Studies in History and Philosophy of Science 47 (2014) 42–50

Contents lists available at ScienceDirect

Studies in History and Philosophy of Science

journal homepage: www.elsevier.com/locate/shpsa

Model uncertainty and policy choice: A plea for integrated subjectivism

Alistair M. C. Isaac

School of Philosophy, Psychology, and Language Sciences, The University of Edinburgh, Dugald Stewart Building, 3 Charles Street, Edinburgh EH8 9AD, United Kingdom

A R T I C L E I N F O

Article history:
Received 19 December 2013
Received in revised form
22 May 2014
Available online 9 July 2014

Keywords:
Uncertainty;
Modeling;
Public policy;
Economics;
Decision theory;
Climate science

A B S T R A C T

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.

When citing this paper, please use the full journal title *Studies in History and Philosophy of Science*

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 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

E-mail address: a.m.c.isaac@ed.ac.uk.

http://dx.doi.org/10.1016/j.shpsa.2014.05.004
0039-3681/

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.[1] 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 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 heterogenous 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.[2] 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 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.[3]

[1] 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.

[2] This is a general feature of scientific inquiry: we understand some aspect of nature by crystallizing 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 & Morrison, 1999, esp. chap. 2; or Weisberg, 2013, chap. 6.)

[3] 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).