

Nonepistemic Values and the Multiple Goals of Science

Kevin C. Elliott and Daniel J. McKaughan*†

Recent efforts to argue that nonepistemic values have a legitimate role to play in assessing scientific models, theories, and hypotheses typically either reject the distinction between epistemic and nonepistemic values or incorporate nonepistemic values only as a secondary consideration for resolving epistemic uncertainty. Given that scientific representations can legitimately be evaluated not only based on their fit with the world but also with respect to their fit with the needs of their users, we show in two case studies that nonepistemic values can play a legitimate role as factors that override epistemic considerations in assessing scientific representations for practical purposes.

1. Introduction. It is widely recognized that nonepistemic values have a legitimate role to play in many aspects of scientific reasoning, including the choice of research projects and the application of scientific results (Lacey 2005; Elliott 2011). However, it is much less clear that nonepistemic values have a legitimate role to play in assessing scientific models, theories, or hypotheses. Recent efforts to argue that they do have a role to play typically either reject the distinction between epistemic and nonepistemic values en-

Received February 2013; revised August 2013.

*To contact the authors, please write to: Kevin C. Elliott, Lyman-Briggs College and Department of Fisheries and Wildlife, Michigan State University, East Lansing, MI 48825; e-mail: kce@msu.edu.

†We thank Martin Carrier, Heather Douglas, Ashley Graham Kennedy, Hugh Lacey, Angela Potochnik, and three anonymous reviewers for very helpful comments on previous versions of this article. We also benefited from discussions with the participants at the “Cognitive Attitudes and Values in Science” conference at the University of Notre Dame, the “Society for Philosophy of Science in Practice” conference in Toronto, and the “State’s Stake in Science” conference at the University of Bielefeld. This article was made possible through the support of a grant provided to Daniel J. McKaughan from The Faraday Institute for Science and Religion at St. Edmund’s College, Cambridge, and the Templeton World Charity Foundation as well as by sabbatical leave provided to Kevin C. Elliott by the University of South Carolina.

Philosophy of Science, 81 (January 2014) pp. 1–21. 0031-8248/2014/8101-0005\$10.00
Copyright 2014 by the Philosophy of Science Association. All rights reserved.

tirely (see, e.g., Rooney 1992; Longino 1996) or incorporate nonepistemic values only as a secondary consideration for resolving epistemic uncertainty (Howard 2006; Steel 2010).¹ Critics of the distinction between epistemic and nonepistemic values note that it is unclear whether traditional epistemic values such as simplicity are purely epistemic or whether they also incorporate nonepistemic considerations (Douglas 2009, 90; Steel 2010). Moreover, they argue that some values that are not traditionally regarded as epistemic (e.g., novelty, applicability, and ontological heterogeneity) may in fact serve as alternative epistemic values (Longino 1996). On this basis, Longino (2002) argues that efforts to maintain scientific objectivity should focus on scrutinizing and criticizing values rather than trying to exclude any particular category of values completely.

Those philosophers of science who still attempt to distinguish epistemic from nonepistemic values typically allow nonepistemic values to influence the assessment of models, theories, and hypotheses only as a secondary consideration when epistemic values leave room for uncertainty (for more discussion, see Brown [2013]). There are different ways of conceptualizing this secondary role for nonepistemic values. Heather Douglas (2009) distinguishes epistemic criteria such as predictive competence and internal consistency from the range of other values that can influence scientific reasoning. She argues that other values should not directly serve as reasons for accepting or rejecting hypotheses, but they can act indirectly to influence the standards of evidence that scientists demand when responding to uncertainty. Daniel Steel and Kyle Powys Whyte (2012) argue that Douglas's distinction between direct and indirect roles does not provide reliable guidance for identifying appropriate and inappropriate influences of nonepistemic values in science (see also Elliott 2013). Instead, they argue for a "values-in-science" standard, according to which nonepistemic values should play the role of "tiebreakers" when two conclusions or methodological approaches are equally well supported by epistemic values (Steel 2010; Steel and Whyte 2012). This position has a long and respectable tradition of supporters. As Don Howard (2006) has argued, Pierre Duhem and Otto Neurath also limited

1. For the purposes of this study we take epistemic values to be considerations that are relevant to assessing the extent to which a representation matches the world; they are criteria for theory choice that bear on the likely truth or falsity of a theory—or some other semantic relation of interest such as empirical adequacy (McKaughan 2007). Fit with the empirical data and self-consistency are clear examples of epistemic values, if anything is, whereas the fact that a particular theory flatters our social or political interests often will not be. We should note that we are sympathetic to Justin Biddle's (2013) argument that there are disadvantages to using the term 'values' to refer to the myriad of contextual factors that can influence scientific reasoning. Nevertheless, we have chosen to retain the language of epistemic and nonepistemic values throughout this article for the sake of maintaining terminological continuity with previous literature in this field.

the role of nonepistemic considerations (such as a theory's conduciveness to the achievement of particular social and political ends) to a moment in theory choice after logic and experience have narrowed the range of viable alternatives as much as they can. One of us (Elliott 2011) has recently proposed a slightly different version of this position, according to which nonepistemic values have a legitimate role to play in assessing scientific theories or hypotheses when three conditions are met: (1) epistemic values are insufficient to determine a decision, (2) it would be problematic for scientists to suspend their judgment, and (3) they have ethical reasons for incorporating nonepistemic considerations in making their decision.

While there are differences between these approaches, they all seem to assume that epistemic values should play a privileged role in scientific theory assessment, as long as they can be distinguished from nonepistemic values. At first glance, this is a reasonable assumption, considering that science is typically conceived as a search for true claims about the world, and nonepistemic values are by definition not directly relevant to this enterprise. However, in section 2 we argue that reading a single goal off scientific activity is more complicated than it first appears. As the analyses of scientific representation by Ron Giere (2004, 2006) and Bas van Fraassen (2008) show, representations can be evaluated not only on the basis of the relations that they bear to the world but also in connection with the various uses to which they are put. To the extent that scientists need to evaluate their representations along both dimensions, estimates about the likely truth of a model or theory may be only one of several considerations that factor into decisions about its acceptance. In sections 3 and 4, we illustrate the resulting roles for nonepistemic values in two case studies. Section 5 responds to several objections against our thesis and further analyzes how the roles of epistemic and nonepistemic values can be balanced and evaluated in specific cases.

2. Trade-offs and the Multiple Goals of Science. The complex judgments involved in assessing scientific hypotheses, theories, and models—whether with respect to epistemic or practical goals—often involve weighing the relative importance of a range of considerations, which can sometimes stand in tension. Kuhn famously called our attention—in “Objectivity, Value Judgment, and Theory Choice” (1977)—to cases in which rival theories exemplify different (epistemic) desiderata to varying extents, leaving scientists to decide which values should be given the most weight in a particular context. A very similar point has been made in recent literature on modeling. Scientists constructing a particular model will often find that, other things being equal, increasing one desirable feature of a model, such as precision, compromises another, such as generality or breadth of applicability (Matthewson and Weisberg 2009; Potochnik 2012). But Kuhn's talk of ‘theory choice’

leaves unclear what sort of choice we are being asked to make when deciding between various theories or models. This is unfortunate, because questions about how best to balance trade-offs between various desiderata clearly depend on what our goals are when we make our choices.

One's purposes in making a given calculation may affect, for example, whether treating a plane as frictionless or a star as a point mass will yield approximations useful for the task at hand. As Michael Weisberg points out, it is not uncommon for scientists to rely on different and even incompatible models in a variety of domains:

In ecology, for example, one finds theorists constructing multiple models of phenomena such as predation, each of which contains different idealizing assumptions, approximations, and simplifications. Chemists continue to rely on both the molecular orbital and valence bond models of chemical bonding, which make different, incompatible assumptions. In a dramatic example of MMI [Multiple-Models Idealization, the practice of building multiple incompatible models, each of which makes distinct causal claims about the nature and causal structure giving rise to a phenomenon], the United States National Weather Service employs three complex models of global circulation patterns to model the weather. Each of these models contains different idealizing assumptions about the basic physical processes involved in weather formation. (Weisberg 2007)

In our view, one needs to recognize a place for the aims of the user in order to make sense of why such trade-offs get balanced the way that they do in practice.

Indeed, several of the best recent attempts at developing a general account of the nature of scientific representation, by Ron Giere (2004, 2006) and Bas van Fraassen (2008), have called attention to the importance of explicitly incorporating a role for agents or users (as well as their goals and purposes) as a crucial component of any adequate analysis. For example, Giere describes modeling practices in science using this schema: "Scientists use models to represent aspects of the world for specific purposes" (Giere 2004, 742; see also Giere 2006, 60; van Fraassen 2008, 21). According to this schema, the representational success of models can be evaluated not only in terms of their fit with the world but also in terms of their suitability to the needs and goals of their users. We can ask questions not just about the semantic relations between a representation and the world (e.g., "Do the theoretical terms of the model successfully refer to entities that actually exist?" or "Is what the hypothesis says about the world correct?" or "Is this theory true or at least empirically adequate?" or "How accurately does it represent the world?") but also pragmatic questions about the relations between a

representation and its users (e.g., “Is it easy enough to use this model?” or “Is this hypothesis accurate enough for our present purposes?” or “Can this theory provide results in a timely fashion?” or “Is this model relatively inexpensive to use?”).

Both Giere and van Fraassen develop analyses that apply to scientific representations in general, including theories, hypotheses, and models.² Any object or proposition that is used to represent something else can be analyzed both with respect to its fit with the object to be represented and with respect to its fit with the practical purposes for which it is used. As an example of the role that practical considerations can play alongside epistemic ones in decisions about the construction and choice of representations, consider mapmaking (see Kitcher 2001). A commuter rail map might be designed to convey information about the order of stops along the line without any pretense to accurate representation of relative distances or scale. Notice that a question like “Which map should I choose?” is a goal-dependent question. The commuter rail map will be clearly inferior to other representations of the same territory for many other purposes. Nonetheless, provided that this map fits the world closely enough in relevant respects to help us to get to our destination successfully, practical qualities such as being easy to understand and simple to use provide good reasons for relying on it.

The upshot of thinking more carefully about the multiple goals that scientists have when choosing scientific representations is that it helps us to understand how scientists can sensibly prioritize nonepistemic considerations over epistemic ones in some cases. Scientists need not always maximize the fit between a model and the world; rather, the purposes of the users determine what sort of fit with the world (and therefore what balance between epistemic and nonepistemic considerations) is needed in particular contexts. Scientists use models and theories to represent the world for specific purposes, and if they can serve those purposes best by sacrificing some epistemic features for the sake of nonepistemic ones, it is entirely legitimate for them to do so. For example, it may be easier to achieve certain practical goals if scientists adopt a model or hypothesis that posits less

2. In some of his work, Giere distinguishes hypotheses from models and argues that hypotheses make claims about the specific respects in which models correspond to particular features of the world (see, e.g., Giere 2004). On this view a model is not by itself true or false; what is true or false are the hypotheses that specify in what respects the model is intended to represent particular features of the world. Even when the terminology is regimented in this way, one can still ask questions about how well the hypotheses fit the world as well as how well they achieve practical purposes (e.g., how long does it take to test them, or how easy are they to communicate to others?). Moreover, most figures do not interpret models and hypotheses in such precise ways; they often blur together various representational entities such as theories, models, and hypotheses.

realistic entities but that is easier to use. Of course, the fact that the practical assessment of a model might include aims other than representing the world as accurately as possible need not preclude or compromise efforts to assess the model solely from an epistemic perspective as well.³ But we think that practical assessments of representations play a very important role in actual scientific practice, and thus an adequate account of the roles for values in science needs to take account of these sorts of assessments.

One might object to the practice of allowing pragmatic or nonepistemic considerations to trump epistemic ones (or to the idea of evaluating models not solely based on their fit with the world but also based on their fit with various needs of the models' users) by arguing that this violates the ultimately epistemic goals of science. We offer two points in response. First, it seems relatively clear that, in current scientific practice, the choice of models and theories is governed by a range of practical goals in addition to epistemic ones. It is a descriptive fact of ordinary scientific practice that models represent their targets with varying degrees of success and typically focus selectively on those factors that are necessary in order to achieve the purposes for which they are used. Even our most successful models of nature are often known to be partial, simplified, incomplete, or only approximate. Indeed, the use of idealization, which goes beyond mere abstraction by deliberately employing assumptions known not to be true of the system of interest (e.g., treating bodies as point masses, surfaces as frictionless planes, collisions as perfectly elastic, nonisolable systems as isolated systems) is a pervasive part of model-building methodologies. In some cases, these simplifications or idealizations could assist in achieving epistemic goals (e.g., obtaining more accurate predictions). But scientists also routinely use models that incorporate a range of simplifications for the sake of greater computational efficiency or tractability or for other pragmatic reasons (such as ease of use) that are distinct from epistemic considerations.

Second, worries that strict epistemic assessments of models would somehow be compromised by allowing a role for nonepistemic considerations can be allayed by considering the wide array of cognitive attitudes that scientists can and do adopt toward models and scientific representations more generally (see e.g., McKaughan 2007; Elliott and Willmes 2013). For example, in lieu of believing a theory to be true (perhaps to a higher or lesser degree) one could simply entertain a theory, or one could accept it as worthy of further pursuit, or one could accept it as a basis for policy making, or one could combine such attitudes with other sorts of beliefs about why reliance on

3. We are assuming for the sake of argument that it is indeed possible to perform a purely epistemic assessment of scientific representations. Some philosophers might be skeptical of this possibility. But if one concluded that it is not possible to perform such purely epistemic assessments, it would merely strengthen the main argument of our study.

the theory is useful for realizing one's practical goals in a given context.⁴ As long as scientists are careful to adopt appropriate cognitive attitudes toward their representations and are explicit about their reasons for relying on them in a given context, allowing nonepistemic considerations to factor into one's reasons for acceptance (or 'theory choice') need not preclude strictly epistemic assessments of the representation. Nothing we have said precludes one's epistemic opinion (e.g., beliefs) from being strictly oriented toward anything besides truth. Nor does it require people to violate any plausible epistemic norms. When engineers working on the space program at NASA employ Newtonian mechanics to predict the trajectory of a rocket, this surely does not imply that they believe that Newtonian mechanics provides a true or even empirically adequate account of all physical phenomena. Rather, some more qualified belief that "it's good enough for government work" is in play: they use it with the full awareness that the theory is false, while believing that the predicted values are close enough for the purposes at hand.

In order to make our argument in this section more concrete, we turn in the following sections to two examples of the ways that nonepistemic values can take priority over epistemic ones when assessing scientific representations for practical purposes. Because there is so much room for disagreement about which values are epistemic and which ones are nonepistemic, we have tried to choose examples where the contrast between these two sorts of values is particularly stark. The cases discussed in the following sections focus on conflicts between the epistemic value of accurate prediction versus nonepistemic values such as ease of use or speed of generating results. We do not mean to suggest, however, that epistemic values always require providing maximal predictive details about a phenomenon. In some cases scientists might weigh various epistemic values and conclude that they should prioritize explanatory power over predictive accuracy. Our point is that even in those cases where other epistemic values are prioritized, the reasoning that we highlight in this study justifies choosing a different model or theory than would be chosen based solely on weighing all epistemic considerations.

3. Expedited Risk Assessments. For more than 30 years, risk assessment has played a central role in government approaches to regulating toxic chemicals. Traditionally, it has involved several science-intensive stages, in which researchers identify hazardous substances, develop dose-response curves for their toxic effects, create exposure assessments, and then characterize the resulting risks (NRC 1983). This process is highly time consuming and can be strategically manipulated by regulated entities to slow

4. As we define the cognitive attitude of acceptance, to accept a hypothesis, h , is to presuppose h for specific reasons as a premise in one's deliberations (see Cohen 1992; McKaughan 2007; Elliott and Willmes 2013).

the regulatory process to a standstill (Michaels 2008; Cranor 2011). In response to these difficulties, Carl Cranor (1995) has argued that the social costs of relying on current risk-assessment procedures (which are fairly accurate but very slow) are generally greater than they would be if regulatory agencies relied on less accurate but much quicker methodologies for assessing risks. Thus, his analysis shows that when scientists are accepting a methodological approach or a model for the purposes of guiding regulatory policy in ways that minimize social costs, they sometimes have to sacrifice epistemic values such as accuracy for the sake of nonepistemic values such as the ability to generate rapid conclusions.

The reason why conventional approaches to risk assessment generate such high social costs is that they are too slow to keep up with information about likely carcinogens. As Cranor (1995) explains, once an entity such as the International Agency for Research on Cancer (IARC) or the National Toxicology Program (NTP) identifies a substance as a carcinogen in animals, government agencies have to assess the severity of the risks that they pose in order to develop regulations. For chemicals other than pesticides and pharmaceuticals, the manufacturers do not have to show that their products are safe before putting them on the market in the United States; rather, the burden of proof is on the government to show via risk assessments that the chemicals ought to be restricted or removed from the market. Therefore, because conventional risk assessment methodologies are so time and labor intensive, only 20% or 30% of the animal carcinogens identified by the IARC or the NTP have actually received risk assessments and been regulated. This massive underregulation of likely carcinogens is very costly to society in terms of human and environmental health effects. However, there would also be significant costs to the economy if policy makers attempted to regulate likely carcinogens without collecting information about the severity of the risks that they posed.

The beauty of Cranor's (1995) analysis is that he examines in a precise way the trade-offs between more and less accurate modeling approaches for assessing risks. He notes that in the early 1990s the California Environmental Protection Agency (CEPA) used an expedited risk assessment methodology to estimate the carcinogenic potency of 200 chemicals in an 8-month period, whereas they had been able to analyze only 70 chemicals in 5 years using traditional methodologies. However, the expedited approach was slightly less accurate, yielding roughly 2.7% more major overregulation (which was defined as estimates of potency more than 25-fold higher than traditional approaches), 12% more minor overregulation (defined as estimates of potency between 5- and 25-fold higher than traditional approaches), and 5% more minor underregulation (defined as estimates of potency between 5- and 25-fold lower than traditional approaches).

Cranor then calculated the difference between the social costs of using the CEPA's expedited methodology and the traditional risk-assessment methodology. He calculated the social costs by assuming a scenario (very close to the actual scenario in California when Cranor was writing) in which 400 substances had been identified by IARC as animal carcinogens but only 74 of them had been regulated, because conventional risk assessment methodologies had been too slow to evaluate any more of them. Thus, assuming that the animal carcinogens all turned out to be human carcinogens, the remaining 326 substances counted as regulatory "false negatives" under the conventional approach, insofar as they were not regulated even though they posed some level of risk to the public. In contrast, Cranor assumed that an expedited methodology would make it possible to evaluate risks from all 400 carcinogens, but with the potential for a few more false negatives and false positives because scientists would be employing a less accurate risk-assessment methodology. He calculated the costs of false negatives and false positives using a variety of different estimates, ranging from the typical figure given by economists (a 10:1 ratio between the social cost of false negatives and false positives) to figures that were much less favorable to the expedited approach (a 1:1 ratio).

Unsurprisingly, Cranor found that the CEPA's expedited approach to risk assessment was vastly better than traditional risk assessment approaches if one's goal was to minimize social costs. This finding held for all ratios between the costs of false negatives and false positives (from 10:1 to 1:1). In this case, it was far better to accept a small decrease in accuracy for the sake of avoiding the huge number of regulatory false negatives associated with traditional risk-assessment approaches. But it is particularly striking that Cranor obtained very similar results when he adjusted his assumptions to be much less favorable to the expedited approach. He considered a hypothetical case in which the expedited approach generated 50% more major overregulation than traditional risk assessment approaches and in which only 60% of the animal carcinogens identified by IARC actually turned out to be human carcinogens. Even under these assumptions, the expedited methodology generated fewer social costs when the ratio between the costs of false negatives and false positives fell between 10:1 and 2.5:1. Therefore, this case vividly illustrates the fact that when scientists are accepting modeling approaches for the sake of minimizing social costs in a regulatory environment, the ability to generate information rapidly may be a much higher priority than the ability to generate accurate results.

A critic of our approach to analyzing this case study might attempt to distinguish two kinds of scientific reasoning at work in this example. First, scientists have to determine their goals; in this case, Cranor's goal was to identify which risk-assessment model (i.e., the traditional model or the

expedited model) would minimize overall social costs. Second, scientists have to determine which model or theory will best enable them to achieve their goals; in this case, Cranor concluded that the expedited model would be best for achieving his goal of minimizing social costs. The critic might then argue that neither of these two kinds of reasoning involves nonepistemic values taking priority over epistemic ones. In the first instance (setting goals), nonepistemic values are directly relevant to the decision, so it does not make sense to say that they take priority. In the second instance (determining which model best achieves those goals), nonepistemic values do not appear to have a legitimate role to play; determining which model actually minimizes social costs appears to be a purely epistemic or objective matter.

That second claim might be disputed, for the reasons spelled out by Kuhn (1977) and in section 2. Even in cases where one's goal is purely epistemic (e.g., truth), desiderata relevant to the achievement of that end can stand in tension in ways that arguably make value judgments about how to balance those considerations inevitable. But suppose we grant that, once one sets the aims, the question of which representation most effectively achieves those aims is an entirely epistemic and objective matter. Notice that the critic's way of setting up the issue changes the subject in a subtle but significant way. Having allowed nonepistemic goals free rein in the first type of reasoning (i.e., goal setting), in the second type the critic focuses on trying to determine which model most effectively achieves those goals. Our focus, however, is on deciding how to prioritize the various qualities that representations can exemplify. We are asking: Which set of epistemic and nonepistemic qualities should be considered when choosing a model or theory? And our answer is: It depends on one's aims. Unless one has aims that are purely epistemic, this opens the door to the relevance of practical considerations that facilitate the achievement of the aims.

Our point is that scientists often have goals that are not purely epistemic, so when they engage in the process of choosing a representation that will best enable them to achieve their goals they can legitimately prioritize qualities that are nonepistemic. Even if determining which representation best enables the achievement of certain ends is a purely epistemic and objective matter, it can still turn out that the representation that is best for accomplishing particular practical aims is not the one that would be chosen based solely on considering epistemic desiderata. In the risk-assessment case, for example, scientists can appropriately choose to accept an expedited model if its nonepistemic qualities (e.g., speed) make it better for achieving their goal of minimizing social costs even though the conventional modeling technique appears to have the best qualities from a purely epistemic standpoint.

Nevertheless, the critic might reply that decisions about how to prioritize the various qualities that representations can exemplify should be classified

as part of the goal-setting process, and for that reason it remains misleading to claim that nonepistemic values are taking priority over epistemic values in a case like this one. In other words, once one sets the primary goal of minimizing social costs, scientists are forced to pursue various secondary goals, such as developing models with nonepistemic qualities like speed and ease of use. The critic would thus contend that within the constraints set by these nonepistemic values, scientists can still focus on developing the model with the best epistemic qualities—and therefore it does not make sense to say that epistemic values are being overridden by nonepistemic values.

Our response is threefold. First, as we clarified in section 2, in using the language of “taking priority” or “overriding” we should not be understood as claiming that allowing this sort of role for nonepistemic values would change what we think about the epistemic status of a theory. Instead, what we are claiming is that, in some cases, one can have good reason to accept a theory (i.e., take it as a basis for practical reasoning and action) that is recognizably inferior to a rival from an epistemic point of view.

Second, we think that characterizing the kind of reasoning that we are studying in the risk-assessment case and in the other cases in this article as a form of goal setting or framing can be bought only at the price of a strained or artificial reconstruction of the actual reasoning involved. For example, Cranor and the risk assessors that he studied did not appear to set a particular goal for the speed of their models and then attempt to choose the most epistemically virtuous model that attained the desired level of speed. Instead, the need to act forced them to choose between two models, and Cranor tried to convince reluctant risk assessors that they should sacrifice their desire to use the most realistic model because it was more important to use a speedy model if they really wanted to minimize social costs.

Finally, even if the reasoning that we are discussing in this article can be redescribed as a form of goal setting or framing rather than as a prioritizing of nonepistemic values, it is still significant and worthy of attention. Typically, when philosophers point out that nonepistemic values have a role to play in setting the goals of science, they are thinking about choosing which topics scientists should investigate. We are highlighting a much more subtle role for nonepistemic values, namely, that they must be weighed against epistemic values when choosing scientific representations in order to achieve the goals of scientific activity in a particular domain. Whether one describes this sort of reasoning as a prioritizing of nonepistemic values or instead as part of framing or goal setting, we think that it deserves attention in the literature on values in science.

4. Rapid Assessment Methods for Wetland Banking. Recent efforts to engage in wetland mitigation banking provide another example of how nonepistemic values such as ease of use can sometimes take priority over

epistemic values such as accuracy in making predictions, depending on scientists' particular goals or purposes for using scientific models and methods. Over the past 100 years, society has experienced a 180-degree shift in attitudes toward wetlands. Throughout much of American history, the drainage of wetlands was considered to be the responsible thing to do, both to maximize the productivity of land and to prevent diseases (Meyer 2004). However, over the course of the twentieth century, scientists began to recognize that wetlands performed a host of valuable functions, including water purification, storm protection, nutrient cycling, and animal habitat (Mitsch and Gosselink 2007). As a result, new policies have been put in place to make it more difficult to drain wetlands. Among the most important of these was Section 404 of the Clean Water Act, which requires the Army Corps of Engineers (in consultation with the Environmental Protection Agency) to provide a permit to individuals who wish to dredge or fill wetlands on their property (Hough and Robertson 2009).

In practice, neither the Corps nor the EPA has been eager to deny permits for damaging wetlands, but they have been somewhat more inclined to demand that mitigation measures be taken in response. In practice, the most common approach to mitigation is to preserve or restore one area of wetland in order to compensate for a wetland area that is destroyed elsewhere (Hough and Robertson 2009, 23). In recent years, regulatory agencies, developers, and entrepreneurs have developed a "banking" approach to handling this mitigation effort. Rather than forcing developers to preserve or restore wetlands themselves, regulatory agencies allow developers to purchase mitigation "credits" from specialists who create "banks" of preserved or restored wetlands. By 2005, there were over 350 active banks, 75 sold-out banks, and over 150 banks under review. Annually, almost \$3 billion is spent on compensatory mitigation, primarily for either contractor-operated replacement wetlands or for banked wetland credits (Hough and Robertson 2009, 24–25).

Science becomes important to this mitigation banking process because regulators have to decide whether restored wetlands are sufficiently similar to the destroyed wetlands in order to justify trading the two. Thus, they need to employ models as part of assessment methods that yield comparisons of different wetlands in terms of their key socially relevant features. Geographer Morgan Robertson (2004, 2006) has analyzed these methods and emphasized how the models for characterizing wetlands in the banking context differ from those that one would employ if one were aiming for the most ecologically sophisticated characterization. These differences arise in large part because bankers need assessments to be authoritative, cheap, and quick (Salzman and Ruhl 2000, 665). In other words, in order to develop a flourishing mitigation banking system, these nonepistemic values must sometimes be prioritized over epistemic values like predictive accuracy when

choosing assessment models. The influence of these values affects numerous aspects of the assessment process.

First, in order to decide whether a natural wetland (which is to be destroyed) and a restored wetland should be classified as equivalent, scientists and regulators focus on a specific bundle of features (e.g., particular functions or ecosystem services) that are regarded as most relevant for exchange. Second, they employ “rapid assessment methods” (RAMs), which consist of algorithms that convert a variety of data about a wetland into numerical scores that estimate a wetland’s functional value for providing the features of interest (Robertson 2006, 373). Third, in order to keep the process simple, a wetland is typically represented by one main score rather than a variety of different functional scores. For example, because plants are relatively easy to look at, and because botanists have developed some relatively high-quality analyses of the plants in various sorts of wetlands, the function of “floristic biodiversity” has become a rough indicator for most other wetland functions. Fourth, even the assessment of floristic biodiversity depends on simplifying assumptions. For example, Robertson explains that regulatory guidelines require that the identification of plant species be performed in May or June, when most plants have not yet flowered. Therefore, the assessors have to base their assessments on the observation of plant features that may indicate the identity of particular species but that are not compelling by rigorous botanical standards.

5. Objections. Although the preceding sections provide a number of examples of the ways that nonepistemic values can receive priority over epistemic values in assessing scientific representations for practical purposes, one might still worry that the study’s argument is problematic. We would like to address three potential objections: (1) that the argument is trivial, (2) that the cases considered in the article are examples of bad science, and (3) that more clarity is needed about the conditions under which nonepistemic values can legitimately override epistemic values.

First, one might think that the study’s central claim is trivial. Perhaps it is entirely obvious that as long as scientific representations are being assessed for practical purposes, nonepistemic values can take priority over epistemic ones. The problem with this objection is that this allegedly obvious insight has not been clearly discussed in the science-and-values literature. In fact, there has been relatively little discussion of how to assess models, hypotheses, or theories from a practical perspective. While some figures have acknowledged that hypotheses can be accepted as a basis for action rather than believed to be true, they have not been very clear about what these two cognitive attitudes involve (see Elliott and Willmes 2013). Moreover, some of those who explicitly claim to be talking about the assessment of hypotheses for practical purposes have still argued that nonepistemic values should be

incorporated only in a secondary role (see, e.g., Elliott 2011). In fact, if this study's claims were obvious and trivial, then the view that nonepistemic values should be limited to a secondary role should not be as prevalent as it is. Suppose, for example, that one is engaged in a purely epistemic assessment of a theory; in that case, nonepistemic values appear to be entirely irrelevant and have no legitimate role to play in assessment, even as tiebreakers. But suppose that one is engaged in a practical assessment of a theory for particular purposes; if it were truly obvious that nonepistemic values can take priority over epistemic values in such cases, then nonepistemic values would not need to be limited to a secondary role. Therefore, there appears to be significant confusion about what is involved in assessing scientific representations for practical purposes and when this is legitimate.

Given the apparent confusion in the existing literature, the argument of this study does not appear to be trivial. But a second objection is that the assessment of scientific representations for practical purposes results in bad science that has lost its distinctively epistemic character. One way to express this worry would be to claim that when scientists assess representations for practical purposes (and therefore potentially allow nonepistemic values to receive priority over epistemic values), it results in a sort of hybrid science that is illegitimate.

But there are two problems with this objection. First, numerous commentators have recently argued that scientific practice as a whole has taken on a largely hybrid character in which both theoretical and practical considerations play a central role. This "mixed" scientific practice has been called "technoscience," as well as "mode-2," "entrepreneurial," "postacademic," or "postmodern" science (see, e.g., Nordmann, Radder, and Schiemann 2011). Clearly, there is value to identifying cases where important norms of science are sacrificed for the sake of pursuing these hybrid forms of practice. However, if one immediately rules out all hybrid forms of scientific practice without further justification, one is left with a severely limited picture of scientific reasoning. A second problem with the objection that the prioritization of nonepistemic values results in illegitimate, hybrid science is that, as discussed in section 2, simplifications and idealizations are commonly used even in traditional areas of science for making scientific reasoning easier. Especially in practical contexts, the attitude scientists often take seems to be that a particular idealization, even if it does not pass our most rigorous epistemic standards for evaluation, is "good enough for government work" or "a close enough approximation given our present purposes."

Therefore, as long as scientists are not adopting a cognitive attitude toward their representations that requires the prioritization of epistemic values, it seems to be an open question whether or not it is inappropriate to sacrifice epistemic concerns for other priorities. After all, humans have a number of goals (e.g., promoting health, happiness, and flourishing communities) in

addition to the goals of seeking truths about the world. Although it is generally easier to achieve these other goals when one has highly reliable information, the case studies in the preceding sections show that epistemic concerns need not always be an overriding priority. There can, of course, be examples where regulatory policies are designed in a suboptimal fashion so that too much accuracy is sacrificed in order to meet social goals (as may have happened in the case of wetland mitigation banking). But in other cases, the regulatory environment can be well suited for meeting social goals, in which case there does not appear to be anything wrong with engaging in practical assessments of scientific representations so that those goals can be met.

But this response raises the specter of a third worry. Given the argument of this study, how can one distinguish whether nonepistemic values are receiving too much priority when assessing scientific representations? If one adopts a naturalized approach to the philosophy of science, as we do, it is difficult to formulate an easy or universal answer to this question. The extent to which nonepistemic values should be prioritized over epistemic values arguably depends a great deal on the expectations and conventions of the scientific community in particular contexts. As we indicated in our response to the previous objection, there appear to be many cases in which scientists do choose models that are less ideal from an epistemic standpoint for the sake of other desiderata such as convenience and ease of use. However, this may be more common and easier to justify in some cases, such as when researchers are assessing piecemeal models, rather than others, such as when they are assessing theories or hypotheses. Nevertheless, it is important to recognize that even theories of broad scope such as quantum mechanics and gravitational theory are widely adopted by the scientific community despite being incompatible and thus epistemically problematic.

We think that attention to the case studies presented in sections 3 and 4 suggests two guiding principles that may provide a starting point for determining when and to what extent nonepistemic values can appropriately receive priority over epistemic ones. First, if nonepistemic values are to play an appropriate role in assessing scientific representations, those engaged in the assessments need to be explicit about the goals of their assessments and the roles that nonepistemic values played in the assessments as a result. Second, nonepistemic values should receive priority only to the extent that they advance the goals associated with the assessments that are in play.

The principle that scientists should be as explicit as possible about the sorts of assessments that they are performing and the ways that nonepistemic values are influencing those assessments fits well with other literature that emphasizes the importance of acknowledging how values influence scientific reasoning (e.g., Douglas 2009; Elliott 2011). The value of this explicitness is that if those engaged in assessing a hypothesis or theory

are clear enough about their goals and the criteria that they are employing in the assessment process, those who disagree about those goals or criteria can “backtrack” and adopt their own alternative assessments and conclusions (McKaughan and Elliott 2013). Along these lines, one of the strengths of Cranor’s analysis of risk-assessment models for estimating carcinogenicity (discussed in sec. 3) is that he was so clear about his criteria for assessment. By being explicit about his goal of determining which model minimized social costs (measured in economic terms), he made it particularly easy for others to determine how alternative assumptions and assessment criteria would result in different conclusions. For example, he showed how different assumptions about the economic costs of false positive and false negative errors would yield different assessments of the models. Moreover, he not only clarified how the risk assessment models compared from the practical perspective of minimizing social costs, but he also clarified how they compared from a purely epistemic perspective (i.e., in terms of how frequently the expedited risk-assessment model overestimated and underestimated carcinogenicity relative to the traditional model).

The second principle (namely, that nonepistemic values should receive priority only to the extent that they advance the goals associated with particular assessments) has played an especially obvious role in recent debates about the wetland restoration case discussed in section 4. For example, many ecologists have pointed out that the approaches used for modeling wetland functions have become so crude that they fail to preserve important wetland functions (Froelich 2003). In fact, regulatory agencies have often concluded that even the sorts of floristic biodiversity assessments discussed in section 4 require too much effort, so natural and artificial wetlands have sometimes been compared solely in terms of the number of acres present (Ruhl and Gregg 2001, 381). To the extent that the goal of regulatory agencies is to choose wetland models that promote the preservation of important wetland functions, one can argue that nonepistemic values are inappropriately trumping epistemic values when the methods and models used by these agencies become too epistemically crude to fulfill that goal.

This does not mean that regulators must always use the methods and models that are most accurate from an epistemic perspective, of course. For example, it is plausible that in at least some contexts, some form of market-based approach to environmental regulation may be socially valuable. And in those cases, it will generally not be feasible to use the most accurate ecological models for regulatory purposes. This point is expressed particularly well by an EPA official quoted by Robertson, who joked that efforts to become more and more ecologically precise could lead to markets in “habitat for middle-aged great blue herons who don’t like shrimp, or something. Obviously, I can’t imagine even trying to do that. . . . You can define a unit

so that you're going to have flourishing mitigation banking. You can also define a unit so that, should there ever be one exchanged, it would be environmentally precise. And those are at potentially different extremes" (Robertson 2004, 368). Therefore, in these sorts of regulatory contexts, where scientific precision conflicts with regulatory goals (such as setting up a flourishing wetland banking system), nonepistemic values can appropriately take priority over epistemic values, in the sense that a relatively inaccurate model or method can be chosen for regulatory purposes. But the balance between epistemic and nonepistemic values can be criticized if it results in the choice of models that do not fulfill the purposes for which they were chosen (e.g., if they yield results that are too crude to preserve important wetland functions).

A potential worry about these two principles is that they seem to allow questionable influences of nonepistemic values on theory assessment. Suppose, for example, that a think tank working for the petroleum industry performed a study and was explicit about the fact that the study conclusions were based on simplifications designed to advance their practical goal of underestimating the risks associated with climate change. While these influences of nonepistemic values seem problematic, they might satisfy the two principles elaborated in this section. Faced with these sorts of difficulties, one might question whether there is any reason to prefer the approach to regulating values in this article to the sorts of approaches that relegate nonepistemic values to a secondary role (e.g., Douglas 2009; Elliott 2011; Steel and Whyte 2012).

Our response is that one cannot capture the reasoning exemplified by the cases in this article while relegating nonepistemic values to a secondary role. We have argued that scientific practice often incorporates practical or mixed assessments of scientific representations and that it is legitimate to give nonepistemic values priority when assessing representations in this way. In these cases, nonepistemic values are clearly not playing a mere tie-breaking role, and they are also not playing the sort of indirect role for values discussed by Douglas (2009). She claims that the distinguishing feature of values playing an indirect role is that their influence diminishes as scientific uncertainty is eliminated (Douglas 2009, 97). In the sorts of cases discussed in this article, the influence of nonepistemic values does not diminish along with uncertainty. For example, even if scientists were absolutely sure that conventional risk-assessment models provided epistemically superior results, they might still choose to use a less accurate expedited model if it produced much faster results and thus lessened social costs. Thus, these values are playing a direct role in our case study; scientists regard the production of faster results as a "reason in itself" for choosing the expedited model and not merely as a reason for altering their standards of evidence in response to the

possibility that they could be choosing an erroneous model. In fact, they are choosing a more erroneous model on purpose because of its nonepistemic virtues.⁵

Moreover, we think that our two principles for regulating nonepistemic values are more promising than they might initially appear. We contend that much of the reason why the case of the petroleum-industry think tank seems so problematic is that these sorts of interest groups are almost never truly explicit about the practical goals that influence their work. If the think tank genuinely acknowledged that they were trying to minimize evidence for climate change and made their methodological choices explicit, it would alleviate much of the concern about their activities. Other scientists, policy makers, and journalists could “backtrack” and point out how one would obtain different results if one employed more epistemically reasonable assumptions. We would also point out that once the practical goals of the think tank were made explicit, the scientific community could reflect on whether those practical goals were reasonable enough to merit significant consideration. Based on our case studies, we would suggest that the scientific community is willing to consider practical appraisals of theories and models if there seem to be compelling social reasons to do so, but they are less likely to take them seriously if they are performed primarily for the benefit of a specific interest group. These judgments about what sorts of practical assessments to entertain constitute a significant way in which ethical and epistemological issues intertwine in scientific practice and merit further analysis from philosophers of science.

6. Conclusion. We have shown how nonepistemic values can legitimately influence the assessment of scientific representations for practical purposes not only as secondary considerations in situations of uncertainty but also as factors that can take priority over epistemic values. If a consideration can clearly be identified as nonepistemic (remembering that various figures disagree about how to make this distinction), then it is not relevant to the purely epistemic appraisal of models or theories. However, when scientists develop models or other scientific representations, they aim not only to develop an

5. Although we are claiming that nonepistemic values can legitimately play a direct role in theory appraisal, which Douglas appears to deny, we think that a relatively minor alteration in her position could make it compatible with ours. She already acknowledges that nonepistemic values can play a direct role in numerous aspects of scientific reasoning, including the choice of research projects and methodologies (Douglas 2009, 98–99). Whereas she focuses on appraising theories from a purely epistemic standpoint and argues correctly that nonepistemic values should not play a direct role in this context, we think that she could allow a direct role for nonepistemic values in mixed or practical contexts of theory appraisal.

adequate fit with the world but also to develop an adequate fit with the needs and aims of those who use the representations. A combination of epistemic and nonepistemic considerations will nearly always be relevant to the practical needs of users. This is significant, because a great deal of real world decision making about the status and use of representations in actual scientific practice takes place in these mixed contexts of appraisal (i.e., where both epistemic and practical considerations are relevant). In such contexts, while epistemic values remain relevant, it can make good sense to allow nonepistemic values not simply to serve as secondary factors but indeed to “take priority over” or “override” purely epistemic values. For example, we have seen in cases of risk assessment and wetland mitigation banking that models and assessment methods frequently need to provide rapid results, to be relatively inexpensive to use, and to be applicable to a wide range of cases. In order to achieve these purposes, it is sometimes necessary to sacrifice epistemic considerations such as accuracy.

We showed that scientists can factor nonepistemic values into their reasoning without violating their epistemic responsibilities as long as they adopt appropriate cognitive attitudes toward the representations and conclusions that they develop. So, for example, if scientists are accepting a particular model for the purposes of generating information for policy makers, it is entirely legitimate for them to consider not only whether the model is reasonably accurate but also to consider whether it generates information quickly enough to be useful. This raises the question of what criteria are needed for determining when nonepistemic values are inappropriately trumping epistemic values. Based on our case studies, we suggested two principles as a starting point for further discussion. First, those engaged in practically assessing models, theories, or hypotheses need to be as explicit as possible about the goals of their assessments and the roles that nonepistemic values played in the assessments as a result. Second, nonepistemic values should receive priority only to the extent that they advance the goals associated with those assessments. Our point is that this further discussion about how to distinguish appropriate and inappropriate values in science would do well to focus less on whether values are epistemic and more on whether particular values help to achieve the purposes for which a particular representation is adopted in a specific context.

REFERENCES

- Biddle, Justin. 2013. “State of the Field: Transient Underdetermination and Values in Science.” *Studies in History and Philosophy of Science* 44:124–33.
- Brown, Matthew. 2013. “Values in Science beyond Underdetermination and Inductive Risk.” *Philosophy of Science* 80 (Proceedings), forthcoming.
- Cohen, Jonathan. 1992. *An Essay on Belief and Acceptance*. New York: Oxford University Press.

- Cranor, Carl. 1995. "The Social Benefits of Expedited Risk Assessments." *Risk Analysis* 15:353–58.
- . 2011. *Legally Poisoned: How the Law Puts Us at Risk from Toxicants*. Cambridge, MA: Harvard University Press.
- Douglas, Heather. 2009. *Science, Policy, and the Value-Free Ideal*. Pittsburgh: University of Pittsburgh Press.
- Elliott, Kevin. 2011. *Is a Little Pollution Good for You? Incorporating Societal Values in Environmental Research*. New York: Oxford University Press.
- . 2013. "Douglas on Values: From Indirect Roles to Multiple Goals." *Studies in History and Philosophy of Science* 44:375–83.
- Elliott, Kevin, and David Willmes. 2013. "Cognitive Attitudes and Values in Science." *Philosophy of Science* 80 (Proceedings), forthcoming.
- Froelich, Adrienne. 2003. "Army Corps: Retreating or Issuing a New Assault on Wetlands." *BioScience* 53:130.
- Giere, Ronald. 2004. "How Models Are Used to Represent Reality." *Philosophy of Science* 71 (Proceedings): 742–52.
- . 2006. *Scientific Perspectivism*. Chicago: University of Chicago Press.
- Hough, Palmer, and Morgan Robertson. 2009. "Mitigation under Section 404 of the Clean Water Act: Where It Comes from, What It Means." *Wetlands Ecology and Management* 17:15–33.
- Howard, Don. 2006. "Lost Wanderers in the Forest of Knowledge: Some Advice on How to Think about the Relation between Discovery and Justification." In *Revisiting Discovery and Justification: Historical and Philosophical Perspectives on the Context Distinction*, ed. Jutta Schickore and Friedrich Steinle, 3–22. Dordrecht: Springer.
- Kitcher, Philip. 2001. *Science, Truth, and Democracy*. New York: Oxford University Press.
- Kuhn, Thomas. 1977. "Objectivity, Value Judgment, and Theory Choice." In *The Essential Tension: Selected Studies in Scientific Tradition and Change*, 320–39. Chicago: University of Chicago Press.
- Lacey, Hugh. 2005. *Values and Objectivity in Science: The Current Controversy about Transgenic Crops*. Lanham, MD: Lexington.
- Longino, Helen. 1996. "Cognitive and Non-cognitive Values in Science: Rethinking the Dichotomy." In *Feminism, Science, and the Philosophy of Science*, ed. Lynn Hankinson Nelson and Jack Nelson, 39–58. Dordrecht: Kluwer.
- . 2002. *The Fate of Knowledge*. Princeton, NJ: Princeton University Press.
- Matthewson, John, and Michael Weisberg. 2009. "The Structure of Tradeoffs in Model Building." *Synthese* 170:169–90.
- McKaughan, Daniel. 2007. "Toward a Richer Vocabulary for Epistemic Attitudes: Mapping the Cognitive Landscape." PhD diss., University of Notre Dame.
- McKaughan, Daniel, and Kevin Elliott. 2013. "Backtracking and the Ethics of Framing: Lessons from Voles and Vasopressin." *Accountability in Research* 20:206–26.
- Meyer, William. 2004. "From Past to Present: A Historical Perspective on Wetlands." In *Wetlands*, ed. Sharon Spray and Karen McGlothlin, 84–100. Lanham, MD: Rowman & Littlefield.
- Michaels, David. 2008. *Doubt Is Their Product: How Industry's Assault on Science Threatens Your Health*. New York: Oxford University Press.
- Mitsch, William, and James Gosselink. 2007. *Wetlands*. 4th ed. Hoboken, NJ: Wiley & Sons.
- Nordmann, Alfred, Hans Radder, and Gregor Schiemann. 2011. "Science after the End of Science? An Introduction to the 'Epochal Break Thesis.'" In *Science Transformed? Debating Claims of an Epochal Break*, ed. Alfred Nordmann, Hans Radder, and Gregor Schiemann, 1–15. Pittsburgh: University of Pittsburgh Press.
- NRC (National Research Council). 1983. *Risk Assessment in the Federal Government: Managing the Process*. Washington, DC: National Academy Press.
- Potochnik, Angela. 2012. "Feminist Implications of Model-Based Science." *Studies in History and Philosophy of Science* 43:383–89.
- Robertson, Morgan. 2004. "The Neoliberalization of Ecosystem Services: Wetland Mitigation Banking and Problems in Environmental Governance." *Geoforum* 35:361–73.
- . 2006. "The Nature That Capital Can See: Science, State, and Market in the Commodification of Ecosystem Services." *Environment and Planning D, Society and Space* 24:367–87.

- Rooney, Phyllis. 1992. "On Values in Science: Is the Epistemic/Non-epistemic Distinction Useful?" In *Proceedings of the 1992 Biennial Meeting of the Philosophy of Science Association*, ed. David Hull, Micky Forbes, and Kathleen Okruhlik, 13–22. East Lansing, MI: Philosophy of Science Association.
- Ruhl, J. B., and R. Juge Gregg. 2001. "Integrating Ecosystem Services into Environmental Law: A Case Study of Wetlands Mitigation Banking." *Stanford Environmental Law Journal* 20:365–92.
- Salzman, James, and J. B. Ruhl. 2000. "Currencies and the Commodification of Environmental Law." *Stanford Law Review* 53:607–94.
- Steel, Daniel. 2010. "Epistemic Values and the Argument from Inductive Risk." *Philosophy of Science* 77:14–34.
- Steel, Daniel, and Kyle P. Whyte. 2012. "Environmental Justice, Values, and Scientific Expertise." *Kennedy Institute of Ethics Journal* 22:163–82.
- van Fraassen, Bas. 2008. *Scientific Representation: Paradoxes of Perspective*. Oxford: Oxford University Press.
- Weisberg, Michael. 2007. "Three Kinds of Idealization." *Journal of Philosophy* 104:639–59.

Copyright of Philosophy of Science is the property of Philosophy of Science Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.