Acceptance, Values, and Inductive Risk

Daniel Steel*†

The argument from inductive risk attempts to show that practical and ethical costs of errors should influence standards of evidence for accepting scientific claims. A common objection charges that this argument presupposes a behavioral theory of acceptance that is inappropriate for science. I respond by showing that the argument from inductive risk is supported by a nonbehavioral theory of acceptance developed by Cohen, which defines acceptance in terms of premising. Moreover, I argue that theories designed to explain how acceptance can be guided exclusively by epistemic values suffer from difficulties that do not afflict Cohen's theory.

1. Introduction. Within the last 2 decades, a number of philosophers have revived the argument from inductive risk, which aims to show that nonepistemic values should influence standards of evidence for accepting a hypothesis when significant ethical or practical consequences are associated with errors (Shrader-Frechette 1991; Cranor 1993, 2006; Douglas 2000, 2009; Kourany 2003; Biddle and Winsberg 2010; Steel 2010; Elliott 2011). A common objection to this argument is that it presupposes a behavioral notion of acceptance that is inappropriate for science (Jeffrey 1956; Levi 1967; McMullin 1982; Lacey 1999; Giere 2003; Dorato 2004; Mitchell 2004). In this essay, I answer this objection by showing that the theory of acceptance developed by Jonathan Cohen (1992) supports the argument from inductive risk without interpreting acceptance in a behavioral manner. According to Cohen, to accept a proposition in a particular context is to make it available as a premise for reasoning in that context. Acceptance on this account is cognitive, not behavioral; for instance, one might accept a claim as a premise in reasoning about a theoretical issue that has no implications

*To contact the author, please write to: Department of Philosophy, 503 S. Kedzie Hall, 368 Farm Lane, Michigan State University, East Lansing, MI 48824-1032; e-mail: steel@msu.edu.

†This work was supported by the Intramural Research Grant Program at Michigan State University (08-IRGP-1448) during the 2009–10 academic year. I would also like to thank Kevin Elliott, David Willmes, and Heather Douglas for helpful comments and discussion.

Philosophy of Science, 80 (December 2013) pp. 818–828. 0031-8248/2013/8005-0029$10.00
Copyright 2013 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/673936 Published online by Cambridge University Press

for action. Nevertheless, Cohen's theory supports the argument from inductive risk because it allows that ethical or pragmatic considerations can be relevant when acceptance does have practical implications.

In addition, I argue that Cohen's theory is preferable to alternative theories of acceptance according to which acceptance should be guided solely by epistemic values. These alternatives can be divided into two groups. The first defines rational acceptance in terms of maximizing expected cognitive utility (Levi 1967; Harsanyi 1983; Maher 1993), while the second insists that certainty is a necessary condition for rational acceptance (Levi 1980; Lacey 1999). Theories in the first group claim that what we accept should have no influence on what we do. I argue that this position is highly problematic when one considers the possibility of accepting normative as well as descriptive propositions, which can then jointly function as premises in practical or ethical reasoning. Theories in the second group have the significant shortcoming of linking acceptance to dogmatism. By contrast, Cohen's theory of acceptance suffers from neither of these difficulties.

2. Inductive Risk and Acceptance. The argument from inductive risk can be succinctly formulated in the following way.

1. One central aim of scientific inference is to decide whether to accept or reject hypotheses.
2. Decisions about whether to accept or reject a scientific hypothesis can have implications for practical action, and when this happens, acceptance decisions should depend in part on nonepistemic value judgments about the costs of error.
3. Therefore, nonepistemic values can legitimately influence scientific inference.[1]

For instance, we might think for ethical reasons that it is worse to mistakenly reject the hypothesis that an artificial sweetener, such as saccharin, is carcinogenic than to make the opposite mistake. In that case, we should demand a higher standard of evidence for rejecting the hypothesis than for accepting it.

The argument from inductive risk presumes some perspective on scientific reasoning according to which practical or ethical costs can influence what one should or should not accept. A commonly made objection to the argument, therefore, rejects premise 2 on the grounds that nonepistemic values should not influence acceptance construed in the cognitive, nonbehavioral sense relevant to science (Levi 1960, 1962, 1967; McMullin 1982;

1. See Steel (2010, 17) for a similar formulation of the argument.

https://doi.org/10.1086/673936 Published online by Cambridge University Press

Lacey 1999, 2004; Dorato 2004; Mitchell 2004).[2] According to this objection, premise 2 can be upheld only if acceptance is interpreted behaviorally, that is, if it assumed that decisions about acceptance can only be made in connection with some practical decision (see Levi 1962, 50). But, the objection continues, such an interpretation of acceptance is inappropriate for research science wherein the aim is to discover the truth rather than to achieve some practical objective. The most direct way to answer this objection is to describe a plausible theory of acceptance that is not behavioral (i.e., that does not say that a statement can be accepted only in relation to a practical objective) but that nevertheless supports the argument from inductive risk. That task is taken up in the next section.

3. Accepting as Premising. Cohen's theory of acceptance centers on a distinction between acceptance and belief. According to Cohen,

> Belief that *p* is a disposition, when one is attending to issues raised, or items referred to, by the proposition that *p*, normally to feel it true that *p* and false that *not-p*, whether or not one is willing to act, speak, or reason accordingly. But to accept the proposition or rule of inference that *p* is to treat it as a given that *p*. More precisely, to accept that *p* is to have or adopt a policy of deeming, positing, or postulating that *p*—i.e. of including that proposition or rule among one's premises for deciding what to do or think in a particular context, whether or not one feels that it is true that *p*. (1992, 4)

The contrast drawn here, then, is between belief as a *disposition* to feel that *p* is true and acceptance as a *decision* to treat *p* as a premise in a particular context. According to this perspective, acceptance and belief often coincide. For instance, I see snow falling, which causes me to believe that it is snowing, and this belief is in turn a good reason to accept that it is snowing (e.g., to use this statement as a premise in practical reasoning about how to dress).

But acceptance and belief can also come apart. Since belief on Cohen's account is an involuntary feeling, it is possible for a person to strongly feel that *p* is true in circumstances in which she recognizes that belief in *p* is irrational. For example, imagine a person who suffers from a dog phobia and upon seeing a friendly Boston terrier has a strong feeling that she will be viciously bitten. But she is aware that this belief is entirely unreasonable. In such cases, then, she makes a policy of treating the claim that the dog is harmless as a premise and thus makes every effort to suppress be-

2. A less popular objection rejects premise 1 on the grounds that scientists should not accept hypotheses but only assign probabilities to them (Jeffrey 1956; see Douglas 2009, 53–55, for a critique of this objection).

https://doi.org/10.1086/673936 Published online by Cambridge University Press

haviors that would express the fear and alarm she so acutely feels. In this situation, Cohen's theory would say that the person believes that the dog will bite but accepts that it is harmless. Conversely, it is possible to accept something that one does not believe. For example, one might choose to take a generalization, such as the ideal gas law, as a premise in reasoning in a particular context, despite believing that it is an approximation that is strictly false (Cohen 1992, 90).[3]

Consider how Cohen's theory relates to the argument from inductive risk. Cohen argues that acceptance rather than belief is the appropriate focus of scientific reasoning (Cohen 1992, 86–100), which amounts to a defense of premise 1 of the argument. More significantly, Cohen's theory supports premise 2, which states that decisions about whether to accept a hypothesis can depend on ethical or practical consequences of error (see 12). Although Cohen does not elaborate this point in much detail, the underlying reasoning is not difficult to work out. First, note that according to Cohen's theory, accepting a claim p is a decision (i.e., to adopt a policy of treating p as an available premise for reasoning in the context in question) that can have consequences for ethical and practical matters. The potential relevance of ethical and practical concerns to decisions about what to accept, therefore, ensues immediately from the following innocuous principle.

> *Consequence Principle*: One should take the consequences of one's decisions into account in the process of deliberation.

The Consequence Principle does not say that consequences are the only thing that matters, nor does it specify how those consequences should be assessed. It merely says that consequences of our acts are relevant considerations in deciding what to do. Premise 2 of the argument of inductive risk is concerned with a particular sort of consequence, namely, the consequences of mistakenly accepting a claim. So, to support premise 2 Cohen's theory also needs to allow that a person can accept a claim whose truth is less than certain. This is entirely possible given Cohen's theory—for instance, I may decide to accept a weather forecast as a premise in reasoning about whether to carry an umbrella, while being well aware of the fallibility of meteorology. Generalizing from the specifics of Cohen's theory, then, the Consequence Principle leads to premise 2 of the argument from inductive risk pro-

3. Horwich (1991, 4) argues that in such cases the person does in fact believe the generalization is true in the specified context. But this suggestion is not plausible for generalizations, such as the ideal gas law, that are approximations that are not literally true of any actual situation. In such cases, the rationale for accepting the generalization is not that it is true in the specified context; instead, the rationale is that the generalization is, in that context, a good enough approximation for the intended purposes.

https://doi.org/10.1086/673936 Published online by Cambridge University Press

vided that acceptance is a decision that can have consequences for practical or ethical matters and be made with regard to less than certain claims.[4]

This suggests a neat categorization of three possible ways to reject premise 2. One might (*a*) take issue with the Consequence Principle, for instance, by claiming that science should be given a special dispensation to pursue knowledge for its own sake and without regard for consequences. I will not examine this idea here, as it has already been extensively critiqued (see Kitcher 2001; Douglas 2009, chap. 4). My focus instead will be on two other potential escape routes: (*b*) one could claim that acceptance decisions cannot (or should not) have consequences for practical actions, or (*c*) one could insist that acceptance is reasonable only when one is sufficiently certain as to be unconcerned about error.

4. Cognitive Utilities and Certainty. In this section, I examine two groups of alternative theories of acceptance, one that explicates rational acceptance in terms of maximizing expected cognitive utility and the other that regards certainty in one form or another as a criterion of acceptance. These two approaches to acceptance did not arise independently of the argument from inductive risk. Significant representatives of both—particularly, Levi (1967) and Lacey (1999)—were developed expressly for the purpose of countering Rudner's (1953) argument and explaining how acceptance of scientific claims can be grounded exclusively in epistemic values. Both of these approaches deny premise 2 of the argument from inductive risk, attempting escape by routes *b* and *c*, respectively. I argue that both of these proposals suffer from serious difficulties not faced by Cohen's theory.

4.1. Maximizing Expected Cognitive Utility. In this section, I examine theories that account for acceptance by postulating something called cognitive utilities. Whereas practical utilities are positive or negative values associated with overt actions, cognitive utilities are associated with acceptance. Thus, a dislike for bicycling in the rain and a zest for bicycling in the sunshine are matters of practical utility, while finding intrinsic value in accepting informative truths is an expression of cognitive utility. According to cognitive utility approaches, moreover, maximizing expected utility is the standard for judging the rationality of decisions. Since expected utility is a probability weighted average of utilities, the bifurcation of cognitive and practical utilities results in two separate types of expected utility calculations, both sharing a common probability distribution, but one using a cognitive utility function and the other a practical utility function (Levi 1967; Harsanyi 1983; Maher 1993). This scenario is depicted in figure 1.

4. Cohen's is not the only theory of acceptance that possesses these characteristics (see Harsanyi 1985; Bratman 1999).

https://doi.org/10.1086/673936 Published online by Cambridge University Press

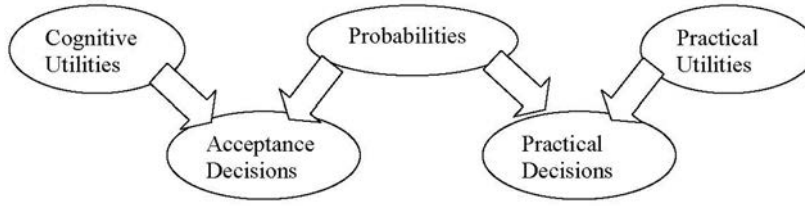


Figure 1. Cognitive utility theories of acceptance.

In figure 1 there is no arrow pointing from acceptance to practical decisions or indeed to anything else. This is a deliberate feature of cognitive utility theories: they claim that what we accept should make no difference to decisions about practical action (Maher 1993, 150). In Levi, this move explicitly arises in response to the argument from inductive risk as advocated by Rudner (Levi 1960, 1967).[5] Advocates of cognitive utility theories do not regard acceptance as totally inert, however, because they hold that acceptance should satisfy deductive cogency (Levi 1967, 26; Maher 1993, 139). Deductive cogency states that, if one accepts a set of sentences A and is aware that B is a logical consequence of A, then one ought to accept B as well or else give up accepting A. But given that one can accept normative as well as descriptive statements, deductive cogency conflicts with the claim that what one accepts should have no implications for practical reasoning. Consider the following example. Suppose I accept that the glass of water before me contains arsenic at a level greater than 50 parts per billion (W). Suppose that I also accept the normative claim that if the glass of water contains arsenic at a level greater than 50 parts per billion, then I should not drink it ($W \rightarrow \neg D$). Obviously, W and $W \rightarrow \neg D$ together logically entail $\neg D$, in other words, that I should not drink the water. Thus, deductive cogency requires that I accept $\neg D$, assuming that awareness of this logical consequence does not make me doubt the reasonableness of accepting W or $W \rightarrow \neg D$. This example illustrates that we can accept normative and descriptive claims that jointly entail further claims about what to do. Given deductive cogency, therefore, it is inescapable that descriptive claims we accept can function as premises in practical reasoning. But if this is so, how can it be that what we accept should have no impact on our actions, as cognitive utility theories assert? Perhaps advocates of cognitive utility theories would deny that we can accept normative statements. Yet this would be a highly implausible position since people frequently do accept norma-

5. Maher takes for granted that practical consequences of error should not influence judgments about what to accept and argues on this basis that what we accept should not influence practical decisions (1993, 150).

https://doi.org/10.1086/673936 Published online by Cambridge University Press

tive claims, for instance, that racism is bad (see Cohen 1992, 15). Moreover, Maher defines acceptance of a statement as the mental state expressed by the sincere and deliberate assertion of it (1993, 132). This apparently allows for the acceptance of normative statements, as people often do make sincere and deliberate assertions about what ought to be done. So, denying the possibility of accepting normative claims is not a promising option. What I think advocates of cognitive utility theories would say is that logical inferences from accepted descriptive and normative claims should not be a standard for rational action. Instead, they would claim, maximizing expected practical utility is the one and only criterion for what it is rational to do (149–50).

Yet the notion that logical consequences flowing from descriptive and normative claims one accepts should not influence one's actions is surely quite strange. Among other things, this would mean regarding practical and ethical reasoning framed in terms of logical inferences from accepted premises as irrelevant to human conduct. A further example may help to make vivid the deeply problematic nature of this position. Consider this famous argument proposed by Peter Singer (1972).

a) Dying from lack of food, shelter, or medical care is very bad.
b) If it is in our power to prevent something very bad without thereby sacrificing anything morally significant, we should do it.
c) It is in our power to prevent people in poorer nations, such as Bangladesh, from dying from lack of food, shelter, or medical care at a small financial cost to ourselves.
d) A small financial cost is not morally significant.
e) Therefore, we should prevent people in poorer nations from dying from lack of food, shelter, and medical care.

Discussions of this argument typically presume that accepting its conclusion should have some impact on our behavior, for example, causing us to support organizations such as Oxfam International or Médecins sans Frontières. If that is right, then whether one accepts the descriptive premise *c* could affect one's actions. But according to cognitive utility theories of acceptance this presumption is mistaken because, they insist, what we accept should make no difference to what we do. As Maher puts it, "*In any situation*, acceptance of *H* is consistent both with acting as if *H* were true and also with acting as if *H* were false" (1993, 149–50; italics added). Thus, according to cognitive utility theories, a rational person could accept premise *c* of Singer's argument, along with the other premises, but deny that this constituted any reason to provide aid on the grounds that accepting *c* is consistent with acting as if it were false. Yet blithely disregarding the moral implications of what one accepts strikes me, and most others I hope,

https://doi.org/10.1086/673936 Published online by Cambridge University Press

as very wrong. Nor can this be claimed as an unavoidable consequence of a Bayesian approach to decision, as some Bayesians have proposed theories of acceptance that do not have this disturbing feature (see Harsanyi 1985).

A possible reply to this objection is that cognitive utility theories are intended to explicate acceptance in purely cognitive contexts, wherein only descriptive claims can be accepted.[6] Limiting the scope of cognitive utility theories of acceptance in this way would block the objection above but at the cost of undermining their relevance to science. For example, accepted norms of methodology are inextricably involved in scientific reasoning. Such norms have an effect on what descriptive claims may be inferred, and new discoveries can have implications for which methods are appropriate. And of course, basic norms of research ethics should be accepted in all scientific contexts. So, it is highly doubtful whether scientific reasoning does, or should, proceed in a context that is purely cognitive in the sense just defined.

Attempts to use cognitive utility theories of acceptance as a basis for rejecting premise 2 of the argument from inductive risk, then, face the following dilemma. To say that accepted claims should not influence action is very problematic, given that it is possible to accept normative as well as descriptive claims and that the conjunction of the two can entail further claims about what ought to be done. However, if accepted claims may influence action, then there is a powerful argument for concluding—on the basis of the Consequence Principle as explained in section 3—that practical and ethical factors can matter to decisions about what to accept.[7] In either case, cognitive utility theories fail to provide a plausible account of acceptance that undermines the argument from inductive risk.

4.2. Acceptance and Certainty. The second premise of the argument of inductive risk presupposes that a decision to accept a scientific hypothesis generally involves a nonnegligible risk of error. Thus, some critics respond to the argument by claiming that acceptance, in a truly cognitive sense, requires a level of certainty that puts inductive risk out of the picture (Levi 1980; Lacey 1999, 71–74; Dorato 2004, 73–75). In this section, I argue that such proposals have the serious drawback of implausibly tying acceptance to dogmatism.

Levi (1980) proposes that accepting a claim as part of our body of knowledge requires judging it to be completely certain. A central function

6. I thank Kevin Elliott for suggesting this response.

7. Kaplan (1981) proposes a Bayesian theory of acceptance not founded on the concept of maximizing expected cognitive utility but that nevertheless requires that what we accept should have no impact on action (330). Thus, his proposal is also subject to this dilemma. See Eells (1983) and Maher (1993, 156–58) for additional critiques of Kaplan's proposal.

https://doi.org/10.1086/673936 Published online by Cambridge University Press

of knowledge, Levi suggests, is to distinguish what are serious possibilities at a given time from what are not (2–5). A claim is a serious possibility at a time just in case it is logically consistent with what we know then. Moreover, Levi says that only serious possibilities have a probability greater than zero (3). Hence, it follows from Levi's proposal that one must regard everything one accepts as knowledge as having probability 1, that is, as being absolutely certain.[8] Levi coins the term "epistemological infallibilism" to label this consequence of his proposal. Epistemological infallibilism would appear, on the face of it, to preclude the possibility that person could envision discovering that something she accepts is false. However, Levi insists that a rational person should, despite being absolutely certain of the complete truth of all she accepts as knowledge, nevertheless admit that some of what she accepts may be shown false by further inquiry (19). Thus, Levi claims that epistemological infallibilism is compatible with rejecting dogmatism and recognizing the corrigibility of what one accepts as knowledge.

Yet Levi's combination of infallibilism and corrigibility is paradoxical. If I am absolutely certain that a claim is true, then rationality requires that I am similarly certain that all new truths discovered in the future will be logically consistent with it. And if I admit that a claim may be disproven in the future, I am expressing uncertainty about its truth. One cannot, in short, have it both ways: either you grant that your knowledge is fallible and hence less than completely certain, or you insist that your knowledge is absolutely certain and thereby imply that it is infallible (see Harsanyi 1985, 22). Thus, Levi has not provided a good explanation of how his position avoids tying acceptance to dogmatism.

Lacey also links acceptance to certainty, although unlike Levi, does not regard complete certainty as a prerequisite for acceptance. Instead, he suggests that only "pragmatic" or "practical certitude" and not "epistemic certitude" is necessary (1999, 14, 71). By practical or pragmatic certitude, Lacey means that the theory "is sufficiently well supported that it need not be submitted to further investigation—where, for example, it is judged that further investigation can be expected only to replicate what has already been replicated many times over and to bring minor refinements of accuracy" (13).[9] But Lacey's retreat from epistemic to practical certainty undermines his objection to the argument from inductive risk. For decisions

8. Let K be our corpus of knowledge at t. Then $\neg K$ is not a serious possibility at t, and hence $p(\neg K) = 0$. Thus, $p(K) = 1$.

9. Lacey's concept of acceptance is similar to Levi's concept of acceptance as evidence (see Levi 1967, 149, 237–38), which also claims that accepting S involves judging further inquiry into S to be pointless.

https://doi.org/10.1086/673936 Published online by Cambridge University Press

about whether to accept a scientific claim in Lacey's sense involve a consideration of how probable or well supported it must be to justify the decision to no longer subject it to further investigation and to treat it as established knowledge. Decisions of this kind can have practical consequences in a world in which scientific knowledge affects decision making on matters of important public concern, such as climate change, regulation of toxic chemicals, or medical practices. The Consequence Principle can then be used to argue, as explained in section 3, that ethical or practical concerns can legitimately influence decisions about what to accept. Therefore, Lacey has provided no reason to undercut the argument that practical and ethical consequences of errors can be relevant to acceptance.

In addition, Lacey's proposal is problematic considered on its own terms because it links acceptance to a dogmatic dismissal of the usefulness of further inquiry. Yet accepting *S* does not require thinking that further research on *S* is pointless. A scientist might accept that anthropogenic greenhouse gas emissions are the primary factor driving climate change, while regarding further documentation of that claim as worthwhile. She could, for instance, judge this research to be useful because it can produce evidence that reduces uncertainty, which may be relevant to her own intellectual curiosity and which may be helpful for persuading others.

5. Conclusion. This essay has responded to a commonly made charge that ethical or practical considerations can appropriately influence decisions about what to accept only if acceptance is construed in a behavioral sense that it is ill suited for science. The response consists in, first, describing a theory of acceptance, developed by Cohen (1992), that supports the argument from inductive risk but that is not behavioral. The second part of the response is a critical examination of theories of acceptance that attempt to explain how acceptance can be guided solely by epistemic values. I argued that any such theory of acceptance confronts a choice among three unpalatable options: (*a*) claim that scientists should be relieved of the normal human responsibility to consider the consequences of one's actions, (*b*) deny that acceptance can (or should) have any implications for practical action, or (*c*) make certainty a prerequisite for acceptance. While option *a* has been critiqued by others (Kitcher 2001; Douglas 2009), this essay explored difficulties inherent in theories of acceptance that opt for *b* or *c*. Option *b* is problematic because it is possible to accept normative as well as descriptive claims, which jointly entail further claims about what ought to be done, while option *c* implausibly links acceptance to dogmatism. Thus, the original objection can be transformed into a challenge for the objectors: provide a plausible theory according to which nonepistemic values should not influence decisions about accepting or rejecting scientific hypotheses.

https://doi.org/10.1086/673936 Published online by Cambridge University Press

REFERENCES

Biddle, J., and E. Winsberg. 2010. "Value Judgments in the Estimation of Uncertainty in Climate Modeling." In *New Waves in Philosophy of Science*, ed. P. D. Magnus and J. Busch, 172–97. New York: Macmillan.
Bratman, M. 1999. *Faces of Intention*. Cambridge: Cambridge University Press.
Cohen, J. 1992. *An Essay on Belief and Acceptance*. Oxford: Clarendon.
Cranor, C. 1993. *Regulating Toxic Substances*. Oxford: Oxford University Press.
———. 2006. *Toxic Torts*. Cambridge: Cambridge University Press.
Dorato, M. 2004. "Epistemic and Nonepistemic Values in Science." In *Science, Values, and Objectivity*, ed. P. Machamer and G. Wolters, 52–77. Pittsburgh: University of Pittsburgh Press.
Douglas, H. 2000. "Inductive Risk and Values in Science." *Philosophy of Science* 67:559–79.
———. 2009. *Science, Policy and the Value-Free Ideal*. Pittsburgh: University of Pittsburgh Press.
Eells, E. 1983. "On a Recent Theory of Rational Acceptance." *Philosophical Studies* 44:331–43.
Elliott, K. 2011. *Is a Little Pollution Good for You? Incorporating Societal Values in Environmental Research*. Oxford: Oxford University Press.
Giere, Ronald. 2003. "A New Program for Philosophy of Science?" *Philosophy of Science* 70: 15–21.
Harsanyi, J. 1983. "Bayesian Decision Theory, Subjective and Objective Probabilities, and Acceptance of Empirical Hypotheses." *Synthese* 57:341–65.
———. 1985. "Acceptance of Empirical Statements: A Bayesian Theory without Cognitive Utilities." *Theory and Decision* 18:1–30.
Horwich, P. 1991. "On the Nature and Norms of Theoretical Commitment." *Philosophy of Science* 58:1–14.
Jeffrey, R. 1956. "Valuation and Acceptance of Scientific Hypotheses." *Philosophy of Science* 23:237–46.
Kaplan, M. 1981. "A Bayesian Theory of Rational Acceptance." *Journal of Philosophy* 78:305–30.
Kitcher, Philip. 2001. *Science, Truth, and Democracy.* Oxford: Oxford University Press.
Kourany, J. 2003. "A Philosophy of Science for the Twenty-First Century." *Philosophy of Science* 70:1–14.
Lacey, H. 1999. *Is Science Value Free?* London: Routledge.
———. 2004. "Is There a Significant Distinction between Cognitive and Social Values?" In *Science, Values, and Objectivity*, ed. P. Machamer and G. Wolters, 24–51. Pittsburgh: University of Pittsburgh Press.
Levi, I. 1960. "Must the Scientist Make Value Judgments?" *Journal of Philosophy* 57:345–57.
———. 1962. "On the Seriousness of Mistakes." *Philosophy of Science* 29:47–65.
———. 1967. *Gambling with Truth*. London: Routledge & Kegan Paul.
———. 1980. *The Enterprise of Knowledge*. Cambridge, MA: MIT Press.
Maher, P. 1993. *Betting on Theories*. Cambridge: Cambridge University Press.
McMullin, E. 1982. "Values in Science." In *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, vol. 1, ed. P. Asquith and D. Nickles, 3–28. East Lansing, MI: Philosophy of Science Association.
Mitchell, S. 2004. "The Prescribed and Proscribed Values in Science Policy." In *Science, Values, and Objectivity*, ed. P. Machamer and G. Wolters, 245–55. Pittsburgh: University of Pittsburgh Press.
Rudner, R. 1953. "The Scientist qua Scientist Makes Value Judgments." *Philosophy of Science* 20:1–6.
Shrader-Frechette, K. 1991. *Risk and Rationality.* Berkeley: University of California Press.
Singer, P. 1972. "Famine, Affluence, and Morality." *Philosophy and Public Affairs* 1:229–43.
Steel, S. 2010. "Epistemic Values and the Argument from Inductive Risk." *Philosophy of Science* 77:14–34.

https://doi.org/10.1086/673936 Published online by Cambridge University Press