Philosophy of Science (2026), **93**, 920–929
doi:10.1017/psa.2026.10215

PSA PHILOSOPHY OF SCIENCE ASSOCIATION

PHILOSOPHY OF SCIENCE

DISCUSSION NOTE

Stuck Between a Rock and a Hard Science: Three Challenges with Putting Epistemic Projection to Work

Kelli Barr[1] and Katharina Bernhard[2]

[1]Department of Philosophy, University of California, Davis, USA and [2]Department of Philosophy, University of St Andrews, UK
Corresponding author: Kelli Barr; Email: barr@ucdavis.edu

(Received 09 August 2025; revised 14 April 2026; accepted 14 April 2026; first published online 26 May 2026)

Is it permissible to implicate non-epistemic values in the 'core' of scientific reasoning? This paper offers an appraisal of Parker's (2024) Epistemic Projection Approach (EPA) prescribing acceptable relations between epistemic and non-epistemic values in research conduct. Non-epistemic values are granted a role in determining epistemic characteristics for a research project but are meant to be 'screened off' from core research decisions via *epistemic projection*. This discussion takes the EPA for a test drive. We identify three challenges to its conceptual and practical viability and show how EPA currently lacks resources to address them, indicating a need for further development.

Project(ion) Management

(Barr and Bernhard speaking together)[1]

The value-free ideal has been declared dead in the literature on values in science. But this has not deterred efforts to construct a more viable version (see Brown 2024 for a recent critical review). Parker (2024) provides what looks like a middle way: an ideal of value-*responsiveness*. Non-epistemic aims and/or values take center stage in planning a research project. Participants – including, but not limited to, professional researchers – engage in a process that aligns epistemic characteristics of the proposed research project with the non-epistemic aims and/or values (moral/ethical, political, social, prudential, or practical/pragmatic) motivating the project[2] in a way that

[1] This discussion grew out of conversations between the authors. It's structure attempts to conserve that spirit, inspired by Lugones and Spellman (1983). Each section indicates whether the authors are speaking together or individually, preserving both points of convergence and divergence between our philosophical interests and perspectives.

[2] Since the EPA does not include a stage or procedure for deliberation about which non-epistemic aims and/or values ought to inform the research, the presumption seems to be that these have been identified and agreed on in advance.

© The Author(s), 2026. Published by Cambridge University Press on behalf of Philosophy of Science Association. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

'screens off' those values from influencing decisions in 'core' epistemic practices, including inductive risk judgements or other trade-off prioritizations, thereby preserving a form of value-freedom. Parker's Epistemic Projection Approach (EPA) prescribes accomplishing this via the process of Epistemic Projection (EP) and three additional criteria (26-30). The point is to pre-register, in a way, those epistemic elements selected for a project to satisfy broader non-epistemic aims. Much like pre-registration constrains researchers' degrees of freedom, EP fixes in advance the relevant epistemic standards and criteria meant to justify design choices and adjudicate decisions while conducting the research. The EPA offers a means for non-epistemic aims and values to permissibly influence researchers' epistemic choices but not to *justify* them.

The EPA divides the research process into discrete stages. First is a planning stage – "Problem Selection" – in which participants specify research questions, epistemic goods (e.g., knowledge, accuracy, understanding), goals, and desiderata, which includes constraints (i.e., minimal standards of conduct) and preferences relevant to achieving the goals within those constraints (e.g., preferences about inductive risk trade-offs). Here, aims and desiderata may be either epistemic, non-epistemic, or mixed. They are then translated into a set of purely epistemic desiderata through the projection process, in which one "reason[s] from a broader aim to an epistemic characteristic of research that *aligns with* that broader aim" (p. 21). Parker's analogy with geometric projection, though still opaque, seems to go like this: Just as projecting a three-dimensional shape onto a plane allows you to see what a shape looks like in fewer dimensions, Epistemic Projection shows what a project that satisfies your non-epistemic aim(s) can 'look like' in terms of its epistemic characteristics only. Epistemic goals and desiderata "align with" the non-epistemic aim(s) in the sense of "being relevant to", or "helpful for," its achievement (22). These epistemic elements constitute an "Epistemic Research Problem," which functions like design specifications for the subsequent "Research" stage, in which the project is further specified and executed. Non-epistemic values are permitted a motivational or justificatory role (Ward 2021) in the planning stage. However, the EP process is meant to ensure that they do not guide research decisions *in situ*. The resulting epistemic research project is based only on projected epistemic goals and desiderata. Thus, the research project is value-responsive but can proceed without reference to non-epistemic aims/values.

Our goal is to take the EPA for a test drive. We are sympathetic to Parker's value-responsiveness ideal, and the brief hypotheticals showing how the EPA could work are tempting. We wish to see more clearly what happens when the rubber hits the road. Specifically: what is gained and lost, both conceptually and practically, by screening off non-epistemic values in this way? As we will see below, on both fronts not much is gained and precious much is lost.

A dilemma walks into a research project . . .

(Barr speaking)

One virtue of scientific inquiry is that it is open ended. One doesn't know in advance what exactly will be learned or which choices will arise, and how, while working on a project. Open-endedness presents a value management problem for the EPA. For the projection method to 'screen off' non-epistemic aims and desiderata

successfully, the pre-specification of epistemic aims and desiderata must be sufficiently relevant *and* decisive in those moments when choices arise during the research execution stage. This raises a procedural question: what is one to do when new information comes to light that requires realignment with existing non-epistemic commitments, or even reveals new ones?

To explore this question, let's consider the following historical case from Worrall (2008). Persistent pulmonary hypertension of the newborn (PPHN) is a relatively rare condition (about 2/1000 infants per year in surveyed populations), but one for which the conventional treatment in the 1970s had only about a 20% survival rate. Researchers in the UK working on extracorporeal membranous oxygenation (ECMO) had strong observational evidence for its effectiveness in treating PPHN by 1982: 80% of infants receiving the new treatment lived. Yet, they designed an RCT to establish its efficacy against the conventional treatment in evidential terms that they thought would warrant changing the treatment standard. A total of 12 infants were enrolled using a modified 'play the winner' randomization, which resulted in a single non-ECMO infant death. However, controversy over the randomization procedure prompted another trial using unweighted randomized assignments and a stopping rule of four infant deaths. All four infants who died were assigned to the conventional treatment. Still, statisticians at the time argued that the resulting sample size was insufficient to achieve statistical power by conventional gold standards. A third RCT was performed in the early 90s, this time with a broader recruitment pool of infants with PPHN or a related respiratory condition. In total, 185 infants were recruited, but the trial was ended prematurely because of what was seen as a high number of infant deaths – again, overwhelmingly in the conventional treatment group.

How many infants is enough? There are risks in both directions: too few infants compromises the rigor of the study, which risks more future deaths if the conclusions turn out not to be true; but too many infants risks needless additional deaths. Importantly, the timing of such decisions can matter just as much as their justification, for one does not know in advance how many will actually die. This is just what happened in the final ECMO trial. The study was ended before reaching the predetermined recruitment threshold because of likely harms to additional infant recruits, even though one epistemic goal of the trial was to secure a larger, more statistically robust sample size. Researchers (and research regulators) learned information in the research stage that, because of their non-epistemic values (chiefly: preventing preventable infant deaths), changed their assessment of what the relevant epistemic constraints ought to be and just when they had been satisfied.

Challenge 1: Turning the Projection Crank Invites a Dilemma

Would ending the trial prematurely to prevent additional infant deaths be a permissible decision under the EPA framework? Parker is clear that the process of EP alone will not get us "epistemically adequate science" (27), meaning science that produces "knowledge and other epistemic goods" (27).[3] The EPA provides three additional criteria such that research is both value-responsive and high quality. I will

[3] Some examples Parker gives of basic epistemic constraints: the research is "genuinely probative", "data should not be fabricated" or "discarded simply because they fail to support a preferred conclusion," and "conclusions should be informed by the evidence collected" (27). Thanks to an anonymous reviewer

focus on the second: that "scientists, qua scientists, will need to *respect basic constraints on epistemically adequate science*" (29, emphasis original).[4] Non-epistemic, epistemic, and mixed desiderata must respect constraints for doing rigorous work; only then is it permissible to project them into the epistemic research program, which determines what choices are made while executing the research project. Fulfilling this criterion requires that, in the research stage, "researchers will not justify their methodological choices by appeal to non-epistemic aims or values; they will justify them as a means to solving the study's epistemic research problem, maintaining consistency with the VFI" (29). I will refer to this as *an implied practical prohibition*. It is crucial for the success of the EPA.

In the ECMO case, staying faithful to pre-registered epistemic constraints would invite a dilemma. You might say that the researchers could simply turn the proverbial projection crank again, returning to the planning stage to re-project their non-epistemic commitments into a new set of epistemic characteristics that would inform redesigning the research project to have a lower recruitment threshold. Then they would be permitted to end the study in accordance with the new set of epistemic constraints. Note, however, that deciding to turn the crank again in situ is (1) a judgment made in the research stage that, in this case, (2) was justified on non-epistemic grounds. But the EPA's practical prohibition *does not permit this decision*.[5] And so, deciding to turn the projection crank again, would violate a key provision of the EPA, regardless of the compelling moral justification.

Could they have avoided such a violation? Perhaps they could deliberate more thoroughly in the planning stage, anticipating many possible forks in the research path and associated kinds of evidence that could give rise to compelling non-epistemic justification for revising decisions about constraints and desiderata (more on non-evidential justification in Bernhard's discussion below). The result would be something like a list of anticipated paths and associated value-aligned epistemic constraints and preferences. Or: they could 'pre-register' in the problem selection stage rules or criteria by which to judge when to re-project. If a problem emerging in the research execution stage is covered by those rules, this provides permissible epistemic grounds to 'turn the crank again'.

But full anticipation of open-ended future events is impossible. In the planning stage, ECMO researchers addressed the question, how many infants must be enrolled in the study? The discovery of how many infants were dying raised a different

for pushing to include these details and clarify their relation to Parker's understanding of scientific quality.

[4] It's unclear from Parker's discussion how we are to evaluate the results of research that violates the implied practical prohibition. If results are potentially biased by wishful thinking (which Parker thinks is the danger posed by non-epistemic value influence), should they be consigned to the dustbin? I thank an anonymous reviewer for raising this question. See footnote 7 for more on this issue; see also Challenges 2 and 3 for elaboration on the difficulties with what exactly should count as wishful thinking.

[5] Parker claims that the EPA permits revisions to the epistemic research problem "if the targeted [non-epistemic] 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" (2024, 28). The ECMO case is different because executing the research project yields new information in situ that requires the researchers (not the stakeholders) to revisit previously made value-assessments at the time such information comes to light.

question: how many infant deaths is enough? The risk of unanticipated discoveries will always, and indeed *must*, remain. Anticipatory planning does not get you omniscience. But also, as Havstad and Brown (2017) show in greater detail, decisions like the one ECMO researchers faced cannot be deferred until after the research has concluded. There is no general, *a priori* way to predict when one might need to turn the projection crank again. Absent omniscient planning, the dilemma stands. Call this the *fidelity dilemma:* the EPA framework would require in this case choosing between a non-epistemically prescribed but practically impermissible violation of value-freedom ('practical' as in: according to EPA's implied prohibition) and a practically prescribed but non-epistemically impermissible obedience to value freedom.[6]

The Fidelity Dilemma Undermines EPA's Value-Responsiveness

Unfortunately, the EPA does not provide resources for breaking out of the dilemma. Achieving the value-responsiveness ideal rests on the practical and epistemic viability of the practical prohibition that non-epistemic values shall never be used to justify decisions in the research execution stage. Had ECMO researchers been operating with the EPA framework, the moral demand to end the trial early would have run afoul of the practical prohibition, forcing a choice between saving rigor and saving lives. Remaining faithful to the EPA's screening-off requirement secures value-freedom in the research execution stage, but at the expense of making value-sensitive judgments when relevant contingencies arise, judgments that are necessary to achieve the research project's multiple, mixed purposes.[7] This invites the fidelity dilemma and is a significant practical limitation of the framework. The EPA's preservation of value-freedom is presented as a virtue. The worry, however, is that value-freedom comes at too high a price: one is prohibited from acting on pertinent ethical, moral, or prudential justifications for a change of course in the research project, even when that change of course would best satisfy the motivating non-epistemic aims or values.

. . . and the researcher asks, "what'll you have?"

(Bernhard speaking)

Arguably, the Fidelity Dilemma is a direct consequence of precluding non-epistemic value use from the core of scientific research. But are there good epistemic reasons for such a preclusion? Parker claims that the reasons are "straightforward" (33). I take the following three assumptions to summarize the idea (see esp. 33):

[6] The ECMO case might seem extreme, but Havstad and Brown (2017) show that the dilemma would stand even in cases where the relevant decision involves modifying epistemic constraints and desiderata rather than ending the research project prematurely. Thanks to an anonymous reviewer for suggesting this connection.

[7] Both authors believe that, on the EPA, violations of the practical prohibition in the ECMO case are best understood as undermining what the framework specifically holds to as minimal epistemic adequacy – perhaps unlike with pre-registration, where epistemically justified deviations do not necessarily undermine the quality or value of the results. Bernhard argues below that it is a problematic assumption on the EPA that value-influences must come at epistemic cost (Challenge 2 and 3) and concludes that the EPA does not generate the promised epistemic benefits Parker claims result from screening-off non-epistemic values. Even if one accepts the assumption about costs, Bernhard shows that the EPA cannot epistemically offset the practical limitations of Barr's Challenge 1. Thanks to an anonymous reviewer for showing that this connection should be highlighted.

(1) Non-epistemic values "reflect what one wishes or desires to be the case".
(2) In virtue of (1) they are irrelevant to choices pertaining to what *is* the case.
(3) When occurring in the research/core stage of science, they are not only irrelevant for believing or accepting some claim P, but potentially harmful, introducing unacceptable bias and distortion.

One straightforward way to explicate the idea behind (1)-(3) is this: Your non-epistemic values should never count as reason for believing/accepting that P in just the same way that your desire that there be coffee left in your mug should never count as reason for believing/accepting that there is. Put differently: how you want/desire/prefer the world to be shall not ever serve as reason for your belief or credence in *that the world is so* (2) because desires just are not the kind of attitudes that can (or should) do that (1).

Arguably, non-epistemic values serving such a role entails that values are desire-like attitudes and that either (a) one uses values *as reasons* in terms of *evidence* (challenge 2 below), or (b) non-epistemic values overwrite one's epistemic grounds (challenge 3 below). I am setting aside whether non-epistemic values (or epistemic values for that matter) always (or ever) present in the form of desires/preferences. For, even assuming they are, it is pivotal to note that – spare some more radical pragmatists – proponents *and* most opponents of the value-free ideal will agree that some interpretation of wishful thinking of the sorts of (a) and (b) are epistemically impermissible in the internal stage of science, and thus agree with some versions of (3)[8]. In fact, such wishful thinking should not be permissible in any epistemic endeavor. We can say more specifically, then, that screening off non-epistemic values from the research stage is intended to forestall either using desires as evidence (a) or having desires overwrite one's epistemic grounds (b), or both. The following two challenges address EPA's performance against the backdrop of these epistemic goals.

Challenge 2: EPA Impermissibly Excludes Epistemically Permissible Uses of Non-Epistemic Values

In terms of graded doxastic attitudes, treating your desire/wish as evidence means that the fact that you desire/wish there was coffee left in your mug directly increases your probability (credence or otherwise)[9] that this is so. Even for many opponents of the value-free ideal such direct value use will be impermissible: desiring coffee to be in your mug indeed is not evidence for that there is. Contrast this with indirect value use, often associated with inductive risk trade-offs: Assume it's in fact more likely than not that you ran out of coffee and your credence that there is *no* coffee left in your mug is, say, 0.79. Is that credence high enough for you to believe/accept that there is no coffee in your mug? There is a 0.21 probability that you have some coffee left. You could suspend judgement on the matter (and investigate your mug).

[8] Part of the trouble with reconstructing Parker's worry is that there is little epistemological substantiation provided for it. I believe (a) and (b) to be charitable and plausible substantiation-attempts for (1)-(3). Thanks to two anonymous reviewers for pushing for clarification of this point.

[9] I presuppose a broadly Bayesian framework here, as does Parker (2014). (See also Parker and Winsberg 2018). I moreover take probabilities to be personal probabilities (credence) subject to acceptance (cf. Steel 2015).

Whatever you decide, there is no magic epistemic/objective threshold that can settle the issue for you.[10] Notably, instead of serving as evidence, your coffee desire (your non-epistemic value) enters in determining a threshold for how high a probability or how good your evidence determining it must be. If your epistemic position puts you above the threshold, believing/accepting that there is no coffee in your mug is justified; if not, suspending judgment on the matter is the epistemically prudent choice. Setting this threshold entails weighing how bad it would be to falsely assume that there is no coffee left in your mug against both how good it would be if you rightly believed/accepted there wasn't and the good and bad consequences of suspending judgement. If the worst that can happen is that you take a sip from an empty cup, 0.79 seems plenty credence for belief/acceptance. However, if I offer to pour orange juice into your mug, then a 0.21 probability of ending up with a ghastly blend of coffee and orange juice might seem plenty for suspending judgement. Non-epistemic values thus can have epistemic impact: they can make the difference between believing/accepting and suspending your coffee claim. But the (un-)desirability of outcomes crucially need not affect the credence you have *directly*. It merely affects whether your credence is judged as epistemically sufficient for acceptance. The effect is *indirect*.

One may debate whether such indirect use is merely causal or really justificatory (pace Ward 2021) – one's epistemological taste will be decisive here. Either way, it is this kind of indirect, threshold determining use of non-epistemic values that Parker's EPA explicitly excludes. The coffee example is analogous to inductive risk cases which Parker seeks to transform, through EP, into predetermined epistemic prioritizations. It also captures the kind of trade-offs in climate modelling that Parker discusses in her "Alex" example, where deciding in the planning stage between two model parameters for representing physical variable X involves a trade-off between increasing accuracy on X and maintaining the model's performance on physical variable Y (20). Crucially, this is a choice between epistemically permissive options. On the EPA, one option will have been pre-selected as adequate-for-purpose in the planning stage. But, while epistemically consequential, things are not intrinsically epistemically better or worse whether Alex' decision is based on *prior agreement* to prioritize X over Y (EPA), on *the desire for accuracy* of X over Y (epistemic value use), on *the desire to prevent future harm* (non-epistemic value use), or on the result of *tossing a coin*. If (a), using non-epistemic values as evidence, was the worry, then the EPA is too exacting. Only direct uses are epistemically problematic in the sense of (1)-(3). Yet the EPA prohibits indirect use all the same. And it is not exacting enough, since it allows for direct epistemic value use.

Challenge 3: EPA Impermissibly Includes Problematic Uses of Epistemic Values

What about (b), value-use that overwrites one's epistemic grounds? Your craving for coffee is so overwhelming that you cannot help but attach an insanely high

[10] Douglas (2017), too, argues that epistemic values cannot determine an evidential sufficiency threshold. But my point is more specific: Even a 0.5 credal threshold typically does not hold any default epistemic significance for acceptance/belief thresholds. See Hempel (1962), Levi (1967), Harsanyi (1985), and Wilholt (2009). This is typically so for accounts where thresholds are *derived* and suspension is variable or an interval – regardless of whether pragmatic or epistemic values/utilities determine that threshold. Recent discussions in epistemic utility theory further attest to this (e.g. Easwaran 2016). I am grateful to two anonymous review comments pushing for clarification.

desirability/utility to believing/accepting that there is coffee left in your mug and insanely high undesirability/disutility to falsely believing/accepting no coffee is left (while caring little about remaining outcomes). Then, no matter how low your credence that you have coffee left, the threshold for believing/accepting that there is coffee left in your mug will be so low that you can't help but impermissibly believe/accept it. Your desires are still not directly evidential, but they outstrip the significance of your epistemic grounds. You end up in a situation comparable to wishful thinking – you end up in case (b). Pascal's Wager is perhaps the most discussed case of that kind. Cases like this are characterized by the extreme utilities used across some outcomes combined with low probability states such that utilities swamp credence. The riddles associated with extreme utility attributions arguably are a consequence of the expected utility framework that underpins it. Being primarily formal artifacts, formal constraints are thus the obvious containment measures. And so if (b) was the worry behind (1)-(3), it would be based on a generic bugbear for expected utility theory that rendered indirect value use problematic.[11] Additionally, even if such cases were realistic troublemakers, they would be so for non-epistemic and epistemic value-use alike: Whether you wish p to be true because of the sheer explanatory power p would have or because justice would be served if p was true, again, makes no difference. Both can drive extreme cases.

In summary, challenge 2 shows that in attempting to avoid (a), EP precludes too much (prohibiting epistemically permissible *indirect* value use). If EPA attempted to avoid (b), it would respond to mere theoretically extreme case of indirect value use. Moreover, both (a) and (b) fail to identify epistemic value-use as equally worrisome. Crucially, the kind of value-use Parker's EPA excludes on little convincing epistemic grounds are the kind of value-use that, if prohibited, can give rise to the fidelity dilemma.

Maybe this result should not surprise. The idea that excluding non-epistemic values to salvage epistemic purity and save science from wishful thinking has been debunked more often than can be referenced here. Douglas' (2009) case for "indirect value use" is a particularly prominent example and Anderson's (1995) discussion of significance in inquiry a particularly forceful one. Challenges 2 and 3 are just variations on these all too familiar discussions.

Conclusion: On Having the Right Tools for a Worthwhile Job

(Barr and Bernhard together)

It's noteworthy that Parker's EPA screens off not only non-epistemic values but all value trade-offs. Complex value considerations restricted to the planning stage are translated through EP into simple 'epistemic prioritization' rules of the kind: 'If confronted with a choice between increasing accuracy on variable X at the expense of Y or the reverse, choose to increase accuracy on X'. Put more simply: 'If either A or B,

[11] Pascal's Wager type cases (see e.g. Hájek 2003) and, relatedly, the St. Peterburg Paradox (see Peterson 2023 for overview) test the boundaries of expected utility theory's framework itself. That is: absent appropriate formal bounds, the legitimacy of any and all uses of the expected utility machinery are equally affected, which includes vindicating probabilism. This global challenge persists independent of local questions around epistemic or non-epistemic value-use. Extreme utility attributions arguably are ill-suited to substantiate the worry that non-epistemic value-use thus leads to wishful thinking. I am grateful to an anonymous reviewer for patiently pressing that this be clarified.

prefer A'. This translation might ward against Pascalian-type cases in research, since it in principle prevents such a wager from emerging in the research stage. Perhaps, then, EPA's primary epistemic function is not to screen off non-epistemic values, but to translate complex value considerations into simple epistemic priority rules and, by this means, foreclose extreme cases of value use.

But if that's the motivation behind the EPA, are the presumed epistemic threats worth all that is lost in translation? We hope this discussion (esp. challenge 3) has shown that it is not. Precluding any complex reasoning that might involve non-epistemic values from the research stage is ill-motivated both practically (challenge 1) and epistemically (challenge 2). On the epistemic side: Engaging with *how much* one ought to prefer one option over another can be epistemically permissible when the impact of values is indirect. Merely knowing *that* one ought to prioritize one option over another need not (and cannot) replace having to make a choice (cf. Havstad and Brown 2017). As a matter of principle, no simple ordinal ranking of the sort 'If either A or B, prefer A' will generate value-free choices through mere pre-selection (per challenge 2).

The worry about information loss carries over to the practical side: There is a need to consider non-epistemic matters during research execution and to ask 'how much' questions in situ, in meaningful, context-sensitive ways. Again, simpler rules will not be informative enough to be decisive. If the sole job of value-free research was to ward off extreme cases of wishful thinking, then not only is EP an unsuitable tool, but: executing research following the EPA may result in a morally dubious job. For, asking *how much* questions is required to discharge one's moral responsibility during research (Douglas 2009). Consider the last round of RCTs in the ECMO case. The doctors performing conventional treatment on the control group were aware that ECMO promised a much higher rate of survival. Arguably, a scientist treating an infant in the control group must be able to ask: "Is this infant's likely death still morally permissible?" And their answer to that question must be able to inform their course of action. Yet, the EPA would prohibit acting on this question during research. One could only ask: "Does this infant's likely death satisfy the agreed upon epistemic prioritizations?" Moral responsibility is thus made subservient to safeguarding epistemic purity, but on unconvincing epistemic grounds and in a way (as challenge 1 shows) that undermines the value-responsiveness ideal.

We want to end this discussion by highlighting benefits of what Parker's approach is meant to accomplish, even if they do not outweigh our estimation of EPA's shortcomings. EP aims to systematize and harmonize value alignment throughout research in a transparent, and thus comparable, way. There are benefits to asking how values may articulate in epistemic priorities, especially if one is concerned with explicit deliberation about alignment between epistemic and non-epistemic commitments. The idea of projection is not without merit for promoting value coordination. Thus, we mean this discussion to encourage further work on EPA-like accounts and other efforts at philosophically robust attention to the role of values in research design and execution.

Acknowledgements. For their helpful conversations and comments on the project, we thank: the Arché Research Centre, University of St Andrews; the "Measurement & Values in the Human Sciences" workshop at Cambridge HPS, Nov. 20, 2024; Wendy Parker, Dejan Makovec, Cristian Larroulet Philippi, Ahmad Elabbar, Torsten Wilholt, Marina DiMarco, Jade Fletcher, Kevin Elliott, and Matthew Brown. K. Barr thanks especially Jim Griesemer and Elihu Gerson for their ongoing support. We also thank our anonymous reviewers for their very helpful questions and suggestions.

Funding Statement. Barr's work was funded in part by the John Templeton Foundation (#62385). Opinions expressed here are hers and not those of the JTF; correspondingly, opinions expressed by the JTF are theirs and not hers.

Declarations. None to declare.

References

Anderson, Elizabeth. 1995. 'Knowledge, Human Interests, and Objectivity in Feminist Epistemology'. *Philosophical Topics* 23 (2): 27–58. https://doi.org/10.5840/philtopics199523213.

Brown, Matthew. 2024. 'For Values in Science: Assessing Recent Arguments for the Ideal of Value-Free Science'. *Synthese* 204: 112–142. https://doi.org/10.1007/s11229-024-04762-1.

Douglas, Heather. 2009. *Science, Policy, and the Value-Free Ideal*. University of Pittsburgh *Press*. https://doi.org/10.2307/j.ctt6wrc78.

Douglas, Heather. 2017. 'Why Inductive Risk Requires Values in Science.' In: K. Elliott, and D. Steel. *Current Controversies in Values and Science*. Routledge. https://doi.org/10.4324/9781315639420-6.

Easwaran, Kenny. 2016. 'Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probabilities'. *Noûs* 50 (4): 816–53. https://doi.org/10.1111/nous.12099.

Hájek, Alan. 2003. 'Waging War on Pascal's Wager'. *The Philosophical Review* 112 (1): 27–56. https://doi.org/10.1215/00318108-112-1-27.

Harsanyi, John C. 1985. 'Acceptance of Empirical Statements: A Bayesian Theory without Cognitive Utilities'. *Theory and Decision* 18 (1): 1–30.

Havstad, Joyce, and Matthew Brown. 2017. 'Inductive Risk, Deferred Decisions, and Climate Science Advising.' In: K. Elliott, T. Richards (eds.) *Exploring Inductive Risk*. New York: OUP. https://doi.org/10.1093/acprof:oso/9780190467715.003.0006.

Hempel, Carl G. 1962. 'Deductive-Nomological vs. Statistical Explanation'. *Minnesota Studies in the Philosophy of Science* III: 98–169. https://doi.org/10.22201/iifs.18704905e.1967.33.

Levi, Isaac. 1967. *Gambling with Truth*. Cambridge: MIT Press.

Lugones, María and Elizabeth Spellman. 1983. 'Have We Got a Theory for You! Feminist Theory, Cultural Imperialism, and the Demand for the 'Woman's Voice''. *Women's Studies International Forum* 6 (6): 573–581.

Parker, Wendy. 2014. 'Values and Uncertainties in Climate Prediction, Revisited'. *Studies in History and Philosophy of Science Part A*, Values and Norms in Modeling: The Progress of Science, vol. 46 (Supplement C): 24–30. https://doi.org/10.1016/j.shpsa.2013.11.003.

Parker, Wendy. 2024. 'The Epistemic Projection Approach to Values in Science'. *Philosophy of Science* 91 (1): 1–26. https://doi.org/10.1017/psa.2023.107.

Parker, Wendy, and Eric Winsberg. 2018. 'Values and Evidence: How Models Make a Difference'. *European Journal for Philosophy of Science* 8 (1): 125–42. https://doi.org/10.1007/s13194-017-0180-6.

Peterson, Martin. 2023. 'The St. Peterburg Paradox'. In *The Stanford Encyclopaedia of Philosophy*, edited by Edward N. Zalta and Uri Nodelman.

Steel, Daniel. 2015. 'Acceptance, Values, and Probability'. *Studies in History and Philosophy of Science Part A* 53: 81–88. https://doi.org/10.1016/j.shpsa.2015.05.010.

Ward, Zina. 2021. 'On Value-Laden Science'. *Studies in History and Philosophy of Science Part A* 85 (February): 54–62. https://doi.org/10.1016/j.shpsa.2020.09.006.

Wilholt, Torsten. 2009. 'Bias and Values in Scientific Research'. *Studies in History and Philosophy of Science Part A* 40 (1): 92–101. https://doi.org/10.1016/j.shpsa.2008.12.005.

Worrall, John. 2008. 'Evidence and Ethics in Medicine'. *Perspectives in Biology and Philosophy* 51 (3): 418–431. https://doi.org/10.1353/pbm.0.0040.

Cite this article: Barr, Kelli and Katharina Bernhard. 2026. "Stuck Between a Rock and a Hard Science: Three Challenges with Putting Epistemic Projection to Work." *Philosophy of Science* 93 (4):920–929. https://doi.org/10.1017/psa.2026.10215