Contents lists available at SciVerse ScienceDirect



Studies in History and Philosophy of Science

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

State of the field: Transient underdetermination and values in science

Justin Biddle

School of Public Policy, Georgia Institute of Technology, 685 Cherry Street, Atlanta, GA 30332, USA

ARTICLE INFO

ABSTRACT

Article history: Received 30 September 2011 Received in revised form 1 July 2012 Available online 22 November 2012

Keywords: Transient underdetermination Science and values This paper examines the state of the field of "science and values"—particularly regarding the implications of the thesis of transient underdetermination for the ideal of value-free science, or what I call the "ideal of epistemic purity." I do this by discussing some of the main arguments in the literature, both for and against the ideal. I examine a preliminary argument from transient underdetermination against the ideal of epistemic purity, and I discuss two different formulations of an objection to this argument—an objection that requires the strict separation of the epistemic from the practical. A secondary aim of the paper is to suggest some future directions for the field, one of which is to replace the vocabulary of values that is often employed in the literature with a more precise one.

© 2012 Elsevier Ltd. All rights reserved.

Studies in Hist and Philosoph

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

1. Introduction

In the first half of the twentieth century, the philosophy of science community in Europe and North America thought very deeply about the interactions between science and society. Topics such as the influence of values in scientific reasoning and the place of science in a democratic society were at the center of the philosophical agenda. By around the middle of the century, however, the agenda had changed dramatically, the result of which was a significant decline in social and political engagement (Howard, 2003; Reisch, 2005). Philosophers of science during this period, for example, recognized that science is in fact a social and value-laden enterprise, but they banished such considerations from the context of justification to the context of discovery; as such, they were thought to be a matter for sociologists and psychologists, not for philosophers (e.g., Popper, 1959 [1934]; Reichenbach, 1938). Fortunately, the pendulum has shifted once again; the turn of the century has witnessed a renewed engagement with social and political concerns, Led by feminist philosophers of science concerned with the influence of sexist presuppositions on the content of science and with the role that science sometimes plays in reinforcing gender stereotypes (e.g., Longino, 1990; Nelson, 1990; Okruhlik, 1994), philosophers are once again examining the role of values in science and, more generally, of science in society.

This shift in the agenda of the philosophy of science community should be welcomed for a number of reasons—not the least of which is that it has the potential to lead to a better philosophy of science. Science is an activity that takes place in real social, political, cultural, and economic contexts, and it both shapes, and is shaped by, these contexts. Recognition of this basic fact raises questions that philosophers of science have often neglected, such as how science should impact public policy making (Kitcher, 2011), and it has the potential to lead to a reconceptualization of the nature of science itself. Such a reconceptualization could have implications for philosophy of science more broadly, including debates over realism and anti-realism, confirmation, and explanation (e.g., Longino, 2002).

This paper examines the state of the field of "science and values"—and in particular, of the status of the ideal of value-free science, or what I call the "ideal of epistemic purity." Since the middle of the twentieth century, the standard view among philosophers of science has been that certain aspects of scientific research can and should be value free—or, more specifically, free from all "non-epistemic," or "contextual" factors, such as moral and political values.¹ Some contextual factors, of course, appropriately influence certain aspects of scientific research; moral and political considerations rightly influence the choices of problems to address, the decisions of how research should be applied, and the setting of

E-mail address: justin.biddle@pubpolicy.gatech.edu

¹ In this paper, "contextual factors" refer to any factor that is traditionally regarded as non-epistemic—that is, any factor that falls outside of the domain of logic, evidence, and epistemic values. For a discussion of epistemic values, see Kuhn (1977), McMullin (1983), and Laudan (1984).

constraints on experimental practices. The traditional view, however, maintains that contextual factors can and should be excluded from the "internal" workings of science—i.e., the epistemic appraisal of theories, models, and hypotheses.² More specifically, this view maintains that (1) the proper application of scientific methods will, as a matter of fact, always screen out all contextual factors, and (2) scientists ought to apply scientific methods properly, thereby screening out all contextual factors (cf. Ruphy, 2006, p. 190). This traditional view is often referred to as the "ideal of value-free science;" for reasons that I will discuss later, I will refer to it as the "ideal of epistemic purity."

While a variety of different arguments have been put forward against the ideal, the most common, and most influential, is the argument from underdetermination.³ One version of this argument begins by stating that *all* theories (or hypotheses, models, etc.) are underdetermined by logic and *all possible* evidence;⁴ this leaves a gap between logic and evidence, on the one hand, and theory choice, on the other, which is inevitably filled by contextual factors. As is apparent, this argument depends upon a very strong version of the underdetermination thesis (following Kitcher (2001), let us call it the thesis of *global* underdetermination): not surprisingly, the standard response to this argument is to question this version of the thesis—a strategy employed by Laudan and Leplin (1991), Norton (2008), Kitcher (2001), and others. While there is a large and important literature on the thesis of global underdetermination, I will not discuss it here.

The reason for this is the possibility that a weaker version of the underdetermination thesis-the thesis of transient underdetermination—suffices to undermine the ideal of epistemic purity (e.g., Howard, 2009; Nelson, 1990; Potter, 1997; Rolin, 2002; Wray, 1999). This thesis states merely that some theories, hypotheses, and models are underdetermined by logic and the currently available evidence. This thesis, moreover, is undoubtedly true. The primary purpose of this paper is to examine the implications of the thesis for the ideal of epistemic purity by reviewing some of the most important arguments in the literature. In Section 2. I discuss a prominent defense of the ideal of epistemic purity by Philip Kitcher (2001), and in Section 3. I examine an argument against the ideal by Heather Douglas. Following this, I discuss two different formulations of an objection to this argument, one by Richard Jeffrey (Section 4.1) and the other by Sandra Mitchell (Section 4.2). In examining this debate, I will highlight the price that one must pay in order to maintain the ideal of epistemic purity in light of the thesis of transient underdetermination.

This discussion will set the stage for a secondary aim of the paper, which is to suggest some further lines of research. While there are many future avenues that could be taken by participants in the debate over "science and values," one important one is to clarify and make more precise some of the central concepts employed within the debate. As noted, critics of the ideal of epistemic purity typically employ the vocabulary of "values," and the central question within the science and values literature is whether values especially moral and political values—appropriately influence theory or hypothesis appraisal. There are a number of reasons, however, to think that the terminology of values is, at best, misleading, and that the notion that they are attempting to criticize is better characterized as epistemic purity rather than value free**dom**. The end of the paper draws upon a neglected essay by Otto Neurath in order to illustrate one way the terminology of values might be re-examined.

2. Underdetermination: Transient, permanent, and global

One of the goals that Kitcher sets for himself in the first part of *Science, Truth, and Democracy* is to defend the ideal of epistemic purity, or what he calls the "ideal of objectivity." To do so, he identifies three different formulations of the thesis of underdetermination—transient, permanent, and global—and argues that, of those formulations that are plausible, none poses a difficulty for the ideal of objectivity. As he defines it, the thesis of transient underdetermination states that *some* theories are underdetermined by logic and the *currently available* evidence. The thesis of permanent underdetermination states that *some* theories are underdetermined by logic and *all possible* evidence, or at least all evidence to which scientists will ever have access. Finally, the thesis of global underdetermined (Kitcher, 2001, pp. 30, 31).

He acknowledges that much cutting-edge research is transiently underdetermined, in that we do not have sufficient evidence available to affirm the truth of one particular theory or hypothesis. Yet he dismisses this type of underdetermination immediately, asserting without argument that it is "familiar and unthreatening" (2001, p. 30). Critics of his conception of objectivity, he maintains, must argue for a stronger thesis, the thesis of global underdetermination:

The underdetermination thesis obtains its bite when permanent underdetermination is taken to be rampant. The global underdetermination thesis, which contends that all (or virtually all) instances of what scientists treat as transient underdetermination are, when properly understood, examples of permanent underdetermination, poses a genuine threat to the ideal of objectivity (2001, p. 31).

The first step of Kitcher's defense of the ideal of objectivity, thus, is to dismiss the relevance of transient underdetermination.

Kitcher is hardly alone in making this move; those who defend the ideal of epistemic purity by criticizing the thesis of underdetermination focus almost exclusively on the global version of this thesis. To take just one more example, John Norton, in his contribution to a recent volume on science and values, restricts his attention to the thesis of global underdetermination-he explicitly disregards transient underdetermination (or what he calls "merely de facto underdetermination")--and argues that no reasonable account of induction supports the global thesis (Norton, 2008, p. 20).⁵ In certain contexts, of course, it is perfectly legitimate to focus exclusively upon the global version of the underdetermination thesis; whether all theories or hypotheses are undetermined by all possible evidence is an interesting and significant philosophical question. However, if the context within which one is working is the debate over science and values, including the defenses and criticisms of the ideal of value freedom or epistemic purity, it is a mistake to ignore transient underdetermination (Howard, 2009; Nelson, 1990; Potter, 1997; Rolin, 2002; Wray, 1999).

² For defenses of this view, see Giere (2003), Kitcher (2001), Koertge (2000, 2003), McMullin (1983), Pinnick (2003), and Ruphy (2006).

³ Arguments against the ideal of epistemic purity include arguments from the value-ladenness of scientific concepts (Dupré, 2007), arguments from inductive risk (Douglas, 2000, 2009; Rudner, 1953; Wilholt, 2009), and arguments from underdetermination (Howard, 2006; Howard, 2009; Kourany, 2003a, 2003b; Longino, 1990, 2002; Neurath, 1983 [1913]). I view the first two types of arguments as special cases of the third. An alternative argument, which one might call the "naturalistic argument against the ideal of epistemic purity," can be found in Solomon (2001).

⁴ This version of the thesis is due to Quine (1951). The other classic presentation of the thesis is in Duhem (1954 [1906]), though Duhem's version is significantly weaker than Quine's.

⁵ Whether Norton's argument could apply to cases of transiently underdetermined research is discussed briefly in footnote 9.

3. The preliminary argument

One might develop a preliminary argument from transient underdetermination against the ideal of epistemic purity-which I will simply call the Argument from Transient Underdetermination (ATU)—as follows. There are many areas of current, cuttingedge science that are underdetermined, at least transiently, by evidence. In many of these areas, especially those relevant for public policy making, hypotheses must be evaluated quickly, before all of the evidence is in. Suppose, for example, that we need to determine whether a particular chemical used in pesticides is sufficiently safe or has acceptable environmental impact, or whether a drug that is currently on the market should be taken off the market. Suppose, furthermore, that the available evidence does not unambiguously determine which hypothesis should be accepted, or which decision should be made. In these cases, we do not have the luxury of waiting until all of the evidence is in, as postponing a decision could result in severe environmental degradation, loss of life, or other adverse effects. In situations such as this, there is a gap between evidence and hypothesis choice, and this gap is inevitably filled by contextual factors. Thus, in these cases, the ideal of epistemic purity fails.

As a way of illustrating this argument, it is helpful to discuss a case study developed by Heather Douglas on the carcinogenic effects of dioxins in laboratory rats (Douglas 2000). She argues that in three different parts of the research process—the choice of methodology, the gathering and characterization of data, and the interpretation of data—"non-epistemic values" such as ethical values play an essential role. In this section, I will focus on her critique of the value-free character of the interpretation of data, as it provides a good illustration of ATU and sets the stage for what I take to be the most prominent objection to ATU, an objection that I will consider in the next section.

In the research that Douglas examines, there is extensive data on the effects of relatively high doses of dioxins in laboratory rats, but there is very little data on the effects of low doses. As a result, it is necessary to choose a model for extrapolating data to lowdose regions. Within the toxicology community, there is a vigorous debate over which model should be chosen. For much of the twentieth century, the "threshold model" was adopted, according to which there is a threshold below which chemicals are entirely safe. In the latter part of the century, however, this assumption came under increasing scrutiny and, in certain cases, was rejected in favor of a "no-threshold model," according to which no dose is completely safe. For example, scientists studying cancer in the 1960s discovered that any amount of radiation, no matter how small, can cause a cell to mutate, which in turn can cause cancer (Douglas, 2000, p. 574). Of course, the probability of this occurring is very low; nevertheless, it is not zero, which implies that the threshold model is, in this case, false.

The question of which of these models best applies to the dioxin case is an important one, not only epistemically but also politically, as the different models have very different regulatory implications. Given that the threshold model implies that some doses are completely safe, it suggests that relatively weaker regulation of dioxins is acceptable; the no-threshold model, on the other hand, suggests tighter forms of regulation. Unfortunately, the available evidence underdetermines the choice of extrapolation model; within the toxicology community, there is a controversy over which model best applies to the dioxin case, and this controversy reflects the insufficiency of evidence to decide the issue (Douglas, 2000, pp. 575–577). If we grant, then, that the choice of model is underdetermined, and if we grant that a decision must be made before all of the evidence is in, then we must conclude that contextual factors will play an inevitable role.

Douglas does not couch her argument in terms of transient underdetermination, but rather in terms of inductive risk. The theoretical basis of her argument is provided in Richard Rudner's wellknown 1953 paper, "The Scientist *Qua* Scientist Makes Value Judgments." Rudner's argument can be outlined as follows:

- P1 The scientist qua scientist accepts or rejects hypotheses.
- P2 No hypothesis is ever completely (with 100% certainty) verified.
- P3 The decision to accept or reject a hypothesis depends upon whether the evidence is sufficiently strong.
- P4 Whether the evidence is *sufficiently* strong is "a function of the *importance*, in a typically ethical sense, of making a mistake in accepting or rejecting the hypothesis" (Rudner, 1953, p. 2, emphasis in original).
- C Therefore, the scientist qua scientist makes (ethical) value judgments.

To illustrate this argument, Rudner writes:

If the hypothesis under consideration were to the effect that a toxic ingredient of a drug was not present in lethal quantity, we would require a relatively high degree of confirmation or confidence before accepting the hypothesis—for the consequences of making a mistake here are exceedingly grave by our moral standards. On the other hand, if say, our hypothesis stated that, on the basis of a sample, a certain lot of machine stamped belt buckles was not defective, the degree of confidence we should require would be relatively not so high. *How sure we need to be before we accept a hypothesis will depend upon how serious a mistake would be* (Rudner, 1953, p. 2, emphasis in original).

Douglas accepts Rudner's argument and adapts it to her own study to show that, in this particular case, value judgments play an inevitable role in the choice of a model for extrapolating data. Given the uncertainties associated with the choice of an extrapolation model, there is a significant chance that the wrong choice will be made. Moreover, the social, moral, and economic costs of wrongly adopting each model are very different. If we wrongly chose the threshold model, many people would get sick and die prematurely of cancer; the moral cost, in this case, is very high, not to mention the economic costs of treating these individuals. On the other hand, if we wrongly chose the no-threshold model, the worst that would happen is that corporate profits would be slightly reduced. Because of this, Douglas concludes that in this area of research, non-epistemic values, and in particular ethical values, should not only influence the decision, but should lead us to adopt the no-threshold model. In regulatory science and other areas of science that have important implications for policy making, "non-epistemic values are a required part of the internal aspects" of scientific reasoning" (Douglas, 2000, p. 559, emphasis added).

Douglas strengthens her argument for the value-laden character of scientific reasoning by arguing that contextual factors influence a much earlier stage of research than the interpretation of data, namely the characterization of data (Douglas, 2000, pp. 569–572). The data relevant to the study on the effects of dioxins was produced via examinations of slides of rat livers exposed to different levels of dioxins. The rat liver slides are themselves highly ambiguous, and different groups of scientists examining the very same slides produced very different data sets. In some cases, scientists working in the same groups could not agree upon whether a particular slide indicated a tumor; in these cases, the groups produced their data points on the basis of majority votes. Douglas argues that, in this research area, contextual factors play an inevitably role in the characterization of data. This example, however, provides the beginnings of an argument for permanent, rather than merely transient underdetermination; because I am restricting the scope of this paper to transient underdetermination, I will not emphasize it here.

4. On separating the epistemic and the practical

The most common and most forceful objection to ATU states that ATU conflates two different domains that can and should remain separate: the epistemic and the practical. According to this objection, while contextual factors such as moral and political values play a legitimate role in the practical evaluation of researchthat is, in the evaluation of how a given area of research should lead us to *act*—they can and should be excluded from the *epistemic* evaluation of research. One version of this objection has been given by Sandra Mitchell in direct response to Douglas's argument. A different formulation of this same objection is given by Richard Jeffrey in direct response to Rudner's argument from inductive risk. In this section, I will examine the different ways in which Jeffrey and Mitchell spell out this objection, and I will discuss the difficulties that both formulations face in providing us an account of how contextual factors can always be excluded from the epistemic appraisal of transiently underdetermined research.⁶ On the leffreyan-and, more broadly, Bayesian-account, it is simply impossible to screen out all contextual factors from the epistemic appraisal of transiently underdetermined research. On this account, it might be possible to maintain the ideal of epistemic purity in the long run of inquiry, but it is impossible to do so in transiently underdetermined contexts. On the Mitchell account, the situation is not so straightforward. It is possible, on this account, to maintain the ideal of epistemic purity even in transiently underdetermined contexts; doing so, however, requires that one pay a high price.

4.1. The Jeffrey formulation

Jeffrey attempts to undermine Rudner's argument by questioning one of its central premises, and from there he attempts to explicate how a distinction between the theoretical and the practical can be drawn (Jeffrey, 1956). The premise to which Jeffrey objects is P1: the scientist qua scientist accepts or rejects hypotheses. On Jeffrey's view, the scientist qua scientist does not accept or reject hypotheses but merely assigns probabilities to them. In particular, scientists should assign initial prior probabilities to hypotheses and, when new evidence arises, update these probabilities via Bayesian conditionalization. Values, Jeffrey believes, play no part in this process. Of course, values play an important part in determining how to act on the basis of these probabilities. Values, or utilities, must be assigned to the possible outcomes of the actions under consideration, and then a decision-theoretic calculation is undertaken to determine the optimal course of action. But in Jeffrey's view, values are restricted to the decision of how to act-a decision that does not fall within the domain of the scientist qua scientist but of the policy maker (or the scientist qua policy maker).

In his 1953 paper, Rudner anticipates and responds to Jeffrey's objection (Rudner, 1953, pp. 3–4). Rudner argues that, even on Jeffrey's account, one cannot avoid inductive risk; all that Jeffrey has done is to shift the inductive risk back one step, to accepting or rejecting that a given hypothesis has a given probability, as opposed to accepting or rejecting the hypothesis as true or false. Thus, if we include as hypotheses statements of the form *hypothesis H* has probability *p*, then Rudner's argument, as laid out in the previous section, is untouched by Jeffrey's objection.⁷

But perhaps one is not satisfied with this response; in this case, let us examine Jeffrey's objection in more detail. If Jeffrey's argument is to undermine the ideal of epistemic purity, it must establish that Bayesian conditionalization provides a means for generating probabilities in a manner free from contextual factors. I will represent Bayes's Theorem in simplified fashion:

$$P(H|E) = \frac{P(H)P(E|H)}{P(E)}$$

If his argument is to constitute a strong objection to ATU, one of the following must be true: (1) all of the terms on the right hand side of the equation—the likelihood of the hypothesis *H*, the prior probability of *H*, and the expectedness of the evidence *E*, can always be assigned in a manner free from contextual factors, or (2) if they cannot, the contextual factors involved in their assignment will always wash out in cases of transient underdetermination. It seems highly unlikelv that either the likelihood of H or the expectedness of E could be always be assigned in a manner free from contextual factors. Consider, again, in the dioxin research discussed earlier: neither the threshold hypothesis nor the no-threshold hypothesis appear to entail anything about which evidence we will find; the expectedness of the evidence seems even more difficult to specify. However, let us assume for the sake of argument that the likelihood of H and the expectedness of E can always be assigned in a manner free from contextual factors and focus our attention upon the assignment of the prior probability of H. In determining whether initial prior probabilities can be assigned in a manner free from contextual factors, or whether these factors will always wash out in the context of transiently underdetermined research, it will be necessary to distinguish between two broad conceptions of probability: subjectivism and objectivism.

Suppose that we conceive of probability along subjectivist lines, such that probabilities are nothing more than subjective degrees of belief. On this view, there is only one constraint upon the assignment of initial prior probabilities: internal consistency. As a result, a scientist is free to allow any sort of consideration—subjective preference, ethical or political value, etc.—to determine the assignment of initial priors (assuming, again, that her probability assignments are all internally consistent).

In another context, Wesley Salmon has criticized subjectivism, or personalism, for allowing such a broad swath of considerations to affect prior probability assessment:

The frightening thing about pure unadulterated personalism is that nothing prevents prior probabilities (and other probabilities as well) from being determined by all sorts of idiosyncratic and objectively irrelevant considerations. A given hypothesis might get an extremely low prior probability because the scientist considering it has a hangover, has had a recent fight with his or her lover, is in passionate disagreement with the politics of the scientist who first advanced the hypothesis, harbors deep prejudices against the ethnic group to which the originator of the hypothesis belongs, etc. (Salmon, 1990, p. 183).

Though Salmon clearly finds this consequence problematic, subjectivists will defend their view by noting that, over time, the influence of these "idiosyncratic and objectively irrelevant considerations" will diminish and eventually disappear. Though in the beginning, scientists' probability assessments might diverge radically from one another as a result of these contextual factors, the acquisition of more and more evidence and the updating of probabilities on the basis of this evidence will lead to a convergence of probabilities. Thus, while subjectivists have no basis for excluding contextual factors from affecting the assignment of initial prior probabilities, we might hold out hope that such factors will wash out over time.

⁶ See Biddle & Winsberg (2010) for a response to this objection in the area of climate modeling.

⁷ This objection is also discussed by Douglas (2000, 2009).

Suppose, alternatively, that we were to reject subjectivism and defend a form of objectivism about probability, as Salmon does. In his well-known attempt to marry logical empiricism with Kuhnian historicism, Salmon explicates confirmation in science in terms of Bayesian conditionalization and demands that initial prior probability assessment be constrained by three kinds of "objective" criteria: pragmatic criteria, formal criteria, and material criteria. Pragmatic criteria "have to do with the circumstances in which a new hypothesis originates" (Salmon, 1990, p. 184). They take into account, for example, that scientists are fallible and that most hypotheses are unsuccessful; as a result, no hypothesis should be given an extremely high prior probability. Furthermore, hypotheses advanced by "serious scientists" who are working in their fields of expertise tend to be much more likely to succeed than hypotheses advanced by "scientific cranks" or scientists working far outside their fields (Salmon, 1990, p. 184). Formal criteria pertain to the consistency of scientific theories or hypotheses, both internally and externally (i.e., with surrounding scientific theories, hypotheses, or beliefs that are widely accepted). Hypotheses that fail the tests of internal or external consistency should be given a relatively low prior probability assessment. Finally, material criteria "have to do with the actual structure and content of the hypothesis or theory under consideration" (Salmon, 1990, p. 185). Simplicity, symmetry, and analogy are examples of these criteria.

While many of Salmon's criteria are quite reasonable, it is not at all obvious that all are truly objective. For example, it is hard to see how one can explicate the notions of "serious scientists" and "scientific cranks" in a completely objective fashion. Of course, past scientific success is relevant to this evaluation, but it isn't the entire story; if it were, one would be unable to assign prior probabilities to hypotheses developed by young scientists with little track record. Moreover, if one examines the actual practice of science, it is clear that the attribution of crank status to scientists can have a significant political dimension; for example, in scientific controversies that are transiently underdetermined, it is not uncommon for scientists to label opposing scientists as cranks, frauds, or incompetent—not because they really are, but because they weigh or interpret evidence differently, extrapolate evidence differently, or employ different criteria for judging a hypothesis (e.g., Ruse, 2005).

Similarly, it is also far from obvious that the criteria of external consistency and simplicity can be construed as completely objective, in the sense of involving no contextual factors. As Kuhn emphasized, both of these criteria can be interpreted differently (Kuhn, 1977). Many hypotheses are consistent with some neighboring hypotheses or beliefs and inconsistent with others; different scientists will legitimately hold different views about whether a hypothesis satisfies the criterion, depending upon which neighboring hypotheses or beliefs they emphasize. And simplicity is notoriously difficult to characterize, so much so that many have suggested that there is an ineliminable subjective element to it (cf. Quine & Ullian, 1970). Salmon himself acknowledges that the criterion of simplicity cannot be applied across the board—that it "varies from one scientific context to another" (Salmon, 1990, p. 186). He states that scientists can determine when this criterion is applicable on the basis of "training and experience," though no argument is given for this or for how "training and experience" are always sufficient to exclude all contextual factors. Longino, moreover, argues that the choice of simplicity as an epistemic value has been done for socio-political reasons (Longino, 1996).

However, even if we ignore all of these considerations and assume that his pragmatic, formal, and material criteria are truly objective, there is still significant room for subjective factors to operate, because these criteria are insufficient to determine prior probability assessment. Salmon himself acknowledges this:

One point is apt to be immediately troublesome. If we are to use **Bayes's theorem to compute values of posterior probabilities**, it would appear that we must be prepared to furnish numerical values for the prior probabilities. Unfortunately, it seems preposterous to suppose that plausibility arguments of the kind we have considered could yield exact numerical values (Salmon, 1990, p. 187).

Because objective factors are insufficient to determine prior probability assessment, other kinds of factors will inevitably play a role. Salmon argues that these factors will not compromise epistemic integrity, however, because they will, again, wash out over time:

[B]ecause of a phenomenon known as 'washing out of the priors' or 'swamping of the priors,' even very crude estimates of the prior probabilities will suffice for the kinds of scientific judgments we are concerned to make (Salmon, 1990, p. 187).

Thus, like subjectivism, Salmon's objectivism does not allow for assessing initial prior probabilities in a manner free from contextual factors, but it does provide hope that such factors will be screened out over time,⁸

These brief discussions of subjectivism and objectivism are sufficient to allow us to assess the prospect of the Jeffrey formulation to respond to ATU. Both accounts succeed only on the condition that there is sufficient evidence to yield a convergence in the outcomes of Bayesian conditionalization. Both allow contextual factors to influence the setting of initial prior probabilities, and both argue that at the end of the day, when sufficient evidence has been gathered, these contextual factors will be screened out. Yet, in areas of transiently underdetermined research, we are by definition not at the end of the day; we are rather in the very middle of it. In these situations, contextual factors will have a significant influence upon the outcomes of Bayesian conditionalization, with the potential result of a wide divergence in probability assessments. Thus, on the Bayesian account, it is *impossible* to screen out all contextual factors from the epistemic appraisal of transiently underdetermined research.

As a brief illustration, let us return to a slightly idealized version of the dioxin case discussed earlier. In the debate between those who defend the threshold model and those who defend the nothreshold model, both sides have access to the same evidence, and yet they assign very different probabilities to each respective **model.** Suppose that half of the community (call this half C_1) assigns a probability of .75 to the threshold model (M_T) and .25 to the no-threshold model (M_{N-T}) , and suppose that the other half of the community (call it C_2) assigns a probability of .25 to M_T and .75 to M_{N-T} . This is clearly an idealization, but not an extreme one; an examination of the controversy makes clear that each side assigns a fairly high probability to its preferred model and a fairly low probability to the other. Suppose, furthermore, that the differences in probability assessments between C_1 and C_2 are not due to one side simply making a mistake or to committing fraud or misconduct; rather, this is a genuine case of transient underdetermination. This is not an idealization. On the Jeffreyan account, the most plausible explanation of the very significant difference between probability assessments is that each side assigns to the models very different initial prior probabilities (and perhaps also to the likelihoods and expectedness of the evidence). Given the argument of this section, the most plausible explanation of the differences in

⁸ Salmon's objectivism is not the only form of objectivism. Regarding the influence of contextual factors, however, I assume that the best that other forms will be able to offer is to screen out such factors in the long run, and that they will fare no better than Salmon's in providing a way of screening them out in transiently underdetermined contexts.

prior probability assessments is the influence of different contextual factors. All of this shows that, in the Jeffreyan account, contextual factors not only influence probability assessment in transiently underdetermined research, but that this influence can be very significant.⁹

4.2. The Mitchell formulation

Mitchell provides an alternative formulation of an objection to ATU. Similar to Jeffrey, she argues that while contextual factors inevitably influence the practical evaluation of research, they can and should be screened off from epistemic evaluation. In response to Douglas's argument from inductive risk, Mitchell writes that the argument involves a "conflation of the domains of belief and action [that] confuses rather than clarifies the appropriate role of values in scientific practice" (Mitchell, 2004, p. 250). Mitchell acknowledges that some scientists, as a result of the particular role that they play as scientific advisors for policy-making, must consider questions of both belief and action. Scientists involved in policymaking have different sets of obligations that arise from their different institutional roles, namely their roles as scientists and as government advisors. However, these roles-and the norms that are distinct to each—must be kept separate; if they are not, she argues, we will have no basis upon which to protect against the distortion of science by prevailing political interests (Mitchell, 2004, p. 251). Thus, while the scientist qua policy advisor should bring ethical considerations to bear on her research, Mitchell maintains that the scientist qua scientist can and should exclude such considerations from her reasoning.

Mitchell's account of how scientists can separate the epistemic from the practical is not Bayesian. Instead, she distinguishes between the aims of truth and warranted acceptability and argues that scientists, by applying epistemic values, can generate warranted beliefs without the introduction of contextual factors. "While evidence underdetermines the truth of a causal claim, the other epistemic values... or cognitive values... can be harnessed to generate a judgment of acceptance" (Mitchell, 2004, p. 249). Epistemic values are rational bases for preferring a theory or hypothesis, which transcend narrow evidential support; they include predictive accuracy, problem-solving ability, breadth of scope, and simplicity (Mitchell, 2004, p. 249). This list of epistemic values is similar to those provided by Kuhn (1977), McMullin (1983), and Laudan (1984). Mitchell argues that epistemic values are sufficient to close the gap between logic and evidence, on the one hand, and the warranted acceptance of hypotheses, on the other; because of this, contextual factors need play no role.

Mitchell's argument is open to an objection discussed briefly in the previous sub-section. Many epistemic values are open to different interpretations. Clearly, simplicity can be, and has been, understood in many different ways. Moreover, hypotheses are often consistent with some neighboring theories and inconsistent with others; because of this, judgments of external consistency often depend upon which neighboring theories or hypotheses one emphasizes. Similarly, judgments of problem-solving ability typically depend upon which problems one takes to be important and which problems one is willing to ignore. In addition to these problems of interpretation, epistemic values can often conflict with one another, with the result that some values must be prioritized over others; the values of simplicity and breadth of scope, for example, are often in tension with one another. Kuhn himself emphasizes both the vagueness of epistemic values and the need to prioritize them with respect to each other, and he argues that in many cases, different interpretations and weighting schemes lead to very different assessments of theories. As a result, he argues, theory choice is still underdetermined, not just by logic and evidence narrowly construed, but also by logic, evidence, and a vaguely construed set of epistemic values (Kuhn, 1977).¹⁰

Mitchell is aware of this objection, and she raises and responds to it in her critique of Douglas's argument:

Employing these additional [epistemic] norms does not *uniquely* determine even the acceptability of a claim. As Kuhn pointed out early on, individuals might prioritize different members of this set of values or interpret their application in different ways (Mitchell, 2004, p. 249, her emphasis).

Despite this, she still argues that contextual factors can and should be screened off from the epistemic appraisal of research. Immediately following the passage quoted above, she writes: "Nevertheless, if the scientists in the advisory committees to policymakers make judgments employing only these broadly epistemic or cognitive values, then there is no necessity for values outside this set to enter the process," While this sentence is in the form of a conditional, she clearly believes that the antecedent can be satisfied. For example, she writes elsewhere: "Scientists can certainly disagree about the risk of a potential carcinogen when they disagree about the significance of different epistemic values or the degree of their satisfaction—without nonepistemic values entering in" (Mitchell, 2004, p. 253). Thus, even when research is underdetermined by evidence and epistemic values (vaguely construed), she still maintains that contextual factors, or "non-epistemic values," can and should always be screened off.

In assessing Mitchell's argument, let us consider, again, the debate between those who defend M_T and those who defend M_{N-T} . As we have seen, evidential considerations, narrowly construed, cannot decide between them, as both are consistent with the available evidence. In the previous sub-section, we saw that the most probable explanation for the disagreement on a Bayesian account would include differences in initial prior probability assignment, which are inevitably influenced by contextual factors. On Mitchell's account, the explanation would be different; the controversy is due to the competing sides of the debate, C_1 and C_2 , adopting different interpretations and/or weightings of epistemic values; contextual factors, she argues, can and should always be screened off from this. Suppose that the epistemic values relevant to this dispute are simplicity (*S*) and problem-solving ability (PA) and that there are different interpretations of both of these values $((S_1, S_2, ..., S_n)$ and $(PA_1, PA_2, ..., S_n)$ PA_n)) on offer. (One can enlarge the set of values relevant to the dispute as far as one likes; doing so will not affect the outcome of the argument.) Suppose, furthermore, that C_1 adopts the weighted set of interpreted values (S_2 , PA_1) in support of M_T and that C_2 adopts the weighted set of interpreted values (PA_2 , S_1) in support of M_{N-T} . At this point, we still have the question of why each group chose the weighted set of interpreted values that it did. Can we answer this question by appealing only to epistemic values? No-and Mitchell

⁹ Norton (2008) argues that the thesis of global underdetermination is problematic because it assumes "an impoverished account of induction" (Norton, 2008, p. 17). He discusses three less-impoverished accounts—inductive generalization, hypothetical induction, and probabilistic accounts—and argues that none supports the thesis of global underdetermination. He does not address the issue of whether they provide a way of excluding contextual factors from all instances of transiently underdetermined research. It seems clear, however, that they do not. The preceding discussion illustrates why probabilistic accounts of induction cannot perform this task. "Exclusionary accounts" of hypothetical induction cannot, because there are many instances of transient underdetermination (such as the one discussed by Douglas) in which it is not the case that, if *H* were false, *E* would very likely not obtain. Determining whether inductive generalization could perform this task would require a more precise specification of the account in question; it is difficult, however, to see how any could do so satisfactorily.

¹⁰ This has led many to question the cogency of the distinction between epistemic and non-epistemic values. See, for example, Douglas (2009), Howard (2006), Longino (1990, 1996, 2002), and Rooney (1992).

herself has acknowledged as much. For if we could justify $[S_2, PA_1]$, for example, solely in terms of epistemic values, then the appeal to epistemic values would settle the controversy in favor of M_T . But Mitchell has already acknowledged that appeals to epistemic values cannot, in general, uniquely determine theory or model choice (Mitchell, 2004, p. 249). Given that we cannot justify the choice of one weighted set of interpreted epistemic values solely in terms of epistemic values, the explanation of why this set was chosen over the others must appeal to non-epistemic, or contextual, factors.

One might argue that the question of why one group chose the weighted set of interpreted epistemic values that it did is not a legitimate question—or, in any case, that it is not a question that the defender of the ideal of epistemic purity needs to answer. After all, epistemic inquiry must begin somewhere; why can it not simply begin with a weighted set of interpreted epistemic values? It is true that inquiry must begin somewhere, and it is certainly not the case that scientists must always seek to provide firm rational foundations for the epistemic standards that they employ. At the same time, the ideal of epistemic purity requires that scientists always be able to screen out contextual factors from the epistemic evaluation of research, and part of the epistemic evaluation of research is the choice of epistemic standards. Thus, if research is to be epistemically pure, it must be the case that the epistemic standards to be employed can be chosen in a manner free from contextual factors. In many areas of transiently underdetermined research, however, epistemic standards cannot be chosen in such a manner; the example of dioxin research is just one example of this.

Like the Bayesian account, the objection provided by Mitchell might offer a way to screen out contextual values in the long run of inquiry. It might be that, once sufficient evidence has been gathered, *any* interpretation of, and *any* weighting scheme for, epistemic values might lead to the same choice. In this case, one would not need to decide between competing interpretations and weighting schemes. But again, the question at issue here is whether the ideal of epistemic purity succeeds in the context of transiently underdetermined research.

The Mitchell formulation does allow for the *possibility* of screening out all contextual factors from transiently underdetermined research. Defenders of this formulation could argue that, in the long run of inquiry, contextual factors can be screened out via the application of epistemic values, and scientists should simply remain agnostic, with respect to belief, until this has occurred. In transiently underdetermined contexts, different scientists apply different interpretations of epistemic values, or weight epistemic values differently; in these contexts, scientists should remain completely agnostic with respect to the epistemic appraisal of the hypotheses or models in question. It is only when underdetermination vanishes—for example, when all interpretations and weightings of epistemic values lead to the same choice (assuming that this ever happens)—that scientists should accept these hypotheses or models.

The strategy of remaining agnostic in the face of underdetermination does allow for the possibility of epistemically pure science, but it does so at significant cost. This strategy, if employed, would represent a radical deviation from the way in which science is actually practiced. Scientists do not remain agnostic about a theory or hypothesis until there is overwhelming evidence for it. As Kuhn (1962) and others have argued, if there is such a thing as definitive evidence for a theory, that evidence is typically not acquired until long after the scientific community has accepted that theory; some of the most celebrated instances of theory change, including the acceptances of the heliocentric model and Newtonian mechanics, proceeded in this way. Do we really want to fault the scientists who initiated these revolutions and who accepted their theories before there was definitive evidence? We could, of course, bite this bullet and argue that, if the ideal of epistemic purity is violated in the actual practice of science, even in the most celebrated episodes, then so much the worse for the actual practice of science. This, however, is a high price to pay; it seems reasonable to think that our best accounts of scientific reasoning should allow us to see the great episodes of the actual practice of science for what they are: great episodes.

There are also epistemic costs to the strategy of remaining agnostic in the face of underdetermination (Biddle, 2009). As Kuhn (1962), Lakatos (1970), and others have argued, scientific progress often requires that scientists advocate and defend the paradigms or research programs within which they work.¹¹ Progress requires that scientists attempt to work out a paradigm or research program, overcome anomalies and effect a fit with nature. On these accounts, the typical attitude of scientists is not one of agnosticism but rather of commitment to the paradigm or research program within which they work-and this commitment is epistemically beneficial. If scientists were truly agnostic about a theory-if they were, for example, just as likely to respond to anomalies by seeking a new theory as by attempting to overcome the anomalies-scientific progress would slow, if not grind to a halt. In this case, the prescription to remain agnostic until overwhelming evidence has been acquired might very well inhibit the very attitudes that typically lead a community to acquire overwhelming evidence at all.

The discussion of Section 4 has shown that, on the Bayesian account, it is *impossible* to screen out all contextual factors from the epistemic appraisal of transiently underdetermined research, while on the Mitchell account, it is possible though costly to do so. My own view is that these costs are too high to pay, and as a result, I conclude that, on the Mitchell formulation, it is *unreasonable* to screen out all contextual factors from the epistemic appraisal of transiently underdetermined research. Still, this formulation does provide a possible way of defending the ideal.¹²

5. On ideals, relativism, and underdetermination

One might attempt to defend the ideal of epistemic purity by noting that it is just that—an ideal—and by arguing that even though the ideal is unattainable in transiently underdetermined contexts, we should still attempt to come as close as possible to achieving it. There are a couple of points that need to be made in response to this objection. Firstly, the prescription to follow the evidence as far as it leads is shared by almost all critics of the ideal of epistemic purity. While there are a few who appeal to arguments from underdetermination to support a radical relativism, according to which nature provides virtually no constraints on scientific reasoning, most philosophically-oriented critics do not¹³ They argue that, although logic and evidence do not determine theory choice uniquely, they constrain it significantly; the influence of contextual factors should be minimized, so as to allow evidential considerations as decisive a role as possible (Frank, 1954; Howard, 2006; Longino, 1990; Longino, 2002; Neurath, 1983 [1913]).¹⁴ This

¹¹ Even Feyerabend (1970) allows a role for what he calls the "principle of tenacity."

¹² Some have argued that establishing that it is either impossible or unreasonable to screen out all contextual factors from the epistemic evaluation of research is insufficient to undermine the "ideal of value-free science" and that, in addition, it is necessary to show that contextual factors play a *legitimate* role in the epistemic evaluation of research (e.g., Intemann, 2005). However, if it is impossible or unreasonable to screen out all contextual factors from scientific inquiry, then assuming that we continue doing science, we must acknowledge them as legitimate.

¹³ The early Harry Collins is an example of a defender of such a radical relativism: "Our school... embraces an explicit relativism in which the natural world has a small or nonexistent role in the construction of scientific knowledge" (Collins, 1981, p. 3). It should be noted, however, that most of the sociologists of science who previously held such views—including Collins—have adopted more moderate positions of late (e.g., Collins & Evans, 2007).

¹⁴ Koertge, for example, asserts that defenders of underdetermination-style arguments believe that "science has no deserved epistemic authority" (2003, 225).

point is missed by some defenders of the ideal, such as Koertge (2003) and Pinnick (2003). The ideal of epistemic purity, however, states not only that we should strive to minimize the influence of contextual factors, but also that the proper application of scientific methods will always result in our ability to screen them out entirely. Both Jeffrey and Mitchell maintain this formulation of the ideal, and while many contemporary defenders of the ideal do not define the ideal precisely, most if not all of them take the ideal to imply that the proper application of scientific methods will, as a matter of fact, always screen out all contextual factors (e.g., Ruphy, 2006, p. 190). Those who deny the ideal of epistemic purity maintain that, in transiently underdetermined contexts, even properly applied scientific methods will not screen out all contextual factors. In these contexts, then, the ideal is potentially problematic

We might consider reformulating the ideal as follows: we ought to minimize the influence of contextual factors in the epistemic appraisal of research, even though, in many contexts, we will not be able to eliminate them. While it might be true that we ought to strive to minimize the influence of contextual factors, there are reasons against calling this an ideal of epistemic purity. The first is the well-known dictum that *ought* implies *can*; to claim that we ought to strive to achieve epistemic purity is to imply that we can achieve epistemic purity; if epistemic purity is impossible or unreasonable to achieve, however, then we should not maintain that we ought to strive to achieve it. On this account, the ideal of minimizing the influence of contextual factors is a legitimate one; the ideal of epistemic purity is not.

Furthermore, as many critics of the ideal of epistemic purity have emphasized, the ideal can have the dangerous consequence of masking the influence of contextual factors (e.g., Longino, 1990). In many areas of research, such factors have an influence throughout the research process, including the characterization of data, the choice of methodologies (including the choice of a level of statistical significance), the choice of models and research subjects, and so on. Allowing some factors rather than others to have an influence can systemically skew research in such a way as to raise the probability of achieving a desired outcome (Wilholt, 2009). Over the past thirty years at least, we have seen such tactics employed by the pharmaceutical industry, the chemical industry, and others in order to make their products appear safer and more effective than they really are (e.g., Biddle, 2007; Brown, 2008). Denying the ideal of epistemic purity in transiently underdetermined contexts calls attention to the fact that contextual factors will likely be operating, thus emphasizing the need to scrutinize those factors to ensure that they are not operating in unethical ways.

6. Value freedom and epistemic purity

While the preceding discussions illustrate a number of the difficulties that defenders of the ideal of epistemic purity face, they also suggest some potential new directions for the field of "science and values." One of these concerns the language of "values" that is often employed in the debate. As noted earlier, I have been discussing an ideal that I call the "ideal of epistemic purity," rather than the "ideal of value free science." I have made this terminological change because there are many contextual factors that might influence the epistemic appraisal of research, not just values. Consider, again, Salmon's criticism of subjectivism, in which he argues that subjectivists have no basis for excluding "objectively irrelevant considerations" from influencing prior probability assessment, and consider one such objectively irrelevant consideration—having a hangover. Clearly, defenders of the "ideal of value-free science"—and defenders of most any ideal for science, for that matter—wish to exclude factors such as having a hangover from influencing the appraisal of theories; and yet it is equally clear that having a hangover is not a value. Many critics of the "ideal of value-free science" wish to argue that ethical or political values can, in certain situations, legitimately influence theory or hypothesis appraisal; at the same time, most of these critics maintain that, in many situations, contextual factors that are not accurately described as values will invariably influence theory choice. None of this is to say that the question of the appropriate role of values in science isn't important; it most definitely is. But the philosophical issues at the heart of the "science and values" debate extend beyond this question, and the terminology employed in this debate should reflect this fact.

In this regard, it is perhaps helpful to note that one of the one of the earliest critics of what is now called "the ideal of value-free science," Otto Neurath, did not employ the term "values" at all.¹⁵ In "The Lost Wanderers of Descartes and the Auxiliary Motive," Neurath discusses Descartes' treatment of theoretical and practical reasoning in the Discourse on Method, Descartes imagines a wanderer lost in the forest and argues that, in the absence of complete insight regarding how to get out of the forest, the wanderer should simply pick a direction and walk, without reconsidering the initial decision; in this case, the wanderer is assured of getting out of the forest eventually (assuming that the forest is not too large), avoiding the possibility of continually walking in circles. Descartes argues that, in the case of practical reasoning, insight is incomplete and incapable of determining uniquely a course of action; as a result, we should adopt practical guidelines for how to proceed in the face of partial ignorance. Theoretical reasoning, he argues, is different; in that sphere, complete knowledge is possible, and we should only proceed on the basis of certain knowledge.

Neurath argues that Descartes' characterization of theoretical reasoning is mistaken, and that, in theoretical reasoning just as in practical reasoning, "insight" is insufficient to determine uniquely what to believe: some other factors, which he called "the auxiliary motive," will inevitably play a role. The auxiliary motive refers to any factor that fills the gap between "insight" and decision; some of these factors might properly be characterized as values, while others clearly are not. One example that Neurath gives of an auxiliary motive is tradition; in this case, one chooses the result that is compatible with both insight and those norms and beliefs that are characteristic of the time. This motive might be properly viewed as a value-although, again, Neurath does not use that term. Other examples that he provides are drawing lots and merely doing "something or other," waiting to see "which resolution, after some hesitation will come out on top, as if leaving the decision to exhaustion" (Neurath, 1983 [1913], pp. 4-5). Clearly, these motives are not properly characterized as values. Neurath's point is not to dictate which motives should operate in a given circumstance but rather to emphasize that such an auxiliary motive is present. He refers to those who pretend as if they are making decisions on the basis of insight alone, and who either mask or ignore the presence of the auxiliary motive, as "pseudorationalists." True rationalism, in contrast, "sees its chief triumph in the clear recognition of the limits of actual insight" (Neurath, 1983 [1913], p. 8).¹⁶

One of the lessons that can be drawn from Neurath's work for the current discussion of "science and values" is that there are a

¹⁵ For further discussion of Neurath's views on this matter, see Howard (2006).

¹⁶ Miriam Solomon is an example of a contemporary scholar who eschews the term "values," employing instead the notion of a "decision vector," which is any factor (value, preference, etc.) that influences a decision. In this respect, her terminology approaches Neurath's. Solomon goes further, however, by distinguishing between empirical decision vectors, which are "causes of preference for theories with empirical success, either success in general or one success in particular," and non-empirical decision vectors, which are "other reasons or causes for choice" (Solomon, 2001, 56). The holism defended by Neurath, and later Quine (1951), would preclude distinguishing so sharply between the empirical and the non-empirical.

wide range of different factors that can fill the gap between "insight" (i.e., logic, evidence, and epistemic values broadly construed) and decision making in science. These factors might be sociological in nature (e.g., in one particular scientific sub-discipline, one set of norms are typically employed for evaluating research, as opposed to some other set, which are employed in some other sub-disciplines); they might be consciously adopted ethical or political values; they might be unconsciously held subjective preferences or ideological assumptions, and so on. Some of these factors, again, are properly described as values, while others are not. One of the tasks for philosophers of science working in this area should be to develop a more fine-grained understanding of the kinds of factors that can operate here, so as to specify more clearly which factors play a legitimate role in science and which do not (cf. Douglas, 2009).

7. Conclusion

This paper has examined the state of the field of "science and values"-particularly regarding the implications of the thesis of transient underdetermination for the ideal of value-free science, or what I have called the "ideal of epistemic purity." I have discussed a preliminary argument from transient underdetermination against the ideal, and I have examined two formulations of an objection to this argument. On the Jeffreyan-and, more broadly, Bayesian-formulation, it is impossible to screen out all contextual factors from the epistemic appraisal of transiently underdetermined research. The Mitchell formulation, on the other hand, allows for the possibility of screening out all contextual factorsbut it does so at significant cost. A secondary aim of the paper has been to suggest a new direction for the field of "science and values." The most fundamental issue in the debate over the ideal of epistemic purity is not really whether science can or should be value free but rather whether it can or should be free from all contextual factors. On many accounts of scientific reasoning (for example, those provided by both Jeffrey and Mitchell), it is quite possible for the epistemic appraisal of scientific research-even transiently underdetermined research-to be free from values properly understood (e.g., moral and political values). The situation is rather different, however, with respect to contextual factors. Given this, an important future project for the field would be to clarify and make more precise the central concepts employed within the field, so as to specify more clearly which factors play a legitimate role in science and which do not.

The debate over the ideal of epistemic purity brings a number of additional questions to the forefront. If the ideal of epistemic purity is inadequate, then which ideal should take its place? Competing answers to this question have been given by a number of different commentators, including Douglas (2009), Elliott (2011a, 2011b), Longino (1990, 2002), Kourany (2010), and Solomon (2001). Relatedly, if it is either impossible or inadvisable to attempt to screen out all contextual factors from transiently underdetermined research, which kinds of contextual factors should be allowed, and which should we seek to exclude? A number of philosophers of science have begun to address these questions, but much more remains to be done.

Acknowledgements

I am indebted to many people for their comments on earlier versions of this paper. Thanks especially to Matt Brown, Martin Carrier, Heather Douglas, Kevin Elliott, Don Howard, Paul Hoyningen-Huene, Philip Kitcher, Janet Kourany, Ulrich Krohs, Cornelis Menke, Eric Oberheim, Torsten Wilholt, David Willmes, Eric Winsberg, and two anonymous reviewers.

References

Biddle, J. (2007). Lessons from the Vioxx debacle: What the privatization of science can teach us about social epistemology. Social Epistemology, 21(1), 21–39.

- Biddle, J. (2009). Advocates or unencumbered selves? On the role of Mill's political liberalism in Longino's contextual empiricism. *Philosophy of Science*, 76, 612–623.
- Biddle, J., & Winsberg, E. (2010). Value judgements and the estimation of uncertainty in climate modeling. In P. D. Magnus & J. Busch (Eds.), *New waves in philosophy of science* (pp. 172–197). Basingstoke, Hampshire: Palgrave MacMillan.
- Brown, J. R. (2008). The community of science[®]. In M. Carrier, D. A. Howard, & J. Kourany (Eds.), *The challenge of the social and the pressure of practice: Science and values revisited* (pp. 189–216). Pittsburgh: University of Pittsburgh Press.
- Collins, H. (1981). Stages in the empirical programme of relativism. Social Studies of Science, 11, 3–10.
- Collins, H., & Evans, R. (2007). Rethinking expertise. Chicago: University of Chicago Press.
- Douglas, H. (2000). Inductive risk and values in science. *Philosophy of Science*, 67, 559–579.
- Douglas, H. (2009). Science, policy, and the value-free ideal. Pittsburgh: University of Pittsburgh Press.
- Duhem, P. (1954 [1906]). The aim and structure of physical theory. Princeton: Princeton University Press. (Translation of La théorie physique: Son objet, sa structure. Paris: Marcel Rivière & Cie, 1906.)
- Dupré, J. (2007). Fact and value. In H. Kincaid, J. Dupré, & A. Wylie (Eds.), Value-free science? Ideals and illusions (pp. 21–41). Oxford: Oxford University Press.
- Elliott, K. (2011a). Direct and indirect roles for values in science. Philosophy of Science, 78, 303–324.
- Elliott, K. (2011b). Is a little pollution good for you? Incorporating societal values in environmental research. New York: Oxford University Press.
- Feyerabend, P. (1970). Consolations for the specialist. In I. Lakatos & A. Musgrave (Eds.), Criticism and the growth of knowledge (pp. 197–230). Cambridge: Cambridge University Press.
- Frank, P. G. (1954). The variety of reasons for the acceptance of scientific theories Idem. In *Validation of scientific theories* (pp. 3–17). Boston: Beacon Press.
- Giere, R. N. (2003). A new program for philosophy of science? Philosophy of Science, 70, 15–21.
- Howard, D. A. (2009). Better red than dead—Putting an end to the social irrelevance of postwar philosophy of science. *Science and Education*, *18*, 199–220.
- Howard, D. A. (2003). Two left turns make a right: On the curious political career of North American philosophy of science at midcentury". In G. Hardcastle & A. Richardson (Eds.), Logical Empiricism in North America (pp. 25–93). Minneapolis: University of Minnesota Press.
- Howard, D. A. (2006). Lost wanderers in the forest of knowledge: Some thoughts on the discovery-justification distinction. In J. Schickore & F. Steinle (Eds.), *Revisiting discovery and justification: Historical and philosophical perspectives on the context distinction* (pp. 3–22). Dordrecht: Springer.
- Jeffrey, R. C. (1956). Valuation and acceptance of scientific hypotheses. Philosophy of Science, 22, 237–246.

Kitcher, P. (2001). Science, truth, and democracy. New York: Oxford University Press.

- Kitcher, P. (2011). Science in a democratic society. New York: Oxford University Press.
- Koertge, N. (2000). Science, values, and the value of science. *Philosophy of Science*, 67, S45–S57.
- Koertge, N. (2003). Feminist values and the value of science. In C. Pinnick, N. Koertge, & R. Alemander (Eds.), Scrutinizing feminist epistemology: An examination of gender in science (pp. 222–233). Rutgers, NJ: Rutgers University Press.
- Kourany, J. (2003a). A philosophy of science for the twenty-first century. Philosophy of Science, 70, 1–14.
- Kourany, J. (2003b). Reply to Giere. Philosophy of Science, 70, 22-26.
- Kourany, J. (2010). Philosophy of science after feminism. Oxford: Oxford University Press.
- Kuhn, T. S. (1962). The structure of scientific revolutions. Chicago: University of Chicago Press.
- Kuhn, T. S. (1977). Objectivity, value judgment, and theory choice. In The. Essential (Ed.), Tension: Selected Studies in Scientific Tradition and Change (pp. 320–339). Chicago: University of Chicago Press.
- Lakatos, I. (1970). Falsification and the methodology of scientific research programs. In I. Lakatos & A. Musgrave (Eds.), *Criticism and the growth of knowledge* (pp. 91–196). Cambridge: Cambridge University Press.
- Laudan, L. (1984). Science and values. Berkeley: University of California Press.
- Laudan, L., & Leplin, J. (1991). Empirical equivalence and underdetermination. Journal of Philosophy, 88, 449–472.
- Longino, H. (1990). Science as social knowledge: Values and objectivity in scientific inquiry. Princeton: Princeton University Press.
- Longino, H. (2002). The fate of knowledge. Princeton: Princeton University Press.
- Longino, H. (1996). Cognitive and non-cognitive values in science: Rethinking the dichotomy. In L. H. Nelson & J. Nelson (Eds.), *Feminism, science, and the philosophy of science* (pp. 39–58). Dordrecht: Kluwer.
- McMullin, E. (1983). Values in science. In P. D. Asquith & T. Nickles (Eds.). PSA 1982 (Vol. 2, pp. 3–28). East Lansing, Michigan: Philosophy of Science Association.
- Mitchell, S. D. (2004). The prescribed and proscribed values in science policy. In P. Machamer & G. Wolters (Eds.), *Science, values, and objectivity* (pp. 245–255). Pittsburgh: University of Pittsburgh Press.

- Nelson, L. H. (1990). Who knows? From Quine to feminist empiricism. Philadelphia: Temple University Press.
- Neurath, O. (1983 [1913]). The lost wanderers of Descartes and the auxiliary motive (on the psychology of decision). In R. S. Cohen, & M. Neurath (Eds.), *Philosophical Papers* 1913–1946 (pp. 1–12). Dordrecht: Reidel. (Translation of Die Verirrten des Cartesius und das Auxiliarmotiv. Zur Psychologie des Entschlusses. Jahrbuch der Philosophischen Gesellschaft an der Universität Wien 1913 (pp. 45–59). Leipzig: Johann Ambrosius Barth).
- Norton, J. D. (2008). Must evidence underdetermine theory? In M. Carrier, D. A. Howard, & J. Kourany (Eds.), *The challenge of the social and the pressure of practice: Science and values revisited* (pp. 17–44). Pittsburgh: University of Pittsburgh Press.
- Okruhlik, K. (1994). Gender and the biological sciences. *Biology and Society*, 20, 21-42.
- Pinnick, C. (2003). Feminist philosophy: Implications for feminist philosophy of science. In C. Pinnick, N. Koertge, & R. Alemander (Eds.), Scrutinizing feminist epistemology: An examination of gender in science (pp. 20–30). Rutgers, NJ: Rutgers University Press.
- Popper, K. (1959 [1934]). The logic of scientific discovery. London: Hutchinson. (Translation of Logik der Forschung. Tübingen: Mohr.).
- Potter, E. (1997). Underdetermination undeterred. In L. H. Nelson & J. Nelson (Eds.), Feminism, science, and the philosophy of science (pp. 121–138). Dordrecht: Kluwer.

Quine, W. V. O. (1951). Two dogmas of empiricism. *Philosophical Review*, 60, 20–43. Quine, W. V. O., & Ullian, J. S. (1970). *The web of belief*. New York: Random House.

- Reichenbach, H. (1938). Experience and prediction. Chicago: University of Chicago Press.
- Reisch, G. (2005). How the cold war transformed philosophy of science. To the icy slopes of logic. Cambridge: Cambridge University Press.
- Rolin, K. (2002). Is 'science as social' a feminist insight? *Social Epistemology*, 16, 243-249.
- Rooney, P. (1992). On values in science. Is the epistemic/non-epistemic distinction useful? In D. Hull, M. Forbes, & K. Okruhlick (Eds.). *PSA 1992* (Vol. 2, pp. 13–22). East Lansing: Philosophy of Science Association.
- Rudner, R. (1953). The scientist qua scientist makes value judgments. Philosophy of Science, 20(1), 1–6.
- Ruphy, S. (2006). 'Empiricism all the way down': A defense of the value neutrality of science in response to Helen Longino's contextual empiricism. *Perspectives on Science*, 14, 189–213.
- Ruse, M. (2005). Evolutionary biology and the question of trust. In N. Koertge (Ed.), Scientific values and civic virtues (pp. 99–119). New York: Oxford University Press.
- Salmon, W. (1990). Rationality and objectivity in science, or Tom Kuhn meets Tom Bayes. In C. W. Savage (Ed.), *Scientific theories: Minnesota studies in the philosophy of science, XIV.* Minneapolis, MN: University of Minnesota Press.
- Solomon, M. (2001). Social empiricism. Cambridge, MA: MIT Press. Wilholt, T. (2009). Bias and values in scientific research. Studies in History and Philosophy of Science, 40, 92–101.
- Wray, K. B. (1999). A defense of Longino's social epistemology. *Philosophy of Science*, 66, 538–552.