasizes collective as well as individual responsibilities of scientists in considering inductive risks. And in response to the democratic objection, Intemann points to democratically legitimate mechanisms through which values can influence climate science (2015, pp. 228–29). However, the above tripartite division suggests a further argument that levels 1 and 2 tend to constrain the role of values at level 3.

Individual scientific researchers are constrained by the methodological standards and conventions of their fields (level 2), which in turn may be shaped by broader political and social structures in which those fields are embedded (level 1). Consequently, it is difficult for scientists to adopt methods or approaches that significantly diverge from these standards. For to do so would risk having their work ignored, for instance, by being rejected for publication or disregarded by regulators. Thus, even when established disciplinary conventions or regulatory standards are problematic, insisting that individual scientists depart from them is likely to be ineffectual. Hansen's (2007) complaints about the "scientific reticence" of the IPCC are a good illustration. Hansen would apparently prefer that the IPCC and climate scientists generally exercise less scientific reticence when communicating risks of climate change to the general public, and he appears frustrated that they have not done so. And while Hansen is free to abandon scientific reticence, if he does so, his claims are unlikely to be endorsed by IPCC assessment reports or to make the final cut of the SPM.

In explaining the "reticence" of the IPCC, Brysse and colleagues suggest, "Restraint is a community norm in science, and it tends to lead many scientists (*ceteris paribus*, with some individual exceptions) to be cautious rather than alarmist, dispassionate rather than emotional, understated rather than overstated, restrained rather than excessive, and above all, moderate rather than dramatic" (2013, p. 328). This explanation focuses on level 2, that is, norms and conventions of a scientific community operating within a broader social and political structure. But it is also plausible that level 1 factors are at play. After all, member states direct the process that

selects the IPCC Bureau and must “accept” IPCC assessment reports and “endorse” the SPM. This structure might create a felt pressure to generate reports that would not be seen as “alarmist” by influential member states, hence prompting scientists to “err on the side of least drama” (Brysse et al. 2013). Hansen also suggests that fear of being denied research funding is a motive for climate scientists exercise “scientific reticence” (Hansen 2007, p. 2).

In contrast, the democratic objection portrays scientists as autonomous agents liable to impose their values upon society unless blocked from doing so by the value-free ideal. It is questionable whether this picture is accurate. As the IPCC case illustrates, scientists as individuals are constrained by the norms of scientific communities of which they are members. And science as an institution lacks autonomy because it is dependent on external sources—especially, governments and private industry—for the funds required for research. The democratic objection is also sociologically naïve insofar as it paints policy makers as passive recipients of potentially value-tainted science. But the recent history of climate science in the United States and Canada illustrates the ability of political leaders to seek scientists whose views support their policy agendas, to defund lines of research that may produce undesired results, and to restrict communication between government scientists and the public (Harris 2014; Mooney 2005). In sum, the democratic objection rests on a doubtful picture of scientists as independent advice givers to passive government actors.

Finally, the above discussion raises the question of what role values should have at level 3, that is, at the level of methodological decisions made by individual scientists or research teams. Willholt (2009) argues against values operating at level 3, and proposes instead that decisions about how to balance inductive risks should be governed by conventions adopted by scientific communities. Otherwise, he argues, consumers of scientific research would be left to guess at the evidential standards employed by individual scientists, which would undermine trust in science (Willholt 2009, pp. 98–9). However, I think it would be a mistake to exclude applications of AIR at level 3. Not all ethically significant aspects of methodological decisions can be addressed by socially accepted conventions or standards. Moreover, there might be good reasons for challenging some widely adopted conventions and standards, so scientists should be permitted to try alternatives—for instance, that involve public or stakeholder participation—if they choose to do so. Finally, the fact that individual scientists and research teams are constrained by conventions and standards at levels 1 and 2 suggests that values operating at level 3 are unlikely to undermine democratic processes.

4. Conclusions

The aim of this paper has been to articulate and defend a generalized, normative and structural version of AIR. The interpretation of AIR presented here is generalized in that it explicitly includes the option of suspension of judgment along with acceptance and rejection, and considers costs of prolonged failure to draw any informative inference. This is relevant to explaining what is problematic about extremely strict standards of evidence, such as those advocated by Betz, for accepting claims about climate change. While both descriptive and normative formulations of AIR exist, the difference between these interpretations does not always appear to be recognized. Moreover, previous discussions have not, to my knowledge, clearly articulated the central flaw of descriptive interpretations (i.e., that decisions are not always explained by their ethically significant consequences). Finally, I suggest a three-level framework for analyzing how AIR can be applied at structural and individual levels within science. I use this framework to argue that the democratic objection to AIR, according to which value judgments by scientists would constitute an unacceptable infringement on the democratic process, rests on a socio-logically naïve picture of the relation between science and governments.

References

- Betz, G. 2007. "Probabilities in Climate Policy Advice: A Critical Comment." *Climatic Change* 85: 1–9.
- Betz, G. 2009. "What Range of Future Scenarios Should Climate Policy be Based on? Modal Falsificationism and its Limitations." *Philosophia Naturalis* 46: 133–158.
- Betz, G. 2010. "What's the Worst Case? The Methodology of Possibilistic Prediction." *Analyse und Kritik* 32: 133–158.
- Betz, G. 2013. "In Defence of the Value Free Ideal." *European Journal for Philosophy of Science* 3: 207–220.
- Betz, G. 2015. "Are Climate Models Credible Worlds? Prospects and Limitations of Possibilistic Climate Prediction." *European Journal for Philosophy of Science* 5: 191–215.
- Biddle, J., and E. Winsberg. 2010. "Value Judgments and the Estimation of Uncertainty in Climate Modeling." Pp. 172–197 in *New Waves in Philosophy of Science*. Edited by P. D. Magnus and J. Busch. Basingstoke, Hampshire: Palgrave MacMillan.
- Braithwaite, R. 1953. *Scientific Explanation*. New York: Harper & Row, Publishers.
- Brysse, K., N. Oreskes, J. O'Reilly, and M. Oppenheimer. 2013. "Climate Change Prediction: Erring on the Side of Least Drama?" *Global Environmental Change* 23: 327–337.

- Cain, D., G. Loewenstein, and D. Moore. 2005. "The Dirt on Coming Clean: Perverse Effects of Disclosing Conflicts of Interest." *Journal of Legal Studies* 34: 1–25.
- Churchman, C. 1956. "Statistics, Pragmatics, Induction." *Philosophy of Science* 15: 249–268.
- Cranor, C. 1993. *Regulating Toxic Substances: A Philosophy of Science and the Law*. Oxford: Oxford University Press.
- Dietz, T., and P. Stern. 2008. *Public Participation in Environmental Assessment and Decision Making*. Washington, DC: National Academies Press.
- Douglas, H. 2000. "Inductive Risk and Values in Science." *Philosophy of Science* 67: 559–579.
- Douglas, H. 2007. "Rejecting the Ideal of Value-Free Science." Pp. 120–139 in *Value-Free Science?* Edited by Kincaid, H., J. Dupré, and A. Wylie. Oxford: Oxford University Press.
- Douglas, H. 2009. *Science and the Value-Free Ideal*. Pittsburgh, PA: University of Pittsburgh Press.
- Douglas, H. 2014. "The Moral Terrain of Science." *Erkenntnis* 79: 961–979.
- Duhem, P. 1991. *The Aim and Structure of Physical Theory*. Princeton, NJ: Princeton University Press.
- Elliott, K. 2011. *Is a Little Pollution Good for You?* New York: Oxford University Press.
- Elliott, K. and D. McKaughan. 2014. "Nonepistemic Values and the Multiple Goals of Science." *Philosophy of Science* 81: 1–21.
- European Environment Agency. 2013. *Late Lessons from Early Warnings: Science, Precaution, Innovation*. Luxembourg: Publications Office of the European Union.
- Frigg, R., S. Bradley, H. Du, and L. Smith. 2014. "Laplace's Demon and the Adventures of His Apprentices." *Philosophy of Science* 81: 31–59.
- Frigg, R., L. Smith and D. Stainforth. 2013. "The Myopia of Imperfect Climate Models: The Case of UKCP09." *Philosophy of Science* 80: 886–897.
- Hansen, J. 2007. "Scientific Reticence and Sea Level Rise." *Environmental Research Letters* 2: 1–9.
- Harris, M. 2014. *Party of One: Stephen Harper and Canada's Radical Makeover*. Toronto: Penguin Group Canada.
- Harremoës, P., D. Gee, M. MacGarvin, A. Stirling, J. Keys, B. Wynne, and S. Guedes Vaz. 2002. *The Precautionary Principle in the 20th Century: Late Lessons from Early Warnings*. London: Earthscan Publications.
- Hicks, D. 2014. "A New Direction for Science and Values." *Synthese* 191: 3271–3295.
- Intemann, K. 2015. "Distinguishing between legitimate and illegitimate Values in Climate Modeling." *European Journal for Philosophy of Science* 5: 217–232.

- IPCC. 2013. *IPCC Factsheet: How Does the IPCC Review Process Work?* http://www.ipcc.ch/news_and_events/docs/factsheets/FS_review_process.pdf
- Jeffrey, R. 1956. "Valuation and Acceptance of Scientific Hypotheses." *Philosophy of Science* 23: 237–246.
- John, S. 2015. "The Example of the IPCC Does not Vindicate the Value Free Ideal: A Reply to Gregor Betz." *European Journal for Philosophy of Science* 5: 1–13.
- Joughin, I., B. Smith, and B. Medley. 2014. "Marine Ice Sheet Collapse Potentially Under Way for the Thwaites Glacier Basin, West Antarctica." *Science* 344: 735–738.
- Kincaid, H. 1996. *The Philosophical Foundations of the Social Sciences*. Cambridge, UK: Cambridge University Press.
- Kitcher, P. 2001. *Science, Truth, and Democracy*. Oxford: Oxford University Press.
- Kitcher, P. 2011. *Science in a Democratic Society*. Amherst, NY: Prometheus Books.
- Kuhn, T. 1970. *The Structure of Scientific Revolutions*. Chicago, IL: University of Chicago Press.
- Lacey, H. 1999. *Is Science Value-Free? Values and Scientific Understanding*. London: Routledge.
- Lakatos, I. 1970. "Falsification and the Methodology of Scientific Research Programmes." Pp. 91–196 in *Criticism and the Growth of Knowledge*. Edited by I. Lakatos and A. Musgrave. Cambridge, UK: Cambridge University Press.
- Levi, I. 1980. *The Enterprise of Knowledge*. Cambridge, MA: MIT Press.
- Liu, J., T. Dietz, S. Carpenter, M. Alberti, C. Folke, E. Moran, A. Pell, P. Deadman, T. Kratz, J. Lubchenco, E. Ostrom, Z. Ouyang, W. Provencher, C. Redman, S. Schneider, and W. Taylor. 2007. "Complexity of Coupled Human and Natural Systems." *Science* 317: 1513–1516.
- Loewenstein, G., S. Sah, and D. Cain. 2012. "The Unintended Consequences of Conflict of Interest Disclosure." *Journal of the American Medical Association* 307 (7): 669–670.
- Manson, Neil. 2002. Formulating the Precautionary Principle. *Environmental Ethics* 24: 263–274.
- Mastrandrea, M.D., Field, C.B., Stocker, T.F., Edenhofer, O., Ebi, K.L., Frame, D.J., Held, H., Kriegler, E., Mach, K.J., Matschoss, P.R., Plattner, G.-K., Yohe, G.W., Zwiers, F.W. 2010. *Guidance Note for Lead Authors of the IPCC Fifth Assessment Report on Consistent Treatment of Uncertainties*. Technical Report Intergovernmental Panel on Climate Change (IPCC).
- Mooney, C. 2005. *The Republican War on Science*. New York: Basic Books.
- Morrison, M. 2014. "Values and Uncertainty in Simulation Models." *Erkenntnis* 79: 939–959.

- Nordhaus, W. 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*. New Haven, CT: Yale University Press.
- O'Reilly, J., N. Oreskes, and M. Oppenheimer. 2012. "The Rapid Disintegration of Projections: The West Antarctic Ice Sheet and the Intergovernmental Panel on Climate Change." *Social Studies of Science* 42: 709–731.
- Parker, W. 2010a. "Predicting Weather and Climate: Uncertainty, Ensembles and Probability." *Studies in History and Philosophy of Modern Physics* 41: 263–272.
- Parker, W. 2010b. "Whose Probabilities? Predicting Climate Change with Ensembles of Models." *Philosophy of Science* 77: 985–997.
- Parker, W. 2014. "Values and Uncertainties in Climate Prediction, Revisited." *Studies in History and Philosophy of Science* 46: 24–30.
- Rahmstorf, S., G. Foster, and A. Cazenave. 2012. "Comparing Climate Projections to Observations up to 2011." *Environmental Research Letters* 7 (4): 1–5.
- Rignot, E., J. Mouginot, and B. Scheuchl. 2011. "Antarctic Grounding Line Mapping from Satellite Radar Interferometry." *Geophysical Research Letters* 37: L10504.
- Rudner, R. 1953. "The Scientist Qua Scientist Makes Value Judgments." *Philosophy of Science* 20: 1–6.
- Shrader-Frechette, K. 1991. *Risk and Rationality: Philosophical Foundations for Populist Reforms*. Berkeley: University of California Press.
- Shrader-Frechette, K. 2011. *What Will Work: Fighting Climate Change with Renewable Energy, not Nuclear Power*. New York, NY: Oxford University Press.
- Stainforth, D., M. Allen, E. Tredger, and L. Smith. 2007. "Confidence, Uncertainty and Decision-Support Relevance in Climate Predictions." *Philosophical Transactions of the Royal Society A* 365: 2145–61.
- Stanford, K. 2006. *Exceeding our Grasp: Science, History, and the Problem of Unconceived Alternatives*. Oxford: Oxford University Press.
- Steel, D. 2010. "Epistemic Values and the Argument from Inductive Risk." *Philosophy of Science* 77: 14–34.
- Steel, D. 2011. "Extrapolation, Uncertainty Factors, and the Precautionary Principle." *Studies in History and Philosophy of Biological and Biomedical Sciences* 42: 356–364.
- Steel, D. 2013. "Acceptance, Values, and Inductive Risk." *Philosophy of Science* 80: 818–828.
- Steel, D. 2015. *Philosophy and the Precautionary Principle: Science, Evidence, and Environmental Policy*. Cambridge, UK: Cambridge University Press.
- Steel, D., and K. Whyte. 2012. "Environmental Justice, Values, and Scientific Expertise." *Kennedy Institute of Ethics Journal* 22: 163–182.
- Steele, K. 2012. "The Scientist qua Policy Advisor Makes Value Judgments." *Philosophy of Science* 79: 893–904.

- Stern, N. 2007. *The Economics of Climate Change: The Stern Review*. Cambridge, UK: Cambridge University Press.
- Stocker, T., D. Qin, G. Plattner, M. Tignor, S. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. Midgley. 2013. *Climate Change 2013: The Physical Science Basis: Working Group I Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge, UK: Cambridge University Press.
- Sunstein, C. 2005. *Laws of Fear: Beyond the Precautionary Principle*. Cambridge, UK: Cambridge University Press.
- Weitzman, M. 2009. "On Modeling and Interpreting the Economics of Catastrophic Climate Change." *Review of Economics and Statistics* 91: 1–19.
- Wilholt, T. 2009. "Bias and Values in Scientific Research." *Studies in History and Philosophy of Science* 40: 92–101.
- Winsberg, E. 2010. *Science in the Age of Computer Simulation*. Chicago: Chicago University Press.
- Winsberg, E. 2012. "Values and Uncertainties in the Predictions of Global Climate Models." *Kennedy Institute of Ethics Journal* 22: 111–37.