Philosophy of Science (2024), **91**, 18–36
doi:10.1017/psa.2023.107

PSA PHILOSOPHY OF SCIENCE ASSOCIATION

PHILOSOPHY OF SCIENCE

ARTICLE

The Epistemic Projection Approach to Values in Science

Wendy S. Parker

Department of Philosophy, Virginia Tech, Blacksburg, VA, USA
Email: wendyparker@vt.edu

(Received 23 August 2021; revised 18 January 2023; accepted 18 July 2023; first published online 23 August 2023)

Abstract

I develop a novel account of how nonepistemic aims and values can appropriately influence scientific investigation. At its heart is a process of epistemic projection in which a nonepistemic aim or value is mapped to an epistemic research problem that aligns with that aim or value. Choices in research are then justified as a means of solving that research problem. This *epistemic projection approach* makes research responsive to nonepistemic aims and values yet remains consistent with the value-free ideal; it could be acceptable to parties on both sides of the values-in-science debate. It also promises to be useful in practice.

1. Introduction

Philosophical debate over nonepistemic value influence in science has often focused on what happens in the "internal" or "research" stage of investigation, where scientists perform experiments, build and use models, collect and analyze data, and reach conclusions.[1] According to the *value-free ideal* (VFI), such activities should, as much as possible, be kept free from the influence of nonepistemic values. The least controversial way of understanding the sort of influence that the VFI is meant to prohibit is related to justification: nonepistemic value commitments—such as the positive significance that one attaches to human health, or economic growth, or environmental sustainability—should not serve as *justifying reasons* for choosing some methodological options over others or for accepting hypotheses.[2] The rationale for such a prohibition seems straightforward: nonepistemic values, reflecting what someone wishes or desires to be the case, are irrelevant to methodological choices

[1] Readers who are skeptical of a viable distinction between epistemic and nonepistemic values should feel free to substitute "social and ethical values" or "contextual values" for "nonepistemic values" in what follows.

[2] Ward (2021) discusses different types of value influence. Reiss and Sprenger (2020) distinguish several versions of the VFI.

aimed at finding out what *is* the case, and their influence can bias and distort research in significant ways.

An emerging consensus among philosophers of science, however, is that the VFI should be rejected. According to the argument from inductive risk, nonepistemic value judgments sometimes appropriately figure in the justification of uncertain methodological choices, as scientists try to avoid erring in ways that would have particularly bad practical or social consequences (Douglas 2000, 2009). A second prominent line of argument starts from the observation that research often has joint epistemic and practical aims, where the latter reflect nonepistemic values. It can be entirely appropriate, on this "aims" view, to choose some methodological options over others *because* they promote the joint epistemic and practical (i.e., nonepistemic) aims of research; the point of the research, after all, is to promote or achieve those aims (Elliott and McKaughan 2014; Intemann 2015; see also Brown 2017). According to both lines of argument, nonepistemic value commitments sometimes appropriately constitute pro tanto justifying reasons for choices "internal" to research, contra the VFI.

In a recent paper, Lusk and Elliott (2022) argue that such challenges to the VFI presuppose that scientific assessment can be concerned with more than truth, contrary to the view of VFI proponents. For the challenges to be successful, they contend, an account of how such assessment works is needed. They propose an adequacy-for-purpose account, adapted from scientific model evaluation (Parker 2020). On their account, assessment seeks to determine whether some element of scientific practice (i.e., hypothesis, model, method, etc.) can be used successfully for a purpose of interest. Depending on the purpose, adequacy might require meeting both epistemic criteria (e.g., a specific level of accuracy) and nonepistemic criteria (e.g., giving results quickly) (see also Harvard and Winsberg 2022). According to Lusk and Elliott, even "plain" scientific hypotheses—such as "Earth's climate is warming"—can be appropriately assessed in terms of their adequacy for various purposes, including nonepistemic purposes.

Lusk and Elliott (2022) go on to envision a partial rapprochement among parties to the values-in-science debate, where the parties come to agree that scientific assessment is concerned only with truth, but where the targets of assessment (whose truth is in question) are no longer plain scientific hypotheses—they are hypotheses about the *adequacy* of plain hypotheses and other elements of scientific practice for purposes of interest. As Lusk and Elliott recognize, however, it is not clear that the relevant parties would agree to this. Opponents of the VFI are likely to argue that inductive risk (and thus value) considerations remain relevant when deciding whether to accept adequacy-for-purpose hypotheses, and proponents of the VFI may insist that, really, it is the truth status of plain scientific hypotheses that matters in science.

The present article envisions a less controversial pathway to rapprochement in the values-in-science debate, in the form of a novel account of how nonepistemic values can appropriately influence scientific investigation. At the heart of the account is a process of *epistemic projection*, which involves specifying, in the problem selection stage of investigation, a set of epistemic goals, preferences, and constraints—an *epistemic research problem*—that aligns with a targeted nonepistemic aim or value. Choices made in the research stage of investigation are then justified as a means of solving that epistemic research problem. This *epistemic projection approach* makes research responsive to nonepistemic aims and values, without violating the VFI; it is

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

an approach that could be acceptable to parties on both sides of the values-in-science debate. It also promises to be useful in scientific practice.

The discussion proceeds as follows. Section 2 introduces the process of epistemic projection as a means of formulating a research problem. In section 3, the process of epistemic projection is embedded in an account of how nonepistemic values can appropriately influence scientific investigation: the epistemic projection approach (EPA). Section 4 anticipates and responds to several objections to EPA. Section 5 considers EPA's potential for achieving rapprochement among parties to the values-in-science debate. Finally, section 6 offers concluding remarks.

2. Problem selection and epistemic projection

Suppose a researcher, Alex, adjusts the numerical value assigned to a parameter in a computer model so that the model gives more accurate outputs for physical variable X. She makes this adjustment even though she expects that, as a consequence, the model's performance on physical variable Y will degrade; she accepts this trade-off because her modeling group considers accuracy on X to be a higher priority than accuracy on Y. This prioritization of X, in turn, ultimately stems from the fact that X is closely related to outcomes that are disvalued by society (e.g., flooding, or forest fires, or spread of disease), whereas Y is not.

Examples like this are taken to illustrate that it is sometimes unproblematic for choices "internal" to research to be justified by appeal to nonepistemic values; in this case, the choices concern the numbers assigned to parameters in a computer model. Yet from the brief description of what happened, it is not clear that the purported value influence occurred. Perhaps what occurred instead was that, given X's practical importance, its accurate simulation was recognized as a top priority for the investigation right from the start, in the problem selection stage, and the parameter choices subsequently made in the internal/research stage were justified not by appeal to nonepistemic values but in light of that epistemic priority (see also Elliott and McKaughan 2014, 9–11; Parker and Winsberg 2018, sec. 2). The latter sequence of events is not implausible and is suggestive of a general strategy for taking account of nonepistemic aims/values in scientific investigation—without violating the VFI. To articulate the strategy, however, we first need to take a closer look at the problem selection stage of investigation.

2.1 Problem selection

Oftentimes, problem selection is assumed to involve little more than the specification of one or more research questions or goals. In the second possible sequence of events just described, however, problem selection also involved identifying some research goals as higher priority than others. And the very formulation of the goals—in terms of accurate simulation—reflected an assumption about the type of study that would be undertaken, namely, that it would be a computer simulation study (rather than, e.g., a randomized controlled trial or a qualitative survey). In practice, even in the earliest stages of scientific investigation, there is often a relatively rich, even if implicit, conceptualization of the basic outline of the research to be undertaken.

In the present discussion, problem selection will be understood in a way that leaves room for this richness and allows for making it explicit. In particular, problem

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

selection will be understood as an activity that involves the specification of epistemic goals for research and, optionally, desiderata regarding how those goals are pursued. *Epistemic goals* identify specific epistemic goods (pieces of knowledge, information, understanding) that are sought. Desiderata regarding the pursuit of such goals can take the form *preferences*, which specify features that it is desirable for the research process to have, or *constraints*, which specify features that the research process is required to have. In the preceding example, for instance, there was a preference for prioritizing accurate simulation of X over accurate simulation of various other physical quantities. Together, epistemic goals and desiderata (if any) constitute a *research problem*.[3] A research problem provides a kind of brief for the research to be undertaken: researchers should try to achieve the epistemic goals in a way that satisfies the constraints and preferences.

Constraints and preferences can be nonepistemic, epistemic, or mixed in character. *Nonepistemic* constraints and preferences pertain to nonepistemic aspects of research, such as its side effects, costs, or consequences. For example, a nonepistemic constraint for some archaeological research might be that the research process *not damage or disrespect* the religious artifacts that will be examined. *Mixed* constraints and preferences pertain to epistemic aspects of the research process but also explicitly reference a nonepistemic aim/value. An example would be a preference for prioritizing *accurate simulation* of whichever model variables are most relevant to *policy decision P*. *Epistemic* constraints and preferences pertain to epistemic aspects of the research process and are expressed in a way that makes no reference to any nonepistemic aims/values (though such aims/values can motivate their specification). An example would be a preference for *risking a Type I error* rather than a Type II error when increasing the risk of one or the other at a methodological decision point is unavoidable. If a research problem includes only epistemic constraints and preferences (along with epistemic goals), we will call it an *epistemic research problem*.

2.2 Epistemic projection

Suppose that Alex's modeling group was performing their study in part because they wanted to provide information that would be useful for certain practical decisions. At the start, they might have reasoned as follows: because we want the study findings to be useful for these practical decisions, we should prioritize accurate simulation of the variables that are most relevant to the decisions; the variables that are most relevant are {X, . . .}; so we should prioritize accurate simulation of {X, . . .} if accuracy trade-offs must be made when constructing and tuning our model. This process of reasoning from a broader aim to an epistemic characteristic of research that *aligns with* that broader aim—that is, an epistemic characteristic that helps make the research *relevant to*, *useful for*, or in some other way *well-suited to* the pursuit of that broader aim—will be the essence of the strategy developed here for taking account of nonepistemic aims and values without violating the VFI.

Only somewhat metaphorically, we can characterize the strategy as one of *epistemic projection*. In mathematics, a projection is a mapping of a structure onto a

[3] Steve Elliott (2022) provides a deeper and more detailed analysis of the structure of research problems. The simple analysis herein suffices for present purposes.

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

substructure, for example, a three-dimensional object mapped onto a two-dimensional plane. This geometrical process of projection shows us what the object "looks like" when collapsed to only two dimensions. Similarly, the process of epistemic projection—*the process of reasoning from a broader aim or background value to an epistemic research problem that aligns with that aim/value*—reveals what pursuit of that aim or responsiveness to that value can "look like" when collapsed to a set of choices concerning *epistemic* characteristics of the study to be undertaken, such as its epistemic goals or epistemic priorities. A study characteristic for which there are multiple options in effect defines a "dimension" of the "space" onto which the broader aim is projected; the process of epistemic projection picks out a point (or set of points) in each of these dimensions, corresponding to option(s) that align with (i.e., would be helpful in the pursuit of) the broader aim.[4]

Let's consider the process of epistemic projection in more detail. To identify *epistemic goals* that align with a broader nonepistemic aim, we can consider which epistemic goods (i.e., pieces of knowledge, information, understanding) are *relevant* to the pursuit of the aim. For example, if the nonepistemic aim is to provide treatment to all (and only) people in a population who have a certain serious medical condition that often goes unnoticed, then the epistemic goal of *learning which people in the population have the condition* would be relevant. When multiple epistemic goods are relevant to the pursuit of a nonepistemic aim, any one (or several) of them may be selected when formulating epistemic goals in the process of epistemic projection; the choice ordinarily will be guided by considerations of feasibility and degree of relevance.

When it comes to identifying epistemic desiderata, that is, *epistemic preferences* and *epistemic constraints*, four types in particular are worth considering; in many cases, these will encompass the most significant opportunities for aligning research with a nonepistemic aim:[5]

A. **Inductive risk desiderata** concern evidential standards for a study, such as the statistical significance level at which a null hypothesis will be rejected and/or the type of error in conclusions that should be risked, when some risk must be borne (Douglas 2000, 2009). Example: a preference for risking a Type I error rather than a Type II error when increasing the risk of one or the other is unavoidable. *Inductive risk desiderata can align with a nonepistemic aim when achieving the aim will be impeded more by some types of error in research results than others.*

[4] Clearly the metaphor is imperfect. For instance, unlike mathematical projection, epistemic projection is typically one to many, that is, one nonepistemic aim can be "mapped" to various epistemic research problems. But the metaphor aptly conveys the key idea of *collapsing* or *limiting* the actions that might be taken in the service of a nonepistemic aim to a particular subset, namely, those involved in carrying out an epistemic investigation.

[5] There is substantial overlap between the aspects of research targeted by the types of desiderata listed in what follows and aspects that opponents of the VFI have identified as legitimate targets of value influence (see esp. Elliott 2017). Because the categories were derived semi-independently, the overlap provides some reassurance that the main opportunities for aligning research with nonepistemic aims/values are in fact being captured. But the claim is not that these four types exhaust the possibilities; perhaps other important types can be added to the list.

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

B. **Prioritization desiderata** help to resolve other epistemic trade-offs. When a study has multiple epistemic goals, prioritization desiderata specify which one(s) take priority when trade-offs must be made. Example: a preference for prioritizing accurate simulation of physical variable X over accurate simulation of variable Y, if trade-offs must be made.[6] *Prioritization desiderata can align with a nonepistemic aim when some of the epistemic goals of research are more relevant than others to achieving the broader aim.*

C. **Method desiderata** concern the research strategies, data sets, models, concepts, methods, analysis tools, and other epistemic resources that will be used in a study. Example: a preference for research on the prevalence of long-COVID-19 to employ a particular operationalization of "long-COVID-19." *Method desiderata can align with a nonepistemic aim in various circumstances*, for instance, when the use of some epistemic resources rather than others will enhance the credibility of research among stakeholders whose cooperation is required to achieve the nonepistemic aim.

D. **Form-of-conclusion desiderata** concern the form that research conclusions take, such as their precision, resolution, terminology, or format. Example: a preference for findings about household water quality in a multi-state region to be reported at the county level and accompanied by a color-coded map. *Form-of-conclusion desiderata can align with a nonepistemic aim when formulating conclusions in some ways rather than others will enhance the salience, credibility, or usability of findings for users* (see also Elliott 2017).

As illustrated in these examples, desiderata articulated via epistemic projection pertain to *epistemic characteristics* of research, that is, features of research qua knowledge-seeking activity. They are selected to *align with* a nonepistemic aim but are expressed in a way that *makes no reference to* that (or any other) aim.[7]

Epistemic projection works in much the same way when research does not have nonepistemic aims per se but *background values* are at stake. In Heather Douglas's (2000) well-known dioxin toxicology example, where human health is a background value at stake, epistemic projection could specify an *inductive risk preference* for risking a Type 1 error (say, erroneously concluding that dioxin is carcinogenic) rather than a Type II error when increasing the risk of one or the other is unavoidable, because the consequences of a Type 1 error would be less threatening to human health. For another example, consider an ecologist who wants his research team to investigate the dynamics of a particular ecosystem; he has no further nonepistemic aims in mind—he just wants to know more about the workings of the ecosystem—but he is concerned that the investigation itself will damage the ecosystem's integrity or stability. When performing epistemic projection, he might specify a *method constraint:* the research should employ a method from a particular family of methods.

[6] Typically, such desiderata should indicate *the extent to which* one goal takes priority over another. This complexity is omitted in the present discussion.

[7] Epistemic projection presents no challenge to Heather Douglas's (2017) position that no purely internal (to science) standards can nonarbitrarily resolve how much evidence is sufficient for accepting a hypothesis or whether a Type 1 or Type 2 error should be preferentially risked. In epistemic projection, decisions about such matters *are* informed by nonepistemic aims/values—but the decisions take place in the problem selection stage of investigation, in the process of specifying an epistemic research problem.

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

He specifies this constraint because he believes that methods of this type are least likely to significantly disrupt the functioning of the ecosystem. Here values function as what Matthew Brown (2017) calls *side constraints*: they limit the kind of inquiry one is willing to pursue.[8]

In sum, epistemic projection is an approach to research problem selection; it involves reasoning from a nonepistemic aim or background value to an epistemic research problem that aligns with that aim/value but that is expressed in a way that makes no reference to it.[9] The epistemic research problem then constitutes a brief for the research to be undertaken: try to achieve the epistemic goal(s) while satisfying the epistemic preferences and constraints.[10] Such research could be performed successfully by someone completely unaware of the nonepistemic aims/values operative in the problem selection stage. Epistemic projection *screens off* nonepistemic aims/values from the research stage of investigation.

2.3 An example of epistemic projection

Suppose a group of regional planners is considering which housing development pathways to recommend for their region. Some of the pathways under consideration are attractive from the perspective of economic growth, but the planners are concerned that, if summers get much hotter in the region, these pathways could carry a substantial risk of unmeetable summer energy demand (for air conditioning) and, for people in some living situations, heat-related deaths; the planners want to avoid pathways that carry such risks. To aid their analysis, they ask a group of climate scientists to investigate what future summers in the region will be like; the scientists are offered up to $150,000 to carry out the research and are asked to provide their findings within nine months.

The climate scientists, in consultation with the regional planners, engage in a process of epistemic projection. They structure the process with the help of a series of questions, which are listed here along with a sketch of the reasoning and choices made in response:

1. **Epistemic goals.** *What information about future climate is particularly relevant to the planners' nonepistemic aims (of avoiding development pathways that will lead to unmeetable energy demand and increases in heat-related deaths)?* Information about specific kinds of summer heat events is identified as especially relevant. Three *epistemic goals* are specified: to estimate increases in the frequency of (i) very hot

[8] Basic ethical principles for human subject research are another example of side constraints. The VFI is generally understood to make an exception for such ethical side constraints, permitting them to figure in the justification of choices internal to research. Thus, to meet the standards of the VFI, epistemic projection needn't attempt to specify epistemic constraints that ensure that ethical principles are not violated. But method constraints are one means of doing so.

[9] Although the discussion here focuses on nonepistemic aims/values, epistemic projection can be carried out just the same when research has broader aims/values that are epistemic or cognitive in character, for example, to accelerate discoveries in a novel area of investigation or to resolve an apparent conflict among theories.

[10] Martin Carrier (2022) proposes that, in the context of science-based policy advice, nonepistemic aims/values can function as "separate premises" or "commissions" for the research process. We might think of epistemic projection as a process that transforms such premises/commissions into preferences and constraints for the research process.

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

days, (ii) heatwaves, and (iii) strings of hot nights that will occur in the region by the mid-twenty-first century under a "moderate" greenhouse gas emission scenario.

2. **Inductive risk desiderata.** *Would some errors in findings, more than others, impede pursuit of the nonepistemic aims?* Findings could err by underestimating or overestimating changes in the frequency of these heat events. But underestimating is more likely to lead to the selection of development pathways that lead to unmeetable energy demand and increases in heat-related deaths. So an *inductive risk preference* is specified for risking overestimating (rather than underestimating) increases in heat event frequencies when risk of one type or the other must be taken on.
3. **Prioritization desiderata.** *Are some of the specified epistemic goals more relevant or important than others to the pursuit of the nonepistemic aims?* It turns out that the planners are most concerned about increases in heat-related deaths, and most heat-related deaths occur when people cannot cool enough at night, so a *prioritization preference* is specified for accurate information about (iii) strings of hot nights.
4. **Method desiderata.** *Would the use of some epistemic resources (methods, data sets, etc.) rather than others be advantageous with respect to pursuit of the nonepistemic aim?* The regional planners request that baseline statistics on past heat events be calculated using the national government's official meteorological data set; this is for the sake of consistency with related governmental decisions and will contribute to the political legitimacy of the planning recommendations. Because the scientists consider the government's data set to be of acceptable quality, they adopt its use as a *method constraint*.
5. **Form-of-conclusion desiderata.** *Would it facilitate pursuit of the nonepistemic aims to express findings in a particular form?* Because changes in the frequency of these heat events cannot be predicted with high precision, a set of coarse-grained decision-relevant thresholds are identified, which become the basis for a *form-of-conclusion preference*: it is preferable that scientists classify the expected changes in frequency as: decreasing, increasing less than 50 percent, increasing 50 to 150 percent, or increasing more than 150 percent.

At this point, the epistemic research problem has been specified (see table 1). If we look ahead to the research stage of investigation, we can imagine the scientists making a number of choices to satisfy the specified epistemic preferences and constraints. Given the *inductive risk preference*, the scientists choose to employ a state-of-the-art climate model that they judge less likely than most available models to underestimate these future heat events. Because information about nighttime temperatures is the top *priority*, the scientists invest time adjusting the model's representation of a particular physical process that they believe will increase the accuracy of its simulations of nighttime temperatures. When calculating changes in heat event frequencies, past frequencies are derived from the official government data set per the *method constraint*. Finally, given the study's *form-of-conclusion preference*, the scientists formulate their conclusions in terms of the four categories specified, for example, they report that the frequency of heatwaves is expected to increase by 50 to 150 percent. Each of these choices in the research stage of

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

Table 1. An epistemic research problem identified via epistemic projection

Epistemic research problem components	Specification
Epistemic goal(s)	Estimate changes in frequency of (i) hot days, (ii) heatwaves, (iii) strings of hot nights
Inductive risk preference	Risk *over*-estimating increases in heat events, where inductive risk is unavoidable
Prioritization preference	Highest priority is accurate information about (iii) strings of hot nights
Method constraint	Use official government data set as a baseline
Form-of-conclusion preference	Report frequency changes in terms of the 0/50/150% categories

investigation is justified as a means of satisfying the epistemic desiderata specified for the research, *without* further appeal to nonepistemic aims/values. Indeed, a different team of climate scientists, provided with only the epistemic research problem specified in table 1 and completely unaware of its connection with the nonepistemic aims of the regional planners, could easily make the same (or very similar) choices in research.

A study of future summer temperatures undertaken without the intention of facilitating any particular practical decision, and thus without epistemic projection, might look quite different. It might provide information about changes in average summer temperatures rather than in the frequency of heat events. The computer model and data set used would likely be chosen (from among those considered of decent quality) on grounds of convenience or accessibility and would likely be used without making any changes or improvements. If the regional planners used the results of such a study to aid their decision-making, they might be less likely to succeed in recommending development pathways that avoid unmeetable energy demand and increases in heat-related deaths (see also Parker and Lusk 2019).

3. Values in science: The epistemic projection approach

In the example just presented, the epistemic preferences and constraints of the research problem could be satisfied while still conducting research that meets usual standards of good scientific practice. But imagine a different situation. Suppose that regional planners are strongly antidevelopment and that they want research to show such dramatic future increases in heat events that they can recommend against further development. They specify a method preference—say, a preference for the use of a climate model from a small group of models known to be unreasonably sensitive to changes in greenhouse gas concentrations—to try to ensure that the study that is conducted will predict massive increases in heat events, regardless of whether this is accurate. This preference aligns with their antidevelopment aims, but it is precisely the sort of value influence that even those who reject the VFI would agree is inappropriate. If we are interested in how nonepistemic aims and values can *appropriately* shape scientific investigation, we cannot simply point to the process of epistemic projection.

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

In this section, then, the process of epistemic projection is embedded in an account of how to make research appropriately responsive to broader aims and values.[11] This *epistemic projection approach* involves performing epistemic projection, but with the additional requirement that the epistemic research problem that is specified leave room for research that respects basic constraints on epistemically adequate science. Scientists' job in the research stage of investigation will be to design and carry out a high-quality study that is fit for (i.e., constitutes a solution to) that epistemic research problem, reporting with their findings any choices made in response to epistemic desiderata. These core elements of EPA, along with several other important aspects of its implementation in practice, are discussed in more detail in the following pages.[12]

3.1 EPA: The problem selection stage

According to EPA, when scientific investigation is intended to be responsive to some nonepistemic aim/value, *epistemic projection should be performed at the outset*. Crucially, however, EPA restricts what will count as an acceptable research problem. The restriction stems from the view of science that serves as a foundation for EPA, namely, that science is fundamentally an epistemic endeavor; if science is to contribute to achieving broader practical aims, it should do so via the production of knowledge or other epistemic goods. EPA thus requires that a research problem specified via epistemic projection leave room for a study that respects basic constraints on epistemically adequate science; in that case, the problem will be called *scientifically acceptable* (see figure 1, upper half). EPA is a variety of what Dan Steel (2017) calls *epistemic constraint* approaches to values in science.

One basic constraint on epistemically adequate science is that research must be genuinely probative; it should not be set up such that a particular conclusion will almost certainly be reached, regardless of the truth (Mayo 1996; Anderson 2004; Douglas 2009; Steel and Whyte 2012; Steel 2017). Other basic constraints include that data should not be fabricated, that data should not be discarded simply because they fail to support a preferred conclusion, that conclusions should be informed by the evidence collected, and so on. The research problem specified by the antidevelopment regional planners imagined at the start of this section would *not* be scientifically acceptable: their goal was to obtain a predetermined result, and their method preference left no room for a genuinely probative study but rather was chosen to obtain their desired result.

When it comes to performing epistemic projection, EPA does not require that anyone in particular take part. In practice, who should do so will generally depend on to whose aims/values the research is intended to be responsive. Often a collaboration

[11] EPA is an account of how research can be appropriately responsive to nonepistemic aims and values, once those aims and values are selected. EPA does not address the bigger question of to *which* nonepistemic aims and values research should be responsive. Some leading proposals are that science should be responsive to aims and values that are democratically endorsed (Intemann 2015; Schroeder 2021) or that reflect social and ethical priorities (Elliott 2017); pluralist proposals are also advocated, according to which it is important to have multiple investigations responsive to different values (Lacey 2013; Carrier 2022).

[12] EPA is presented here as an account of how nonepistemic aims/values can appropriately influence scientific investigation, but it seems just as applicable when broader aims/values are epistemic or cognitive in character. Discussion of this must be left for another time.

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

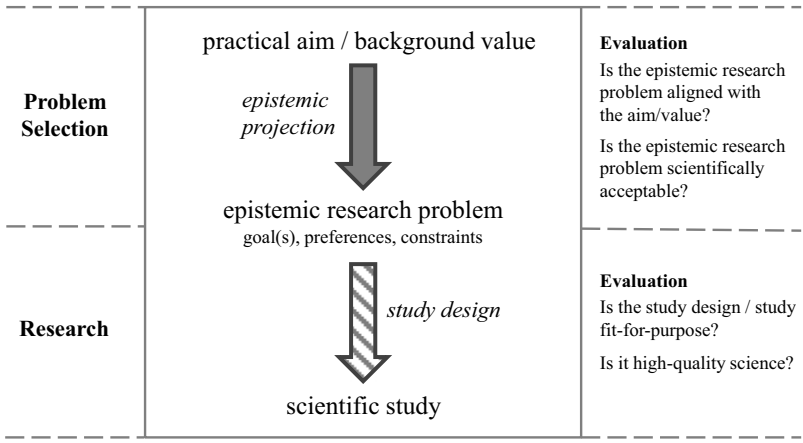

Figure 1. Elements of EPA.

between scientists and relevant decision makers or stakeholders (as in the example in section 2.3) will be the best approach, because scientists will have a sense of which epistemic goals are plausibly achievable and of the type of study that is likely to be undertaken (about which epistemic desiderata might be articulated) but may have limited understanding of the aims and values of decision makers or stakeholders to which the research is meant to be responsive.

Finally, it is important to note that an epistemic research problem specified at the outset of investigation need not be set in stone. EPA allows that a research problem be revised if it is realized, as research unfolds, that the process of epistemic projection could have been done better—for example, if it is realized that there are additional epistemic preferences that would align with the targeted nonepistemic aim/value. An epistemic research problem might also be revised if the targeted aim/value itself changes, for example, if the research is meant to inform a decision process but decision makers' priorities or practical constraints undergo change.

3.2 EPA: The internal/research stage

The epistemic research problem specified in the problem selection stage of investigation constitutes a brief for the research to be undertaken: scientists' job is to design and carry out a study that achieves the specified epistemic goal(s) in a way that satisfies the specified epistemic desiderata. In doing so, scientists, qua scientists, will need to *respect basic constraints on epistemically adequate science*. Indeed, scientists following EPA will be expected not only to respect those basic constraints but to strive to conduct research that is *of excellent epistemic quality*—that is, that uses reliable methods, that controls for errors, that involves sound reasoning, and so on—because doing so can be expected to help them *succeed* in achieving the epistemic goals of the study.

The adequacy-for-purpose framework that Lusk and Elliott (2022) employ, briefly introduced in section 1, is apropos here. But whereas they envision assessing the adequacy or fitness of various elements of scientific practice (hypotheses, theories, methods, etc.), in the context of EPA, the target of assessment will be a scientific study,

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

or a study design, and what is at issue is its fitness for solving an epistemic research problem (see figure 1, lower half). Parker (2020) articulates two varieties of fitness-for-purpose, which can be adapted to the present context as follows. A study is *fit*$_I$*-for-purpose* iff it achieves the research problem's epistemic goals while satisfying all of its epistemic constraints; a study's *fitness*$_I$*-for-purpose* is greater to the extent that it also satisfies the research problem's epistemic preferences. Similarly, a study design is *fit*$_C$*-for-purpose* iff C-type instances of execution of the design are likely to be *fit*$_I$*-for-purpose*; here C refers to features of the context in which the study design will be executed, including the skills, resources, and background knowledge of the scientists involved. A study design's *fitness*$_C$*-for-purpose* increases both as the likelihood of achieving the research problem's epistemic goals increases and as more of the epistemic preferences of the research problem are satisfied. Formulating a study design that has high *fitness*$_C$*-for-purpose* is a good strategy for producing a study whose *fitness*$_I$*-for-purpose* is high.

Note that if epistemic projection has been done skillfully, and the design for the study that is conducted has high fitness-for-purpose, then the methodological options selected in the course of research are likely to include some (and perhaps many) that are the same as those that would be selected by researchers who accept the inductive risk and aims arguments. This is because skillful epistemic projection will have identified inductive risk, prioritization, and/or other epistemic desiderata that align with the targeted aim/value, and the study will have been conducted in a way that satisfies (most or all of) those desiderata. Following EPA, however, researchers will not justify their methodological choices by appeal to nonepistemic aims or values; they will justify them as a means to solving the study's epistemic research problem, maintaining consistency with the VFI.

In some cases, it will not be possible to perform a study that is fit-for-purpose. Epistemic research problems might include goals that are out of reach for today's science, for example, the goal of producing very accurate and precise predictions of stock market behavior. If the insolubility of a research problem is recognized, a good response is to identify a less demanding research problem that aligns with the relevant nonepistemic aim/value. A similar revisionary strategy is recommended if scientists find that an epistemic research problem formulated by others requires methodological choices that, while perhaps not violating basic constraints on epistemically adequate science, nevertheless do not meet the scientists' own (higher) minimum standards for quality; scientists in this situation should propose revisions to the research problem and, if other parties are unwilling to make changes, should feel free to decline to perform the research.

3.3 EPA: Transparency

EPA also includes a transparency requirement: the epistemic research problem, as well as which choices in research were made in response to which epistemic desiderata, should be reported along with study findings. This is the analogue of a transparency requirement often advocated by opponents of the VFI, but they call for researchers to report where methodological choices in research were justified by appeal to nonepistemic values. Under EPA, by contrast, one would report that a

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

particular methodological choice was made in light of an epistemic constraint or preference specified as part of the epistemic research problem being addressed. For example, if a particular analysis technique was chosen to satisfy a method constraint, this would be reported.

There are at least two reasons for requiring such transparency. First, transparency can help consumers of research studies to interpret and weight study findings. When encountering findings from multiple studies, consumers might give greater weight to those in which inductive risk and other preferences were similar to what they would have specified in light of their own aims/values. Second, transparency facilitates scrutiny of choices made. EPA requires that scientists strive to conduct research that is epistemically excellent, but it is of course possible that they will fall short in this regard. For example, a scientist might select a particular method in response to a method preference, when there are alternative methods available that would also satisfy that preference and that would be more likely to achieve the epistemic goals of the study. Of course, relatively poor choices can occur at any point in the research process, and for various reasons, but choices made to satisfy epistemic desiderata may be prime candidates, especially if it is the scientist herself who specified those desiderata in light of her own strongly held values.

4. Objections and replies

This section considers several potential objections to EPA and attempts to forestall them. In brief, these objections charge that EPA is nothing new, that it simply relabels value judgments, that it will be less effective than alternative approaches, that it unhelpfully hides value influences, and that it is vague and overly permissive.

A first objection contends that EPA is just an amalgam of existing views on values in science, rather than a substantively novel proposal. The fact that EPA can be characterized as a variety of epistemic constraint approach and is expected to lead scientists to choose methodological options that likewise would be selected by those who accept the inductive risk and aims arguments seems to bolster this objection.

Response. EPA is distinctive in several important respects. First, EPA proposes that values justify choices only in the *problem selection stage* of investigation, and in doing so, it relies on a richer understanding of research problems than is commonly assumed. Second, EPA provides *concrete guidance* for making research responsive to nonepistemic aims/values: it calls for performing epistemic projection, which includes considering whether particular types of epistemic desiderata (i.e., inductive risk, prioritization, method, form-of-conclusion) would be advantageous for research to have, given a nonepistemic aim/value of interest. Third, because these epistemic desiderata are expressed in ways that make no reference to nonepistemic aims/values, EPA *screens off* nonepistemic aims/values from the research stage of investigation; this is perhaps its most significant innovation. Finally, and because of this screening off, EPA maintains *consistency with the VFI*.

A second, closely related objection charges that EPA achieves consistency with VFI only through sleight of hand, simply by relabeling nonepistemic value judgments as "epistemic" desiderata. According to this objection, though a preference for risking a Type 1 error or for prioritizing accurate simulation of one quantity over another is

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

called an "epistemic" preference in the context of EPA, it is little more than a value judgment in disguise.[13]

Response. This objection seems to rest on a misunderstanding of the nature of epistemic projection. What happens during epistemic projection is not a relabeling of value judgments but a process of *reasoning* from nonepistemic aims/values to epistemic characteristics that it would be advantageous for research to have, given those aims/values. Such characteristics (and preference and constraints regarding them) merit the label "epistemic" because they concern features of research qua knowledge-seeking activity—its methods, the errors to which it is more or less prone, and so on.[14] EPA achieves consistency with the VFI not by relabeling value judgments but by *relocating* value considerations to the problem selection stage of investigation, and in such a way that they are *screened off* from the research stage.

A third objection concerns the effectiveness of EPA. It maintains that the process of epistemic projection will rarely succeed in identifying—at the outset of investigation—all of the epistemic characteristics of a study that would align with some nonepistemic aim/value. It would be more effective to require scientists to attend to nonepistemic aims/values throughout their investigation. That way, they will be less likely to miss opportunities for making research responsive to such aims/values.

Response. First, it is an empirical question how adept anyone is at the process of epistemic projection. Recall also that EPA allows epistemic research problems to be revised if, in the course of conducting research, scientists notice additional significant opportunities to align research with a targeted nonepistemic aim/value. That said, it is plausible that fewer opportunities for such alignment would be missed by an approach that required scientists to select, at each methodological decision point, whichever of the epistemically acceptable options best aligns with the targeted nonepistemic aim/value. Such an approach, however, would violate the VFI, because it would allow for justifying choices *in research* by appeal to nonepistemic aims/values; it would not have EPA's potential for rapprochement (see section 5). It would also be more burdensome for practitioners than EPA.

A fourth concern is that EPA will unhelpfully hide value influences. EPA does not require that scientists report how nonepistemic aims/values shaped the specification of epistemic research problems; it requires only that they report which choices in research were made to satisfy a research problem's epistemic preferences/constraints. Consequently, consumers of scientific studies will have to work harder to determine whether studies are responsive to aims/values that they share.

Response. This is a reasonable worry. One option would be to strengthen the transparency requirement of EPA so that, in addition, scientists are required to report any nonepistemic aims/values that influenced the specification of the study's epistemic research problem. This option merits consideration. On the other hand, there is some evidence that reporting background value motivations for research can (unjustifiably) reduce the credibility of the research in the public eye (see Elliott et al. 2017). It is also

[13] Thanks to an anonymous reviewer for raising this objection.

[14] In the same way, a preference for eating vegan food can be characterized as a "dietary" preference, because it pertains to features of diet. It is a "dietary" preference regardless of whether it stems from health concerns, a desire to fit in with vegan friends, ethical considerations, or another factor.

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

important to note that EPA does not prohibit reporting how nonepistemic aims/values influenced the selection of the epistemic research problem being addressed; scientists are free to report this, and, in applied contexts where sensitivity to stakeholder aims and values is seen as desirable, scientists might often *choose* to do so.

Finally, a fifth objection charges that EPA is vague and overly permissive. EPA, like other epistemic constraint approaches, requires that research respect basic constraints on epistemically adequate science, but one might worry that not enough has been said about what these basic constraints are. Moreover, merely respecting basic constraints might set the bar too low; research might be deserving of criticism, even if it does not violate basic constraints on epistemically adequate science.

Response. In reply to the former concern, it is not difficult to specify some basic constraints on epistemically adequate science, as noted earlier: no fabricating data, no use of methodologies that virtually guarantee a predetermined conclusion, and so on. Nevertheless, as Daniel Steel (2017) notes, it is difficult to draw a sharp line between scientific activities that meet basic constraints on epistemically adequate science and those that don't—there can be borderline cases. There is also the question whether basic constraints vary at all with the research context.[15] So it is fair to say that there is more work to be done in clarifying the nature of, and in articulating, basic constraints on epistemically adequate science. When it comes to the latter concern, it is important to remember that EPA calls for scientists to strive to conduct research that is epistemically excellent. In addition, it is possible that relatively demanding basic constraints on epistemically adequate science can be articulated. Steel, for instance, identifies as a basic constraint the *weak severity criterion*, which requires that a study purporting to test a hypothesis be likely to issue a negative result if the hypothesis is false (see also Steel and Whyte 2012). Finally, research that follows the transparency requirements of EPA can be critically scrutinized (see also Elliott and McKaughan 2014). It is vital that such methodological scrutiny be performed (Mayo and Spanos 2006).

5. Rapprochement, revisited

Following EPA, research will be responsive to nonepistemic aims and values in the ways that many opponents of the VFI consider important, even as consistency with the VFI is maintained. There is thus hope that EPA is an approach that both (some) opponents and (some) proponents of the VFI could accept. Such a rapprochement would be more substantial than the one envisioned by Lusk and Elliott (2022), insofar as parties would be agreeing on a full-fledged account of how nonepistemic aims/values can appropriately influence scientific investigation, not just on a background assumption about the nature of scientific assessment. And it would not require that we abandon the practice of evaluating the truth of "plain" scientific hypotheses, such as "Earth's climate is warming."

Nevertheless, it is clear that EPA will not satisfy everyone. It will not satisfy opponents of the VFI who find epistemic constraint approaches too restrictive and want science to more directly serve social ends. EPA could satisfy those who subscribe to the VFI because they think it is simply *wrongheaded* (or a category mistake) to

[15] Thanks to an anonymous referee for raising this possibility.

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

appeal to nonepistemic aims and values—reflecting what someone wishes or desires to be the case—as justifying reasons for choices made during an activity (research) whose aim is to find out what is the case. But EPA will not satisfy those who subscribe to the VFI because they think that nonepistemic aims and values should *make no difference* to how researchers pursue a given epistemic goal; under EPA, nonepistemic aims and values do still make a difference to how researchers pursue epistemic goals, by shaping the epistemic desiderata specified in research problems in the problem selection stage of investigation. For such make-no-difference proponents of the VFI, however, EPA should be a source of dismay. EPA respects the letter of the VFI yet can deliver research that is just as responsive to nonepistemic aims and values as approaches that permit "internal" value influence; it amounts to a demonstration that the VFI, as standardly formulated, is insufficient for the value-policing work that make-no-difference proponents of the VFI desire.

It is worth reflecting on the source of this failure. A number of critics have argued that the VFI's limited focus on the internal/research stage of investigation is misguided. They show how choices made in the problem selection stage of investigation can impact research in epistemically important (and even problematic) ways. For instance, Kevin Elliott and Dan McKaughan (2009) call attention to ways in which the selection of research questions and projects can shape which data are available to scientists and thus the conclusions they reach. Others have made similar points (see, e.g., Lacey 1999, 2013; Bluhm 2017; Winsberg 2018, chap. 9; Harvard and Winsberg 2022; Carrier 2022). Anke Bueter (2015, 18) offers a diagnosis: "The underlying problem is that the value-free ideal is essentially based on a further premise often not stated explicitly: the idea of a clean division between the inside and the outside of the epistemic process . . . . This premise is false because non-cognitive values influencing decisions traditionally regarded as part of the context of discovery (e.g., decisions on funding, choice of questions, theory pursuit, or significance ascriptions) can have an indirect effect on the evaluation of theories." To be clear, these critics are not arguing that the VFI should be expanded to prohibit value influence even beyond the internal/research stage; indeed, many of them would argue that value influence "internal" to research is sometimes appropriate. Their point is rather that, if the goal is to prevent problematic value influence in science, then the VFI, focused only on what happens in the internal/research stage of investigation, is insufficient.

We can flesh out the diagnosis further with the help of Zina Ward's (2021) analysis of types of value influence. The VFI prohibits appealing to nonepistemic values as *justifying reasons* for methodological choices in the internal/research stage of investigation. It does not prohibit values from being *mere causal contributors* to such choices, because that would make a huge proportion of scientific research illicit, simply because the question or problem under investigation was selected in part for its relevance to nonepistemic concerns. Nearly all research on COVID-19, for instance, would be in violation of the VFI. The problem for make-no-difference proponents of the VFI, illustrated here via EPA, is that strategic choices in the problem selection stage of investigation—mere causal contributors—can sometimes serve nonepistemic aims/values just as well as "internal" methodological choices that are justified by direct appeal to those aims/values. Indeed, strategic choices even farther causally upstream, such as institutional-level policies about how a pot of funding will be distributed, can give rise to science that is biased in

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

just the way that proponents of the VFI (and many others) would like to prevent (Bueter 2015; Holman and Bruner 2017).

Ultimately, then, EPA is more bad news for make-no-difference proponents of the VFI, even as it is an approach that various other parties to the values-in-science debate could find acceptable.

6. Concluding remarks

EPA is a proactive approach to making research responsive to nonepistemic aims and values without violating the VFI. When EPA is enacted successfully, choices made in the "internal" stage of investigation will facilitate pursuit of targeted nonepistemic aims, or be responsive to salient background values, but will not be justified by appeal to those aims/values. Researchers will strive to conduct research that is of excellent epistemic quality and, at a minimum, will respect basic constraints on epistemically adequate science. EPA is thus an approach that both (some) opponents and (some) proponents of the VFI could find acceptable.

Besides this potential for rapprochement, EPA promises to be useful in scientific practice. In some scientific contexts, such as the study of climate change impacts, there is increasing interest in taking account of stakeholder aims and values in research yet significant uncertainty about how to do this, and how to do it appropriately. EPA provides concrete guidance. Most notably, it calls for researchers to consider, at the outset of investigation and in collaboration with stakeholders, whether it would be advantageous for research to have particular epistemic characteristics (related to inductive risk, prioritization, methods, and form of conclusions). In some applied contexts, EPA might even be operationalized as a series of domain-specific questions for researchers to consider each time they begin an investigation. The fact that EPA maintains a firm commitment to the idea that science is an *epistemic* endeavor—limiting how values may shape scientific investigation and encouraging researchers to decline to perform research that they judge will be of insufficient epistemic quality—can be expected to enhance its appeal among practitioners.

Although several objections to EPA were considered, none spoke decisively against it. Given that it is an approach that parties on both sides of the values-in-science debate could find acceptable and, moreover, that it promises to be useful in scientific practice, EPA merits serious consideration as an account of how nonepistemic values can appropriately influence scientific investigation.

Acknowledgments. I thank Casey Helgeson, Eric Winsberg, Sam Fletcher, Deborah Mayo, Gil Hersch, Dan Hoek, Karen Kovaka, and three anonymous reviewers for helpful comments. The article also benefited from discussion with Zoom audiences at talks hosted by Penn State University, Johannes Gutenberg University Mainz, and Stockholm University.

References

Anderson, Elizabeth S. 2004. "Uses of Value Judgments in Science: A General Argument, with Lessons from a Case Study of Feminist Research on Divorce." *Hypatia* 19 (1):1–24. https://doi.org/10.1111/j.1527-2001.2004.tb01266.x.

Bluhm, Robyn. 2017. "Inductive Risk and the Role of Values in Clinical Trials." In *Exploring Inductive Risk: Case Studies in Values in Science*, edited by K. C. Elliott and T Richards, 171–91. New York: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780190467715.003.0010.

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

Brown, Matthew. 2017. "Values in Science: Against Epistemic Priority." In *Current Controversies in Values in Science*, edited by K. C. Elliott and D. Steel, 64–79. New York: Routledge.
Bueter, Anke. 2015. "The Irreducibility of Value-Freedom to Theory Assessment." *Studies in History and Philosophy of Science, Part A* 49:18–26. https://doi.org/10.1016/j.shpsa.2014.10.006.
Carrier, Martin. 2022. "What Does Good Science–Based Advice to Politics Look Like?" *Journal for General Philosophy of Science* 53:5–21. https://doi.org/10.1007/s10838-021-09574-2.
Douglas, Heather E. 2000. "Inductive Risk and Values in Science." *Philosophy of Science* 67 (4):559–79. https://doi.org/10.1086/392855.
Douglas, Heather E. 2009. *Science, Policy, and the Value-Free Ideal*. Pittsburgh, PA: University of Pittsburgh Press. https://doi.org/10.2307/j.ctt6wrc78.
Douglas, Heather E. 2017. "Why Inductive Risk Requires Values in Science." In *Current Controversies in Values in Science*, edited by K. C. Elliott and D. Steel, 81–93. New York: Routledge. https://doi.org/10.4324/9781315639420-6.
Elliott, Kevin C. 2017. *A Tapestry of Values: An Introduction to Values in Science*. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780190260804.001.0001.
Elliott, Kevin C., Aaron M. McCright, Summer Allen, and Thomas Dietz. 2017. "Values in Environmental Research: Citizens' Views of Scientists Who Acknowledge Values." *PLoS ONE* 12 (10):e0186049. https://doi.org/10.1371/journal.pone.0186049.
Elliott, Kevin C., and Daniel J. McKaughan. 2009. "How Values in Scientific Discovery and Pursuit Alter Theory Appraisal." *Philosophy of Science* 76 (5):598–611. https://doi.org/10.1086/605807.
Elliott, Kevin C., and Daniel J. McKaughan. 2014. "Nonepistemic Values and the Multiple Goals of Science." *Philosophy of Science* 81 (1):1–21. https://doi.org/10.1086/674345.
Elliott, Steve. 2022. "Research Problems." *British Journal for the Philosophy of Science* 72 (4):1013–37. https://doi.org/10.1093/bjps/axz052.
Harvard, Stephanie, and Eric Winsberg. 2022. "The Epistemic Risk in Representation." *Kennedy Institute of Ethics Journal* 32 (1):1–31. https://doi.org/10.1353/ken.2022.0001.
Holman, Bennett and Justin Bruner. 2017. "Experimentation by Industrial Selection." *Philosophy of Science* 84 (5):1008–1019.
Intemann, Kristen. 2015. "Distinguishing between Legitimate and Illegitimate Values in Climate Modeling." *European Journal for Philosophy of Science* 5:217–32. https://doi.org/10.1007/s13194-014-0105-6.
Lacey, Hugh. 1999. *Is Science Value Free?* London: Routledge.
Lacey, Hugh. 2013. "Rehabilitating Neutrality." *Philosophical Studies* 163:77–83. https://doi.org/10.1007/s11098-012-0074-6.
Lusk, Greg, and Kevin C. Elliott. 2022. "Non-epistemic Values and Scientific Assessment: An Adequacy-for-Purpose View." *European Journal for Philosophy of Science* 12:35. https://doi.org/10.1007/s13194-022-00458-w.
Mayo, Deborah G. 1996. *Error and the Growth of Experimental Knowledge*. Chicago: University of Chicago Press. https://doi.org/10.7208/chicago/9780226511993.001.0001.
Mayo, Deborah G., and Aris Spanos. 2006. "Philosophical Scrutiny of Evidence of Risks: From Bioethics to Bioevidence." *Philosophy of Science* 73 (5):803–16. https://doi.org/10.1086/518630.
Parker, Wendy S. 2020. "Model Evaluation: An Adequacy-for-Purpose Approach." *Philosophy of Science* 87 (3):457–77. https://doi.org/10.1086/708691.
Parker, Wendy S., and Greg Lusk. 2019. "Incorporating User Values into Climate Services." *Bulletin of the American Meteorological Society* 100 (9):1643–50. https://doi.org/10.1175/BAMS-D-17-0325.1.
Parker, Wendy S., and Eric Winsberg. 2018. "Values and Evidence: How Models Make a Difference." *European Journal for Philosophy of Science* 8:125–42. https://doi.org/10.1007/s13194-017-0180-6.
Reiss, Julian, and Jan Sprenger. 2020. "Scientific Objectivity." In *The Stanford Encyclopedia of Philosophy*, Winter 2020 ed., edited by Edward N. Zalta. https://plato.stanford.edu/archives/win2020/entries/scientific-objectivity/.
Schroeder, S. Andrew. 2021. "Democratic Values: A Better Foundation for Public Trust in Science." *British Journal for the Philosophy of Science* 72 (2):545–62. https://doi.org/10.1093/bjps/axz023.
Steel, Daniel. 2017. "Qualified Epistemic Priority: Comparing Two Approaches to Values in Science." In *Current Controversies in Values in Science*, edited by K. C. Elliott and D. Steel, 49–63. New York: Routledge. https://doi.org/10.4324/9781315639420-4.

https://doi.org/10.1017/psa.2023.107 Published online by Cambridge University Press

Steel, Daniel, and Kyle Poway Whyte. 2012. "Environmental Justice, Values, and Scientific Expertise." *Kennedy Institute of Ethics Journal* 22 (2):163–82. https://doi.org/10.1353/ken.2012.0010.

Ward, Zina. 2021. "On Value-laden Science." *Studies in History and Philosophy of Science* 85 (2):54–62. https://doi.org/10.1016/j.shpsa.2020.09.006.

Winsberg, Eric. 2018. *Philosophy and Climate Science*. Cambridge: Cambridge University Press. https://doi.org/10.1017/9781108164290.

Cite this article: Parker, Wendy S. 2024. "The Epistemic Projection Approach to Values in Science." *Philosophy of Science* 91 (1):18–36. https://doi.org/10.1017/psa.2023.107