



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

## Model uncertainty and policy choice

### Citation for published version:

Isaac, A 2014, 'Model uncertainty and policy choice: A plea for integrated subjectivism', *Studies In History and Philosophy of Science Part A*, vol. 47, pp. 42-50. <https://doi.org/10.1016/j.shpsa.2014.05.004>

### Digital Object Identifier (DOI):

[10.1016/j.shpsa.2014.05.004](https://doi.org/10.1016/j.shpsa.2014.05.004)

### Link:

[Link to publication record in Edinburgh Research Explorer](#)

### Document Version:

Peer reviewed version

### Published In:

Studies In History and Philosophy of Science Part A

### Publisher Rights Statement:

© Isaac, A. (2014). Model uncertainty and policy choice: A plea for integrated subjectivism. *Studies in History and Philosophy of Science Part A*, 47, 42-50. [10.1016/j.shpsa.2014.05.004](https://doi.org/10.1016/j.shpsa.2014.05.004)

### General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

### Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact [openaccess@ed.ac.uk](mailto:openaccess@ed.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.



# Model Uncertainty and Policy Choice: A Plea for Integrated Subjectivism

Alistair M. C. Isaac\*

May 22, 2014

## Abstract

A question at the intersection of scientific modeling and public choice is how to deal with uncertainty about model predictions. This “high-level” uncertainty is necessarily value-laden, and thus must be treated as irreducibly subjective. Nevertheless, formal methods of uncertainty analysis should still be employed for the purpose of clarifying policy debates. I argue that such debates are best informed by models which integrate objective features (which model the world) with subjective ones (modeling the policy-maker). This integrated subjectivism is illustrated with a case study from the literature on monetary policy. The paper concludes with some morals for the use of models in determining climate policy.

## 1 Introduction

A question at the intersection of scientific modeling and public choice is how to deal with uncertainty about model predictions. This “high-level” uncertainty appears to be qualitatively different from the “low-level” uncertainties which occur during model construction. Low-level uncertainties may be reduced systematically through further measurement or experimentation. In contrast, there is no widespread agreement on a systematic procedure for reducing high-level uncertainty. Yet it is model predictions which are relevant for policy choice, and thus a realistic recommendation for the use of scientific models to inform policy must address the fact of high-level uncertainty.

I argue that high-level uncertainty, as well as other values critical for policy choice, should be treated as irreducibly “subjective.” I use this term primarily to indicate properties of a subject, e.g. a scientist or policy-maker, in contrast to “objective,” or subject-independent, properties of the world. The point of emphasizing that high-level uncertainty is always subjective in this sense is to counterbalance a widespread misconception that typical representations of scientific uncertainty are insulated from contentious value judgments. When policy

---

\*University of Edinburgh, contact: a.m.c.isaac@ed.ac.uk.

decisions must be made under conditions of scientific disagreement, this misconception motivates a spurious argument from the lack of scientific consensus to the impossibility of quantifying uncertainty and applying formal decision rules. I defend decision theory in this context, in contrast with those who take such high-level uncertainties to defeat formal methods, for instance proponents of the precautionary principle.

*Integrated subjectivism* is the view that, in order to inform policy choice, a scientific model should be converted to a decision-theoretic one by supplementing it with parameters which represent relevant subjective properties of the policy-maker, e.g. her utility function, priors, or risk aversion. This strategy integrates subjective features into a model which may be interpreted as otherwise objective. I motivate this view with an example from the literature on optimal monetary policy. This example illustrates how including more of a scientific model in the decision-making process than just a probability distribution over outcomes can both constrain rational policy in novel ways and make explicit the loci of disagreement in policy debates. This possibility mitigates to some extent worries about the irreducibly subjective nature of high-level uncertainty when it plays a role in public choice. The paper concludes with some remarks on the application of these ideas to the current debate on climate policy.

## 2 High-Level Uncertainty

Not all “uncertainty” exhibits the same qualitative features. Agreement about how to represent and reduce uncertainty will differ depending upon both the source of the uncertainty and its consequences for prospective action.

For instance, the treatment of uncertainty about the parameter values used when constructing a model appears straightforward: we represent it with error bars, the standard deviation, or some other descriptive statistical technique. This “low-level” uncertainty is typically just variance in the data, and can be reduced by making more measurements, developing more precise measurement procedures, running additional experiments, etc.<sup>1</sup> Suppose two ecologists modeling the growth of the invasive cane toad (*Bufo marinus*) population in Australia disagree about the cane toad birth rate, one of the parameters in the model. Despite disagreement about parameter value, they nevertheless agree on the types of evidence relevant for resolving that disagreement, e.g. additional observations of cane toad breeding in the wild, collection of data on similar toads, experiments on cane toad breeding in various controlled test conditions, etc. Consequently, each modeler knows the actions relevant for convincing her

---

<sup>1</sup>Some parameters cannot be measured directly. Nevertheless, uncertainty about their values can still be reduced by established techniques, for instance statistical estimation. Parker (2010) argues there are cases in which the “best” value to use for a parameter may differ from our best estimate of the quantity it represents—for instance, if deviation from the parameter’s presumed “true” value corrects for imperfections in other parts of the model, improving overall performance (990). In this case, the uncertainty at issue is not about the value of the parameter *per se*, but about the rest of the model, and the pragmatic trade-offs required to improve it.

colleague, and, as evidence accrues, their views should eventually converge on a single value.

The treatment of uncertainty about the predictions of complex models appears much more problematic. This “high-level” uncertainty derives from a heterogeneous set of methodological choices made by the modeler concerning *relevance* and *idealization*. For instance, when building a model of the spread of cane toads throughout Australia, one must decide which parameters to include (predation? rainfall? pond size?) as well as the degree of spatial and temporal granularity of the model (should it partition the continent into square miles? square 100 miles? square feet?). Even the inclusion of “obviously” relevant parameters may be questioned (e.g. Kearney et al., 2008, model the future distribution of cane toads across the continent without including a parameter for the current data on cane toad location, an intuitively relevant value and one typically included in other models).

The question of which parameters are relevant for a model is a specific facet of the more general point that models typically idealize, abstract from, distort, or at least simplify the target system they are intended to represent.<sup>2</sup> Modelers who identify different features of the target system as relevant are idealizing differently, yet there is no consensus theory for evaluating the relative merits of these choices. In the context of “pure” inquiry, long-term empirical success will eventually resolve disagreements about model idealization, but in the context of policy choice, the luxury of waiting for long-term success is typically not available.

Consider, for example, two hypothetical models of cane toad territory expansion. They include some of the same parameters, representing rough geographical features of Australia, yet they calculate changes in cane toad distribution using very different methods. The first is built by a statistician and relies on analysis of trends in past toad movements to predict future cane toad distribution. The second is built by a biologist and relies on an analysis of the cane toad’s physiological and behavioral traits (leap length, daily period of activity, etc.) for generating its predictions. These two models exhibit different virtues. The first can reproduce past data, “predicting” the current distribution of cane toads when fed only their initial location upon introduction in 1935; the second is unable to reproduce this data, but has the virtue of accurately capturing the presumed mechanism of cane toad migration. If our interest in the question of cane toad distribution is purely “academic,” we may tweak and improve these models gradually as we observe actual cane toad spread.

In contrast, if our intent is to make a policy decision, say, how much to spend on culling cane toads this year, then we don’t have the luxury of waiting

---

<sup>2</sup>This is a general feature of scientific inquiry: we understand some aspect of nature by crystalizing out an efficient description of it, for instance in terms of the laws which govern it. A complete and uncondensed representation of nature would be as useless as Borges’ *map the size of the territory*. Although I focus here on models in particular in order to make contact with some specific examples, the considerations raised in this article should apply equally to the use of any aspect of scientific theory in policy decisions. (For an introduction to abstraction, idealization, and distortion in models, see for instance Morgan and Morrison, 1999, esp. Ch. 2; or Weisberg, 2013, Ch. 6.)

for long-term success. We need a method for evaluating *now* the predictions made by each model when they differ on some relevant issue, say, whether or not cane toads will reach Perth if left unchecked. Which virtue should we weigh more: success in reproducing past data or plausibility of mechanism? There is no consensus answer to this question, nor general theory for how to rank the importance of other scientific virtues such as elegance, precision, accuracy, or generality when weighing the merits of incompatible models. Although this example is artificial, it illustrates a general property of time-sensitive model evaluation: it is irreducibly value-laden.

High-level uncertainty is qualitatively different from low-level uncertainty in that there is no consensus on how to represent or reduce it. This is a consequence of the fact that model construction involves trade-offs between competing scientific values (Levins, 1966). Even when a modeling subcommunity exhibits *de facto* agreement over which trade-offs should be made, this “value consensus” does not translate into consensus on how to reduce uncertainty over model predictions. Furthermore, such localized consensus can be misleading, blinding modelers to the substantive contribution made by implicit value choices to overall uncertainty. Arguably, this is the situation in climate science, where the community has chosen to focus on large models of the physical mechanisms of climate, rather than the kind of simple, general models of relations between specific quantities popular in other sciences of complex systems, such as economics or ecology. The point is not that climate science should emulate economics, but rather that *the choice* to focus on models of a particular type is *itself* made under conditions of uncertainty, and thus a substantive, yet easily overlooked, contributor to high-level uncertainty.<sup>3</sup>

In general, scientific predictions that are relevant for public policy are typically subject to some degree of high-level uncertainty. This creates a tension in the way science is discussed in the public sphere, including the pronouncements and actions of scientists themselves when communicating with the public. Typically, policy-relevant issues are publicly discussed in binary terms (is “global warming” occurring (or not)? does lowering interest rates promote economic growth (or not)?), yet the relevant science is more appropriately framed in terms of degrees of certainty or evidential support (what is our degree of certainty in warming? what is the strength of evidence for a human contribution to warming? under what conditions do lower interest rates promote growth; what is our degree of certainty that they obtain?). A public rhetoric of bivalence obscures the nature of the scientific contribution to our knowledge of the world and undermines its effective use in policy choice.

Symptoms of this breakdown between the actualities of scientific modeling

---

<sup>3</sup>Levins (1966) argued that models can result in scientific knowledge, despite trade-offs amongst conflicting values, if they produce results which are *robust*, i.e. obtain across a variety of different simplifying assumptions (423)—in our terminology, high-level uncertainty is reduced when models that prioritize different values agree. This notion of robustness is radically stronger than the kind of robustness on which climate science has focused, i.e. invariant results across changes in parameter value, or across models based on the same commitment to physical mechanism, but with slight differences in implementation (c.f. Parker, 2011).

and public expectation can be found both on the side of scientific practice and on the side of public choice theory. In the climate science literature, for example, there has been a recent surge of interest in techniques for “propagating uncertainty” through a model, converting its low-level uncertainties into an uncertainty value over the model’s predictions. While these procedures serve an important function, they can easily be misinterpreted: a measurement of the *contribution* of low-level uncertainty to high-level uncertainty is not a measurement of high-level uncertainty *itself*. Any complete assessment of high-level uncertainty must also take into account the contributions of idealizing assumptions and corresponding trade-offs in scientific values. In the case of choice theory, a major symptom is found in the precautionary principle, which takes disagreement amongst scientists as an excuse to jettison formal decision theory. I elaborate on this point and defend my own view on how to manage high-level uncertainty for policy choice in the next section; I return to the question of uncertainty management in climate science in Section 5.

### 3 Decision Theory and Integrated Subjectivism

Formal decision theory weighs possible outcomes by their utility (conversely, loss) and the probability they obtain in order to calculate action. The most typical rule here is that of “maximizing expected utility,” or “minimizing loss”:

$$\begin{aligned} &\text{pick the action } a \in A \text{ which for outcomes } O \\ &\text{and loss function } L \text{ minimizes } \sum_{o_i \in O} P(o_i|a)L(o_i) \end{aligned}$$

In the context of the present discussion, I assume a broadly ecumenical version of formal decision theory, which takes action to be calculated *somehow* from probabilities over outcomes and the utility of those outcomes (and perhaps other values as well, for instance degree of risk aversion).

This notion of decision theory subsumes a variety of procedures in addition to maximization of expected utility. For instance, instead of taking expectation, one might take the mode or median of the distribution as most significant for guiding action; instead of assessing outcome probabilities directly, one might work with a more complex function of the probability distribution, say to represent uncertainty aversion. A classic example of such non-expectational decision theory is the prospect theory of Kahneman and Tversky (1979); for some more recent examples, see Thalos and Richardson (forthcoming) and Hahn (2013). Also subsumed under this vanilla notion are formal approaches which employ simpler functions of probabilities, for instance a minimax rule which discards probabilities over outcomes altogether, and simply picks the action which minimizes maximum loss. These diverse procedures are united by their commitment to calculating actions on the basis of a mathematical representation of the decision problem. Rules may then be evaluated in terms of their sensitivity to the mathematical structure of relevance in a given context. For instance, minimax

may make sense when choosing actions in a game against a malicious opponent, but it does not make sense when choosing actions in a natural environment of stable relative frequencies.

The value of upholding decision theory in this sense is twofold. First, I believe it accurately characterizes how humans do indeed make decisions. We value or prefer some outcomes more than others, we believe some outcomes more likely than others, and we choose how to act on the basis of these two weightings. From this descriptive perspective, the functions  $P$  and  $L$  may be thought of as “subjective”: they describe properties of the subject making the decision. Second, translating difficult decision scenarios into quantitative mathematical problems allows us to calculate rational responses when the problem at issue is too complex or unwieldy for easy intuitive analysis. In this context, one might argue that only particular functions  $P$  and  $L$  are rational, for instance, the unique  $P$  which maximizes entropy. From this normative perspective, one might interpret the functions  $P$  and  $L$  as “objective,” i.e. characterizing rational assignments with validity in the world independent of any particular agent’s actual degrees of belief while making a decision.

One may, of course, endorse the mathematical side of formal decision theory, but interpret its parameters as representing subjective properties of a particular agent (e.g. Savage, 1972). Typically, on this *subjectivist* interpretation, the normative constraints on the assignment  $P$  are purely internal, e.g. internal consistency or coherence; no attempt is made to prescribe a single distribution as rational for all agents. Subjectivism in this sense always permits the quantification of high-level uncertainty since, in the absence of consensus, a probability distribution may be generated from an individual’s personal assessment of the relative importance of competing scientific values. A prominent worry about the application of subjectivism in the context of policy choice, however, is that the probability distribution used in the decision calculation is not subject to any external, objective constraints, yet the decision being made will affect society as a whole, not just the individual who defined  $P$ . Nevertheless, a brief examination of the stifling effect of non-formal decision rules on rational debate will illustrate why even subjectively assigned probability distributions are preferable to approaches which abandon mathematical representation altogether.

As a representative case, consider one of the most prominent qualitative rules for policy choice, the precautionary principle. Although there are many formulations of the precautionary principle (Manson, 2002), it typically invokes a necessity to act if an outcome involving irreversible environmental damage is *possible* (regardless of the degree to which it is *probable*). As usually interpreted, the principle violates formal decision theory because: a) rather than weigh costs and benefits across all outcomes, it singles out a distinguished class of outcomes, those involving environmental damage; and b) the relative probabilities of these outcomes are not taken into account, their mere possibility motivates action. Typically, the precautionary principle is triggered by scientific uncertainty about the potential effects of an action and is then used to ban that action, for instance the planting of genetically modified crops (Tait, 2001) or the building of schools with exposure to radio frequency fields (Foster et al., 2000).

However, if the descriptive claim that actual agents make choices on the basis of weights across the probabilities and utilities of outcomes is correct, then invoking the precautionary principle shuts down rational debate. For example, if you and I disagree about whether or not to take a particular action  $a$ , where the potential outcome  $o$  involves environmental damage, then we either disagree about how likely the outcome is,  $P(o|a)$ , how terrible it is,  $L(o)$ , or perhaps in our respective degrees of aversion to risk or uncertainty, which might be quantified in various ways. If you invoke the precautionary principle in our argument, you obscure which of these values we actually disagree over. But this is dangerous to public discourse, because we should take different actions to resolve our disagreement in each case. A disagreement about  $P$  might be resolved by appealing to more data or further modeling, a disagreement about  $L$  will require an elucidation of our respective moral norms, and a disagreement about risk aversion might simply be irresolvable. If these quantities are explicitly represented as distinct components of the decision-making calculation, as they are in formal decision theory, then the locus of our disagreement can be clarified and effectively addressed.<sup>4</sup>

So, policy decisions should be framed in the context of formal decision theory, but the parameters of a decision-theoretic model should be open to debate. This allows for a more effective means of locating disagreement and targeting discussion than attempts to jettison decision theory in favor of value-biased qualitative guidelines such as the precautionary principle. However, this in turn implies that in the context of policy decisions informed by scientific models exhibiting high-level uncertainty, we will need to place a probability distribution somehow over possible outcomes. How can this be done?

The first point to emphasize is that, even if the model (or modelers) offers no recommendation for a probability distribution over outcomes, nevertheless, *some* distribution can be assigned. At the very least, one could assign a distribution subjectively, constrained only by norms of internal coherence. Even a purely subjectivist theory of such probability distributions is constrained by rational norms on probability change (e.g. Bayes' rule). At least in finite contexts, initially divergent probability assignments will converge on underlying "objective" probabilities when updated on enough additional evidence, mitigating the worry that their "arbitrary" character defeats rational decision making.

More importantly, however, we may be able to do even better than this. If the only aspect of our decision-theoretic model which comes from our scientific model is the probability distribution over outcomes, and that distribution was itself assigned arbitrarily due to our high-level uncertainty, we have taken very

---

<sup>4</sup>If the precautionary principle is understood as topic neutral, then it can be interpreted as an instance of a formal decision rule, e.g. maximin (Hansson, 1997), minimax regret (Bretteville, 1999), or even a form of expected utility calculation (Gollier et al., 2000). Once the principle is interpreted formally, however, it invites the kind of transparency of debate advocated here. For instance, since maximin chooses an action without reference to probabilities over states, it is plausible only for strategic interactions against a malicious opponent or under conditions of complete ignorance about nature, preconditions which may fairly be challenged in any environmental policy debate. I am indebted to an anonymous reviewer for helpful insight on this point.



little from the science. In fact, we have only taken the list of possible outcomes, leaving behind information about why these particular outcomes have been identified, or how they relate to the low-level parameters in the model. Yet this information may be relevant for our decision. If, for example, parameter  $\alpha$  has a large effect on whether a particular outcome of interest obtains, while parameter  $\beta$  has a very small effect, it is in our interest to make more careful measurements of  $\alpha$  in order to improve the quality of our decision, even if more precise values for  $\alpha$  do not reduce our overall high-level uncertainty.

The solution I explore in the following section is *integrated subjectivism*. The idea is to integrate subjective parameters relevant for decisions into a scientific model, which may otherwise be interpreted objectively, i.e. as characterizing aspects of the world independent of any agent's attitudes or values. Rather than simply summarize the model's predictions with a bare probability distribution, richer aspects of the scientific model are retained side by side with parameters for the decision-maker's loss function, risk aversion, prior beliefs, etc. The idea is to transform a scientific model into a decision-theoretic one in which objective parameters (about the world) and subjective parameters (about the agent) peacefully coexist. By integrating these subjective features into an objective model, qualitative and quantitative constraints on rational decision making may emerge that are independent of the values of subjective parameters themselves. As they apply to any rational agent, these constraints can then inform policy choice, isolating the particular subjective parameters of greatest relevance for resolving any residual disagreement.

Integrated subjectivism characterizes a general attitude toward the development of models for policy decisions, not a deterministic recipe. The question of how many subjective parameters and how much objective structure to incorporate in an integrated model should be determined pragmatically on a case by case basis. If there is widespread consensus on a model's predictions and the ranking of outcomes by loss, there is no need for an explicit integrated model. If there is consensus on the space of relevant models, though not their relative value, and risk is not a factor, then Bayesian decision theory may be applied. In this case, the "subjective" parameters are the utility function and the prior probability distribution over models, and the "objective" parameters are the likelihoods (conditional probability distributions) associated with each model. A Bayesian model such as this is thus the simplest example of an integrated subjectivist model. However, there are contexts in which scientific models do not provide determinate likelihoods and policy-makers disagree on both utilities and risk—both economic and environmental policy debates plausibly exhibit these features. In these circumstances richer integrated subjectivist models may be able to constructively inform policy decisions. The possibility of rich integrated models and examples of how they may inform policy are illustrated in the following case study from the history of economics.

## 4 An Example: Optimal Monetary Policy

This section examines a sequence of models in the literature on optimal monetary policy.<sup>5</sup> We shall see that, by supplementing models initially intended to characterize objective features of the world with subjective parameters, novel qualitative and quantitative constraints on rational policy emerged and loci of policy disagreements were made more precise.

Friedman (1953) addressed the topic of optimal monetary policy by treating it as a question of how countercyclic action could be used to increase the stability of an inherently fluctuating system. Levels of (total) income fluctuate as a function of both natural pressures and policy interventions, each of which is itself susceptible to fluctuation. This motivates the model

$$Y(t) = Q(t) + U(t), \tag{1}$$

where  $Y(t)$  is the level of total income (e.g. GDP);  $Q(t)$  is the effect of policy at time  $t$  (not an intervention made at  $t$ , but the *effect at  $t$*  of policy whenever implemented); and  $U(t)$  is the fluctuation in income due to factors other than the intervention. The model is general enough to capture any fluctuating economic quantity of interest, and has been influential in the literature on optimal policy considered generally.

Friedman’s insight was that, if monotonic trends are discounted, the stability of  $Y(t)$  can be measured by its variance,  $\sigma_Y^2$ . Then the relevant equation for assessing the stabilizing effects of policy is given by

$$\sigma_Y^2 = \sigma_Q^2 + \sigma_U^2 + 2r_{QU}\sigma_Q\sigma_U \tag{2}$$

where  $r_{QU}$  is the correlation between  $Q$  and  $U$ . This setup allows Friedman to argue for the novel result that not all countercyclic interventions are stabilizing. Of course, if  $r_{QU}$  is positive, i.e. interventions are “in the wrong direction,” they destabilize. The real surprise here is that  $r_{QU}$  may be negative, i.e. interventions push income “in the right direction” (toward the mean of  $Y(t)$ ), yet the overall effect is an increase in  $\sigma_Y^2$ , i.e. instability *increases*.

(2) is intended as a model of objective factors in the world. Friedman is concerned with our level of uncertainty about these factors, but he does not explicitly represent this uncertainty in the model. Friedman’s policy recommendations turn on his doubts about our ability to sufficiently reduce our uncertainty about those factors crucial for making effective policy decisions, such as  $\sigma_Q$  or  $r_{QU}$ . This conclusion is in line with Friedman’s general stance against overly activist / interventionist government policy. However, we can also see this stance reflected in the choices Friedman made concerning the relevance of various objective factors while building this model. For instance, his choice to include the effects of policy, but not policy actions themselves, implicitly downplays any relevance of the particulars of policy choice for his conclusions about the possibility of optimal policy in general.

---

<sup>5</sup>I previously discussed this example in Isaac (2013). This discussion expands and supplements that one, focusing here on norms for policy rather than general norms for science.

An economic tradition with radically different theoretical commitments is the targets and instruments approach of Tinbergen and Theil (e.g. Tinbergen, 1952). This approach embraces an activist attitude toward government policy, and reveals this attitude in models which explicitly relate “targets,” properties of the economy (more generally society) of interest, to “instruments,” policies or political actions. In general, many targets and instruments interact simultaneously, but the simplest model in this approach is

$$y = aP + u, \tag{3}$$

where  $y$  is the target (Friedman’s  $Y$ ),  $P$  is the policy instrument,  $a$  calibrates the effect of  $P$  on  $y$ , and  $u$  represents exogenous influences on  $y$  (Friedman’s  $U$ ). A fundamental result in this approach is the principle of *certainty equivalence*, which states that, if  $a$  is known, then optimal policy when  $u$  is unknown is equivalent to optimal policy when  $u$  is known (Simon, 1956; Theil, 1957).

Brainard (1967) begins the process of reconciling these two approaches by weakening the assumptions necessary for certainty equivalence. Beginning with (3), he considers the possibility of uncertainty about  $a$ . Although  $P$  is known to the policy-maker, if  $a$  is not, then the problem of determining optimal policy is greatly complicated. More generally, Brainard allows for the possibility of interaction between  $u$ ,  $a$ , and  $P$ , capturing the fact that changes in policy may affect the nature of exogenous influences on the target variable. For example, a change in interest rate may alter the attitudes of consumers who do not themselves take loans on the new rate, thereby changing the nature of their influence on GDP. These assumptions motivate the same transformation to (3) that Friedman performed on (1), although Brainard will put it to a different purpose:

$$\sigma_y^2 = \sigma_a^2 P^2 + \sigma_u^2 + 2r_{au}\sigma_a\sigma_u P. \tag{4}$$

This model may be interpreted as a description of either objective features of the economy, or of the policy-maker’s subjective state of uncertainty about these features. Brainard interprets the variances here as uncertainties about the values in (3), yet Friedman interpreted the variances in the mathematically analogous (2) as objective summaries of fluctuations in (1). Crucially, choice of interpretation may be made after the model has been mathematically manipulated. In particular, since (3) explicitly includes a value for the policy-maker’s action, Brainard can combine it with a loss function to derive a decision-theoretic model. Use  $\bar{x}$  to indicate the average of  $x$ ,  $E(x)$  to indicate the expectation of  $x$ , and  $L$  to indicate loss. Following common practice in statistics, Brainard assumes loss is equivalent to mean squared error:

$$E(L) = E(y - y^*)^2, \tag{5}$$

where  $y^*$  is the desired value of  $y$ . This is equivalent to

$$E(L) = (\bar{y} - y^*)^2 + \sigma_y^2. \tag{6}$$

Substituting (3) and (4) into (6) allows Brainard to derive

$$E(L) = (\bar{a}P + \bar{u} - y^*)^2 + \sigma_a^2 P^2 + \sigma_u^2 + 2r_{au}\sigma_a\sigma_u P, \tag{7}$$

which he then differentiates with respect to  $P$  in order to state the relationship between optimal policy  $P^*$  and characteristics of  $u$  and  $a$ :

$$P^* = \frac{\bar{a}(y^* - \bar{u}) - r_{au}\sigma_a\sigma_u}{\bar{a}^2 + \sigma_a^2}. \quad (8)$$

Brainard’s analysis is an instance of integrated subjectivism in the sense defined in the previous section. It unites equation (3), which represents objective aspects of the economy, equation (6), which represents the subjective beliefs and values of the policy-maker, and equation (4), which may be interpreted as either objective or subjective, within a single integrated model. From a mathematical standpoint, the integration of these equations is a matter of straightforward substitution of identities; from a conceptual standpoint, however, models unambiguously representing the world and the policy-maker have been transformed into a single, heterogeneous model, aspects of which equivocate on these two interpretations. For instance, the variables on the right of (8), which were interpreted subjectively in the derivation of (7), can be given objective empirical significance if they are interpreted as measurable properties of the previous economic cycle.

This heterogeneous model allows the derivation of some surprising qualitative and quantitative constraints on optimal policy. For instance, a counterintuitive consequence of this analysis is that it is not always optimal to act to bring  $\bar{y}$  maximally close to  $y^*$ . As the policy-maker adjusts policy on successive economic cycles, closing only part of the gap between expected and desired values of  $y$  may allow for a more efficient convergence on  $y^*$  in the long run. A particularly striking instance of this principle obtains when there is very high positive correlation between  $a$  and  $u$ . Under these conditions, it may actually be optimal in the short term to act to move  $\bar{y}$  *away* from  $y^*$ , i.e. move the expected value of  $y$  further from the desired value. Intuitively, if moving  $\bar{y}$  in the “wrong” direction reduces the variance in  $u$  (because  $r_{au}$  is positive) enough, then the overall loss may be smaller since the increase in distance between  $\bar{y}$  and  $y^*$  is counterbalanced by the reduction in  $\sigma_y^2$  (Brainard, 1967, 416–7).

Although Brainard’s model illustrates the recommendation of the previous section in the sense that it converts a scientific model into a decision-theoretic one, one might nevertheless object that it is inappropriate to call it “subjective.” The notions of “loss” and “expectation” characterized in (5) are only metaphors. “Expectation” is just average, a purely mathematical notion, while the mean squared “loss” function employed is widely agreed upon by statisticians, and chosen for mathematical convenience, not because of any contentious, subject-dependent norms. Thus, these quantities are really better understood as “objective,” i.e. matters of mathematical elegance and widespread agreement which do not depend in any irreducible way on properties of a single agent. I agree with these points to some degree, but they are mitigated by the final step in our example, which demonstrates the constructive value of incorporating irreducibly subjective risk aversion into the model.

Mitchell (1979) completes the reconciliation of the Friedman analysis with the targets and instruments approach by incorporating a new parameter into

equation (6),  $0 \leq \gamma \leq 1$ :

$$E(L) = \gamma(\bar{y} - y^*)^2 + (1 - \gamma)\sigma_y^2. \quad (9)$$

Mitchell calls  $\gamma$  the “risk aversion parameter” and states it “may be subjectively picked by the policy-maker” (913). The value of incorporating this new parameter into the model is that Friedman’s model and Brainard’s generalization of the targets and instruments approach each fall out as special cases. If  $\gamma = 1/2$ , we get Brainard’s model, and variance and deviation from the target receive equal weight in calculating the loss function. If  $\gamma = 0$ , we get Friedman’s model, where the particular target  $y^*$  is irrelevant, and the only goal of policy is to dampen cyclic activity in the economy. As  $\gamma$  grows closer to 1, the policy-maker becomes less and less “risk averse” (as she cares less and less about the predicability (variance) of the target).

The difference in the pre-theoretic commitments of Friedman’s non-interventionist approach and of the more activist targets and instruments approach is now explicitly quantized in the decision calculation. By substituting (3) and (4) into (9), Mitchell can derive qualitative and quantitative results about the relationship between  $\gamma$ ,  $P^*$ ,  $y^*$ , and  $\sigma_y^2$ . For example, as risk aversion grows in a specific (but plausible) range, the strength of its effect on  $\sigma_y^2$  also grows, implying that the more risk averse you are, the more precisely you must specify your degree of risk aversion in order to calculate optimal policy (Mitchell, 1979, 917).<sup>6</sup> Results such as this do not tell the policy-maker what degree of risk aversion to adopt, this is subjectively up to her, but they rationally constrain how that subjective parameter interacts with policy choice, for instance how precisely it must be stated in order to calculate a specific action.

The models of Brainard and Mitchell demonstrate how integrated subjectivism is possible as a strategy for generating decision-theoretic models from scientific models. Such models can be used to generate richer qualitative and quantitative constraints on policy choice than those which simply weight outcomes by a probability distribution. Furthermore, supplementing them with irreducibly subjective parameters such as risk aversion can have the effect of sharpening debate by illustrating the precise effects of contentious subjective value commitments on policy choice. This example by no means constitutes a recipe for applying integrated subjectivism in all cases. Nevertheless, it shows the potential power of the approach for both a) making use of information in a scientific model not captured by its predictions alone, and b) motivating rational constraints on policy choice by examining the interaction between objectively and subjectively interpreted parameters.

---

<sup>6</sup>I omit the full details of the argument in the interests of space, but here’s a sketch: Mitchell substitutes his expression for  $P^*$  back into (4) and differentiates with respect to  $\gamma$  in order to determine the sensitivity of the minimum  $\sigma_y^2$  on  $\gamma$ . Since the second derivative is negative when  $\bar{a} > 2\sigma_a$  and  $\gamma$  is low, the function is concave in this range and, as  $\gamma$  shrinks, minimum  $\sigma_y^2$  changes more and more rapidly.

## 5 Applications in Climate Science

The literature on the management of uncertainty in climate science is vast, and I will not pretend to survey it comprehensively here, nor to provide my own concrete model. Rather, I hope to draw from the above discussion a few small morals on some targeted points in this literature.

First, a quick overview of the nature of the problem. In order to assess long-term trends in climate change, and in particular the relationship between human activities, such as burning fossil fuels, and warming trends, there has been a concerted effort in recent decades to build large computer models of the earth's climate. Many discussions of the uncertainty in these models focus on two types of uncertainty: parametric uncertainty and structural uncertainty (e.g. Tatang et al., 1997; Webster et al., 2003; Biddle and Winsberg, 2010).<sup>7</sup> Parametric uncertainty is just low-level uncertainty: uncertainty about the value of parameters such as ice albedo, quantity of human CO<sub>2</sub> emissions in a given year, rate of oceanic heat uptake, etc. These parameters are related to each other via mathematical equations, and structural uncertainty is uncertainty about the validity of these equations. This includes both uncertainty due to the lack of a confirmed theory about the relationship between particular quantities and uncertainty due to the acknowledged use of approximation methods in the equations in order to make computer simulation feasible. This structural uncertainty constitutes only one contributor to what I have called high-level uncertainty.

A third type of uncertainty taken to be of particular interest arises due to the existence of a plurality of distinct models offering conflicting predictions about future scenarios (Parker, 2006). These models make different structural assumptions and typically have different virtues (e.g. perform differently at reproducing short-term vs. long-term past data; more realistically capture different aspects of the mechanism underlying climate change; were tuned on different data sets; etc.). Just like the conflicting statistical and biological models of cane toad migration posited in Section 2, there is no consensus method for assessing their respective merits. Furthermore this plurality appears to be *irreducible*. Lenhard and Winsberg (2010) argue that irreducible pluralism follows from our inability to analytically determine the relative contributions of different components of the model to its overall outcome. Without such a “diagnosis,” we cannot assess the relevant similarities and differences across distinct models, and thus have no basis for expecting convergence in their predictions (260). Parker (2006) identifies a number of distinct difficulties for picking a single “best” climate model: pluralism of attitudes toward simplifying assumptions (352); difficulty in assessing quality of model predictions (353); lack of consensus on a measure of overall model quality (354); no clearly superior retrodiction of past data (356). The *de facto* effect of these difficulties is climate model pluralism for the foreseeable future. For the sake of a label, call uncertainty about how to reconcile the predictions of distinct models “plurality uncertainty.”

---

<sup>7</sup>Other types of uncertainty frequently discussed include uncertainties about initial conditions, data, and boundary conditions. The first two are typically instances of low-level uncertainty, while the third typically contributes to high-level uncertainty.

Perhaps the greatest insight following from the above discussion for the climate policy debate is that plurality uncertainty is not a qualitatively distinct problem, but is continuous with structural uncertainty. Both contribute to high-level uncertainty, and were treated interchangeably when the concept was introduced in Section 2. The reason is that both types of uncertainty contribute to high-level uncertainty in precisely the same way: assessment of either structural or plurality uncertainty requires (implicit) commitment to a weighting over scientific values,<sup>8</sup> yet there is no consensus theory of how this weighting translates into uncertainty over the predictions of the model. In general, the problem of assessing how value judgments contribute to prediction uncertainty is subject to the same difficulties independent of those judgments' motivations.

This issue is obscured in the climate science literature by a relatively large *de facto* consensus amongst climate modelers to rank the virtue of plausibility of mechanism ("accuracy") over all others when assessing structural uncertainty. Here is a typical example:

Structural uncertainty in models arises because of inaccurate treatment of dynamical, physical, and chemical processes, inexact numerical schemes, and inadequate resolutions. Structural uncertainty in atmospheric models can be reduced, for example, by increasing model resolution, improving parameterization schemes, refining model dynamics, and implementing state-of-the-art numerical methods. (Tatang et al., 1997, 21,925)

But reduction of structural uncertainty by these means does not necessarily translate into a reduction in high-level uncertainty. To believe so would be to commit the map-the-size-of-the-territory mistake: to assume both that the rich detail of the target system may be approached arbitrarily closely and that this approach improves the quality of the model. Despite the *de facto* endorsement of this position by much of the climate modeling community, it cannot be supported as a general precept. Modeling, like scientific inquiry generally, *must* make trade-offs between competing values. Prioritizing "accuracy" to the exclusion of other values is insupportable in the long term, and contentious in the short term.

This becomes clear when the climate case is compared to other sciences of large-scale dynamical systems, e.g. economics, neuroscience, or ecology. In each case, models which prioritize accuracy are complemented by models which prioritize simplicity, generality, or empirical fit. A well-known example is the Lotka-Volterra predator-prey model, which effectively generates complex population dynamics from just two equations.<sup>9</sup> While it is no substitute for a more

---

<sup>8</sup>Biddle and Winsberg (2010) also conclude that proposals to reduce structural uncertainty involve value-laden decisions. Their emphasis is on the role of "non-epistemic" values in this process. I agree with their discussion, but my point is weaker, merely that some values are necessarily prioritized over others in a manner open to rational dispute.

<sup>9</sup>While this is one of the best known such models in philosophy, other examples abound: in economics, supply and demand curves; in neuroscience, idealized models of neurons, such as the McCulloch-Pitts or Hogkin-Huxley, and of learning, such as Hebb's; in physics, idealized

complex and “realistic” model of any ecosystem of interest, it nevertheless is also not replaced by such models. This is because the Lotka-Volterra model provides a concise explanation of effects that are robust across a wide array of conditions, for instance the “Volterra Property”: when both populations are subject to a general decimation, the relative decrease in predators is greater than that of prey (Weisberg and Reisman, 2008). By trading generality against accuracy, the Lotka-Volterra model increases our understanding of the invariant causal structure of predation-driven population dynamics, and with this increase in understanding comes a corresponding increase in certainty—a more “accurate” model that predicts violations of Lotka-Volterra effects should not be trusted in the absence of independent empirical confirmation, or analytic understanding of the source of the violation. In the context of decision making, the transparency and conceptual clarity of a simple model such as Lotka-Volterra may be far more useful than the opaque predictions of a large-scale, even if more “accurate,” one.

The complex role of value judgments in the construction of models ensures that straightforward computational methods for assigning a probability distribution over the outcomes of a single model do not actually reflect high-level uncertainty. Such distributions may be assigned from frequency data on large numbers of model runs (if the model is stochastic), or on model runs with slight changes to parameter values (if it is deterministic). Typically, models are too computationally demanding to allow for so many runs, hence the interest in analytic methods for propagation of uncertainties through a model. Such methods translate probability distributions over parameter values, whether measured by variance in the data (e.g. Webster et al., 2003), or by expert assessment of priors (e.g. Wigley and Raper, 2001), into a probability distribution over model outcomes. This provides a quantified assessment of the contribution to model uncertainty made by low-level uncertainty about parameter values. However, these calculations do not incorporate structural uncertainty or other value judgments contributing to high-level uncertainty. One symptom of this neglect is that possible, but “surprising” climate events, such as an abrupt change in thermohaline circulation, receive no probability mass (Wigley and Raper, 2001, 454; Webster et al., 2003, 298–9). Model construction requires choices about scientific values, and a measure of uncertainty over model predictions adequate to guide action must take these choices into account.<sup>10</sup>

These reflections should clarify why I disagree with those who take plurality uncertainty to pose a unique conceptual problem for policy choice, and in particular for the quantification of high-level uncertainty. Lenhard and Winsberg, for example, argue that one’s attitude toward the possibility of model convergence should inform one’s attitude toward the possibility of assigning a quantitative measure of uncertainty (QMU) over outcomes. They are worried that pragmatists about the assignment of QMUs are not taking the possibility

---

models of springs and oscillators, the Ising model of phase transitions; etc. In all these cases, more “realistic” or “accurate” models of the system of interest do not replace, or undermine the importance of, the simpler model for purposes of explanation, conceptual clarification, and the prediction and understanding of novel results.

<sup>10</sup>See Tebaldi and Knutti (2007, 2054–5) for a similar argument.



that convergence will never obtain (i.e. irreducible pluralism) seriously enough:

But the convergence skeptic worries that [the] QMU pragmatist is overly optimistic about making constructive use of the plurality of models without a deeper understanding of the sources of model variation: QMU methods embed climate simulations into a framework of probabilistic modeling the justification of which is itself uncertain. Inevitably, some relatively strong assumptions have to go into it. (Lenhard and Winsberg, 2010, 260)

But the assignment of probabilities to outcomes from *single* models faces exactly the same problem: strong assumptions will have to be made about scientific value for which there is no ultimate basis in understanding or unimpeachable justification.<sup>11</sup> True, my motivations for supporting the assignment of probabilities are pragmatic: as argued in Section 3, such probabilities are both necessary for rational decision making about climate policy and provide an explicit representation of one locus of possible disagreement during debate. But the suggestion that instead of assigning probabilities one should simply “accept model pluralism as . . . useful information for the decision process” (Lenhard and Winsberg, 2010, 261) *simpliciter* offers policy-makers *less* rather than more of the information available from modeling practice. When policy is decided, outcomes will *somehow* be weighted probabilistically: *if the modelers don't do it, then the politicians will* (c.f. Schneider, 2001; Frame et al., 2007).

Nevertheless, I do support the conclusion that *more* than just a probability distribution over outcomes should be communicated to policy-makers, if possible. But is integrated subjectivism of the form discussed in Section 4 possible in climate science? There is already a concerted effort to integrate economic models with climate models (so-called “Integrated Assessment Models”), and the importance of explicitly representing the subjective values of policy-makers in these models has been defended by Morgan and Dowlatabadi (1996). Schneider (1997) endorses this conclusion and advocates an active role for policy-makers in collaborating with modelers to understand the limits of integrated models while providing feedback on those subjective values which should be included to maximize their efficacy for policy choice. Lempert et al. (2003) articulate an approach employing simple simulation models and weighting plausible policy outcomes with various standardized measures of value, such as the UN’s “Human Development Index.” Their strategy attempts to postpone disagreements over long-term outcome probabilities as long as possible by searching for policies that are satisfactory across a wide array of near-term possibilities.<sup>12</sup>

Typical Integrated Assessment Models are too large and complex to be susceptible to the kind of analytic results discussed in Section 4. Nevertheless, there are some simple models that support analytic conclusions of relevance,

---

<sup>11</sup>c.f. Parker’s (2010) discussion of norms for assigning probabilities under plurality uncertainty; it applies *mutatis mutandis* to high-level uncertainty about individual models.

<sup>12</sup>Although Lempert et al., 2003, explicitly reject a narrow version of “traditional decision theory,” their own approach is still quantitative, and thus subsumed under the ecumenical notion of decision theory adopted in Section 3.

although these tend to reorient debate back onto precise outcome probabilities (e.g. Weitzman’s (2009) analysis of the role of low probability “catastrophes” in expected utility calculations; or Gollier, et al.’s (2000) analysis of the relation between risk aversion and prudent action when information about risk severity changes over time). A successful analysis of the kind discussed above for monetary policy would need to link the most contentious quantities for climate policy (precise probabilities and utilities for low probability events) to more easily measurable (subjective or objective) quantities, though how this should be done is by no means antecedently clear.

In summary, we may draw the following morals for the use of climate models in policy debates:

1. When there are disagreements over utility or risk, uncertainty over climate model outcomes *must* be quantified if policy debate is to remain rational.
2. No mechanical method for quantifying high-level uncertainty on the basis of low-level uncertainties *alone* will ever provide a value-neutral measure of uncertainty.
3. Pluralism uncertainty is not qualitatively distinct from structural uncertainty. This is not to say pluralism uncertainty is not an important topic, rather that the conceptual issues made transparent in the consideration of pluralism uncertainty apply equally in the assessment of high-level uncertainty for a single model.
4. Models that integrate subjective features of the policy-maker with objective features of the climate system have the potential to address the worry that a mere probability distribution over outcomes cannot capture the nuances of scientific uncertainty about climate models.
5. Both Bayesian approaches to model pluralism and Integrated Assessment Models that explicitly represent policy-maker values constitute instances of integrated subjectivism. If analytic results of the kind discussed for optimal monetary policy are to be found, however, these models must be substantially simplified (i.e. trade-offs must be made against the value of “accuracy”).

## 6 Conclusion: A Plea

While we should indeed strive to find the firmest objective basis for the beliefs and values that inform our policy decisions, this is not always possible. Some relevant values, such as our aversion to risk, are irreducibly subjective; others, such as our uncertainty about the future, may require subjective assessment when action is urgent. Acknowledging that available evidence may rest on subjectively assigned values should not be an excuse to abandon decision theory, nor should it forestall the use of science in informing those values, e.g. by assigning probabilities, even when the norms for how it should do so are contentious. In

cases such as this, we should strive to integrate explicit representations of these subjective values with those quantities that greater scientific consensus allows us to treat objectively. Such integrated subjectivism can inform and clarify debate, as well as suggest a path for moving toward rational consensus when science provides less certain knowledge of the world than we desire.

## Acknowledgements

This paper was inspired by a talk given by Wendy Parker at the University of Pennsylvania in March, 2012. Earlier versions were presented at the MCMP's Models and Decisions Conference, April, 2013, and the University of Pennsylvania's Spring HPS Workshop, May, 2013. It has benefitted greatly from conversations with conference participants, particularly a long chat with Jeffrey Barrett during a particularly pleasant spring stroll through Munich, as well as from the comments of three anonymous reviewers. This research was supported in part by NSF grant no. 1028130.

## References

- Biddle, J. and Winsberg, E. (2010). Value judgments and the estimation of uncertainty in climate modeling. In Magnus, P. D. and Busch, J., editors, *New Waves in Philosophy of Science*, pages 172–197. Palgrave Macmillan.
- Brainard, W. (1967). Uncertainty and the effectiveness of policy. *The American Economic Review*, 57(2):411–425.
- Bretteville, C. (1999). Decision Criteria Under Uncertainty and the Climate Problem. Technical report, University of Oslo. Working Paper 10.
- Foster, K. R., Vecchia, P., and Repacholi, M. H. (2000). Science and the precautionary principle. *Science*, 288:979–981.
- Frame, D. J., Faull, N. E., Joshi, M. M., and Allen, M. R. (2007). Probabilistic climate forecasts and inductive problems. *Philosophical Transactions of the Royal Society A*, 365:1971–1992.
- Friedman, M. (1953). The effects of a full-employment policy on economic stability: A formal analysis. In *Essays In Positive Economics*, pages 117–132. University of Chicago Press.
- Gollier, C., Jullien, B., and Treich, N. (2000). Scientific progress and irreversibility: An economic interpretation of the ‘precautionary principle’. *Journal of Public Economics*, 75:229–253.
- Hahn, U. (2013). Modeling human decision making: Some puzzles. Keynote at MST2013: Models and Decisions, Munich; April 11, 2013.
- Hansson, S. O. (1997). The limits of precaution. *Foundations of Science*, 2:293–306.

- Isaac, A. M. C. (2013). Modeling without representation. *Synthese*, 190:3611–3623.
- Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292.
- Kearney, M., Phillips, B. L., Tracy, C. R., Christian, K. A., Betts, G., and Porter, W. P. (2008). Modelling species distributions without using species distributions: The cane toad in Australia under current and future climates. *Ecography*, 31:423–434.
- Lempert, R. J., Popper, S. W., and Bankes, S. C. (2003). *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis*. RAND.
- Lenhard, J. and Winsberg, E. (2010). Holism, entrenchment, and the future of climate model pluralism. *Studies in History and Philosophy of Modern Physics*, 41:253–262.
- Levins, R. (1966). The strategy of model building in population biology. *American Scientist*, 54(4):421–431.
- Manson, N. A. (2002). Formulating the precautionary principle. *Environmental Ethics*, 24(3):263–274.
- Mitchell, D. W. (1979). Risk aversion and optimal macro policy. *The Economic Journal*, 89:913–918.
- Morgan, M. G. and Dowlatabadi, H. (1996). Learning from integrated assessment of climate change. *Climatic Change*, 34:337–368.
- Morgan, M. S. and Morrison, M., editors (1999). *Models as Mediators: Perspectives on Natural and Social Science*. Cambridge UP.
- Parker, W. S. (2006). Understanding pluralism in climate modeling. *Foundations of Science*, 11:349–368.
- Parker, W. S. (2010). Whose probabilities? predicting climate change with ensembles of models. *Philosophy of Science*, 77(5):985–997.
- Parker, W. S. (2011). When climate models agree: The significance of robust model predictions. *Philosophy of Science*, 78:579–600.
- Savage, L. J. (1972). *The Foundations of Statistics*. Dover, 2nd edition.
- Schneider, S. H. (1997). Integrated assessment modeling of global climate change: Transparent rational tool for policy making or opaque screen hiding value-laden assumptions? *Environment Modeling and Assessment*, 2:229–249.
- Schneider, S. H. (2001). What is ‘dangerous’ climate change? *Nature*, 411:17–19.

- Simon, H. A. (1956). Dynamic programming under uncertainty with a quadratic criterion function. *Econometrica*, 24(1):74–81.
- Tait, J. (2001). More Faust than Frankenstein: The European debate about the precautionary principle and risk regulation for genetically modified crops. *Journal of Risk Research*, 4(2):175–189.
- Tatang, M. A., Pan, W., Prinn, R. G., and McRae, G. J. (1997). An efficient method for parametric uncertainty analysis of numerical geophysical models. *Journal of Geophysical Research*, 102(D18):21,925–21,932.
- Tebaldi, C. and Knutti, R. (2007). The use of the multi-model ensemble in probabilistic climate projections. *Philosophical Transactions of the Royal Society A*, 365(1857):2053–2075.
- Thalos, M. and Richardson, O. (forthcoming). Capitalization in the St. Petersburg game: Why statistical distributions matter. *Politics, Philosophy & Economics*.
- Theil, H. (1957). A note on certainty equivalence in dynamic planning. *Econometrica*, 25(2):346–349.
- Tinbergen, J. (1952). *On the Theory of Economic Policy*. North-Holland.
- Webster, M., Forest, C., Reilly, J., Babiker, M., Kicklighter, D., Mayer, M., Prinn, R., Sarofim, M., Sokolov, A., Stone, P., and Wang, C. (2003). Uncertainty analysis of climate change and policy response. *Climate Change*, 61:295–320.
- Weisberg, M. (2013). *Simulation and Similarity: Using Models to Understand the World*. Oxford UP.
- Weisberg, M. and Reisman, K. (2008). The robust Volterra principle. *Philosophy of Science*, 75:106–131.
- Weitzman, M. L. (2009). On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics*, 91:1–19.
- Wigley, T. M. L. and Raper, S. C. B. (2001). Interpretations of high projections for global-mean warming. *Science*, 293:451–454.