An intermediate approach to value management

Wendy S. Parker
Department of Philosophy, Virginia Tech

To appear in *The Routledge Handbook of Values and Science*, edited by Kevin Elliott and Ted Richards

Abstract. The epistemic projection approach (EPA) is an intermediate approach to value management in science. It recognizes that there are sometimes good reasons to make research responsive to contextual values, but it achieves this responsiveness via the careful formulation of a research problem in the problem-selection stage of investigation. EPA is thus an approach that could be acceptable to some parties on both sides of the debate over the value-free ideal. Independent of this, EPA provides practitioners with concrete guidance on how to make research responsive to contextual values. This is illustrated with an example involving air pollution.

1. Introduction

According to the value-free ideal (VFI), choices made in the "research" stage of investigation, as scientists produce and analyze data and draw conclusions, should be kept free from the influence of contextual values as much as possible; ideally, such choices should be based on epistemic considerations alone. The rationale for this prohibition on value influence seems straightforward: contextual values, reflecting what someone wishes or desires to be the case, are irrelevant to choices of methodology that are aimed at finding out what *is* the case, and their influence can bias research in significant ways.

Nevertheless, the dominant view among philosophers today is that the VFI should be rejected. According to the *inductive risk argument*, it can be appropriate, and perhaps even morally required, for uncertain choices faced in the research stage of investigation – such as how to classify ambiguous data or which of several plausible models to employ – to be informed by judgments of how bad the consequences of different errors in conclusions would be (Douglas 2000; 2009). Likewise, according to the *aims argument*, when the aims of scientific investigation are both epistemic and practical (reflecting contextual values), it can be appropriate for methodological choices to be justified by appeal to those joint aims (Elliott and McKaughan 2014; Intemann 2015; see also Brown 2017). These are just some of the arguments leveled against the VFI.

As a consequence, the focus of philosophical discussion has shifted to the question of what constitutes legitimate or illegitimate value influence in research, recently dubbed "the new demarcation problem" (Holman and Wilholt 2022; Resnik and Elliott 2023). A range of positions have been advocated, from permissive ones that see science as a force for social justice, to

restrictive ones that deem appeals to contextual values to be legitimate only when epistemic considerations alone are not decisive.

Another emerging thread of discussion, however, considers the possibility of at least a partial rapprochement among defenders and challengers of the VFI (DiMarco and Khalifa 2019; Lusk and Elliott 2022; Parker 2024). Are there approaches to value management that could be accepted by (some) parties on both sides of the debate? This chapter articulates one proposal, the *epistemic projection approach* (EPA) (Parker 2024). EPA recognizes that sometimes there are good reasons to make scientific research responsive to contextual values, but it achieves this responsiveness via the careful formulation of an epistemic research problem in the problem-selection stage of investigation. A further attractive feature of EPA, independent of its potential for rapprochement, is that it provides practitioners with concrete guidance on value management. Readers primarily interested in this guidance may focus their attention on Section 4.

2. Foundations of EPA

Important insights underlie the inductive risk and aims arguments: scientists *should* take due care to avoid errors that would have particularly bad consequences, and sometimes it *does* make sense for the broader aims of scientific investigation to impact how research proceeds. But in their typical formulations, these arguments are too permissive with regard to value influence.

The aims argument allows that, when research has both epistemic and practical aims, methodological choices may be justified by appeal to those joint aims. This seems right in some cases. For instance, since weather forecasts are produced in order to help people plan their activities, it is appropriate that scientists attend to both accuracy and speed when developing their forecast models; it is no good if highly accurate "forecasts" arrive after the weather has passed. The trouble is that the aims argument doesn't seem to prohibit investigators from attaching much greater importance to practical aims than epistemic ones such that, in some cases, they might be justified in engaging in shoddy science – or even outright fraud – in order to promote those practical aims. This leads Steel (2017) to argue for a "qualified epistemic priority," according to which choices in research may be justified by appeal to joint epistemic and practical aims, but the research must respect basic constraints on epistemically adequate science: investigations must be genuinely probative, data should not be manufactured or cherry picked, conclusions should take account of relevant evidence, and so on.

Qualified epistemic priority is underwritten by the idea that science is first and foremost an epistemic endeavor; if science is to serve some broader practical aims, it should do so by providing epistemic goods, like information and knowledge, that are relevant to the pursuit of those aims, rather than via persuasion, manipulation, etc. (ibid.). Accepting this view of science, we can say even more. Not only should research respect basic constraints on epistemically adequate science, but the default criterion for choosing among routes of investigation should be epistemic quality, i.e., the likelihood that the route will achieve the epistemic goals of the research (if they are achievable at all). This default criterion may be overridden for compelling

ethical, pragmatic, or contextual reasons, such as respecting the autonomy of human research subjects, or managing finite resources, or meeting procedural or legal requirements for research that will inform regulations. But reducing the epistemic quality of research should be viewed as a cost, even when on balance such a reduction is worth it for other reasons. This is one foundational commitment of EPA.

A second foundational commitment concerns inductive risk. According to the inductive risk argument, contextual values sometimes have a role to play in resolving uncertain methodological choices. For example, in a study investigating whether a chemical is carcinogenic, if the foreseeable consequences of a Type II error (concluding that the chemical is not carcinogenic when it is) are judged to be particularly bad, then a researcher might choose to treat ambiguous data, such as tissue samples where there are indicators but not definitive evidence that cancer is present, in a way that minimizes the risk of a Type II error; the researcher might choose to classify those ambiguous cases as cancerous (see Douglas 2000).

But often a better response to uncertain methodological choices like these is to propagate the uncertainty through to the study's conclusions (for variations on this response, see, e.g., Jeffrey 1956; Betz 2017; Menon and Stegenga 2023). Thus, in the example just given, the researcher could calculate how exposure to the chemical affects the risk of cancer, assuming first that ambiguous cases are cancer and then assuming that they are not, arriving at a range estimate rather than a single-valued estimate of the increased risk. Such an approach is better for multiple reasons. It delivers a more accurate picture of researchers' actual epistemic situation vis-à-vis the question of interest and, when findings are intended to inform practical decisions, it allows for those decisions to be made with awareness of extant uncertainty. In addition, when there is value pluralism in society, permitting just one set of values to influence how a research question is addressed can undermine the democratic legitimacy of scientific findings, erode public trust, and lead to an undesirable fragmentation of scientific knowledge (e.g., "conservative science" informed by one set of values and "liberal science" informed by another) (see Betz 2013; Schroeder 2021; Menon and Stegenga 2023).

Consequently, EPA incorporates a default preference for propagating methodological uncertainty through to conclusions, either formally or using expert judgment, rather than choosing one methodological option over the others on value-related inductive risk grounds. Once again, this default preference can sometimes be overridden, since in some cases uncertainty propagation is impossible or undesirable. Perhaps implementing one methodological option destroys or consumes the object of analysis, or perhaps it would take too long or be too expensive to explore the implications of the study's uncertain methodological choices, or perhaps the results are being produced for a narrow set of users whose decision making will not benefit from (and might even be confused by) a careful accounting of uncertainty. It is in situations like these that resolving uncertain methodological choices by appeal to inductive risk considerations (alongside pragmatic ones) becomes more attractive.

In summary, two foundational commitments of EPA are as follows. First, the default criterion for choosing among routes of investigation should be epistemic quality. Second, when there is uncertainty about which methodological choice is best for achieving the epistemic goals of a study, there is a default preference for propagating that uncertainty through to conclusions, rather than choosing one methodological option over the others on value-related inductive risk grounds. These defaults may be overridden for compelling ethical, pragmatic, or contextual reasons. But at a minimum the investigation must respect basic constraints on epistemically adequate science. Otherwise, the activity does not really count as *science*.

3. EPA: Values in problem selection

EPA is primarily an account of *how* to make research responsive to contextual values – whether expressed in the form of broader practical aims, ethical constraints, or in some other way – once it has been determined that at least some responsiveness is appropriate. Challengers of the VFI argue that value considerations sometimes appropriately serve as justifying reasons for choices in the research stage of investigation. EPA, by contrast, limits consideration of values to the problem selection stage. More specifically, EPA proposes that, when there are compelling reasons for research to be responsive to contextual values, this responsiveness should occur by *formulating an epistemic research problem that aligns with those values*. The rest of this section unpacks this proposal.

3.1 Problem selection

The problem selection stage of investigation has received comparatively little attention from philosophers interested in values in science. Often the selection of a research problem is simply equated with the choice of a question to be answered. Yet research problems clearly can, and in practice often do, have a richer structure. For instance, an investigation may be undertaken to address several research questions, some of which are of higher priority for answering than others. Or it might be clear from the outset that only information produced before a particular practical decision must be made will be useful, so that the research problem to be solved is not merely to answer a given question but to do so within a specified time frame.

EPA understands research problems in a way that leaves room for this richer structure and allows for making it explicit. In particular, it understands research problems to include not only epistemic goals (corresponding to research questions), but also, optionally, constraints and preferences regarding how those goals are pursued. *Epistemic goals* specify particular epistemic goods to be obtained in research, such as particular pieces of information or knowledge. For example, an ecologist might have the goal of learning the main ecological roles of three salient species in a novel ecosystem. *Constraints* specify features that research undertaken in pursuit of those goals is required to have. For example, the ecologist might have permission to study the novel ecosystem only insofar as the research does not significantly damage the ecosystem's functioning. *Preferences* indicate features that are desirable, but not required, for research to have. It might be preferred, for instance, that research on the novel ecosystem be conducted

using the instruments and other resources that the ecologist's research team already possesses, rather than requiring the acquisition of new ones.

A research problem is *epistemic* in character if its constraints and preferences (if any) reference only features of research qua epistemic activity—its methods and data, the types of errors to which it is more or less prone, and so on; the research problem's specification does not reference any broader aims or contextual values. The example of a constraint given above, that research on the novel ecosystem *must not significantly damage the ecosystem's functioning*, fails to meet this criterion. However, a constraint that simply *excludes the use of method X* would count as epistemic in character, even if the underlying reason for excluding the method was that it carried a substantial risk of damaging the ecosystem's functioning. (In case the reader doubts that it is appropriate to consider this an "epistemic" constraint, consider by analogy that a preference for vegan food is a "dietary" preference, regardless of the motivations for it.)

A research problem *aligns with* a broader aim (or a contextual value expressed in some other form) to the extent that research that achieves the problem's epistemic goals, while meeting its constraints and preferences, can be expected to produce results that are relevant for, useful to, or otherwise helpful with respect to, pursuit of that broader aim (or the promotion of the contextual value in some other form). Suppose, for example, that the broader aim motivating an investigation is to make infrastructure decisions that will protect a city from deadly flooding from an adjacent river, given concerns that climate change might bring about such flooding. A research problem that aligns with this broader aim might include the epistemic goal of learning the largest plausible changes in the frequency and magnitude of flooding over the next several decades, as well as a preference for risking overestimating these changes, rather than underestimating them, if some inductive risk must be taken on. A research problem including the opposite inductive risk preference would not align with the broader aim, since underestimating the flooding is more likely to lead to infrastructure decisions that inadequately protect the city.

The process of reasoning from a broader aim (or contextual value expressed in some other form) to an epistemic research problem that aligns with that aim/value is at the heart of EPA and is called *epistemic projection* (Parker 2024). It involves identifying (i) one or more plausibly achievable epistemic goals that are relevant to pursuit of the broader aim (or the promotion of the contextual value), and, optionally, (ii) epistemic constraints and/or preferences regarding how those goals are pursued that will help to align the research with the broader aim (or the contextual value expressed in some other way). The terminology of "projection" here is inspired by the process of geometrical projection, which shows what an object "looks like" when collapsed to a reduced set of dimensions, e.g., a 3-dimensional object projected onto a 2-dimensional plane; the process of epistemic projection reveals what pursuit of a broader aim (or protection/promotion of a contextual value expressed in some other way) can "look like" when limited to choices regarding epistemic goals, constraints, and preferences that jointly constitute a research problem.

To make (ii) manageable, EPA calls for consideration of just four types of epistemic constraint/preference, which typically present the most significant opportunities for aligning research with contextual values:

a) *epistemic priorities*: when research has more than one epistemic goal, these constraints and preferences indicate which epistemic goals take priority if trade-offs must be made;
b) *epistemic resources*: these constraints and preferences identify concepts, models, methods, instruments, etc., that should or should not be employed in research;
c) *inductive risk*: these constraints and preferences indicate the standards of evidence that should be used when drawing conclusions and which type(s) of error in conclusions should be preferentially avoided and in what circumstances;
d) *form-of-conclusion*: these constraints and preferences indicate a desired (or undesired) mode of formulating and communicating research findings.

Of course, not every research problem identified via epistemic projection will include all four of these. Sometimes, an investigation will have only one epistemic goal, making (a) irrelevant. Other times, only an inductive risk preference will be important. And so on.

Note, however, that even when there are compelling reasons to try to align research with a broader aim or contextual value, the drawbacks of including particular constraints or preferences in a research problem should be considered: Will meeting the constraint or preference significantly reduce the epistemic quality of the research? Will it be very costly in terms of time, funding, or other resources? Even if including the constraint or preference would help to align research with a contextual value, the foreseeable costs – both epistemic and practical – might be significant enough that, on balance, a decision is made not to include it.

Any constraints or preferences that are included must leave room for research that, at a minimum, respects basic constraints on epistemically adequate science; this is necessary for the research problem to count as *scientifically acceptable*. Thus, for example, a research problem cannot include an epistemic resource constraint – say, requiring the use of a particular methodology – such that research that satisfies that constraint will be virtually guaranteed to produce a particular result, regardless of the truth. Ideally, the research problem will leave room for research that not only respects basic constraints on epistemically adequate science but is of excellent epistemic quality.

3.2 The research stage

A research problem identified via epistemic projection then constitutes a brief for the research stage of investigation: researchers should try to achieve the specified epistemic goals, while satisfying the specified epistemic constraints and preferences (see Carrier 2022 for a similar view in the context of science-based policy advice).

Importantly, the process of epistemic projection will have *screened off* the broader aims / contextual values from this research stage: the research problem will be specified in a way that makes no reference to those aims / values, and choices in the research stage will be justified as a

means of solving that epistemic research problem, not by appeal to any broader aims / contextual values. (For a concrete illustration, see the research problem specification in Table 1, in the next section.) Indeed, in principle the research stage of investigation could be undertaken by researchers who are unaware of the aims / values that played a role in the formulation of the research problem.

If the research findings are shared publicly, then researchers should also report which methodological choices in the research stage were made in response to which preferences and constraints specified in the research problem. This transparency requirement is intended to facilitate healthy scrutiny of research practices. EPA does not, however, require that researchers report the broader aims or values considered during epistemic projection; these broader aims and values are irrelevant to the assessment of the quality of research qua epistemic activity. Of course, reporting such aims or values, as well as the reasoning that was employed in the process of epistemic projection, is not prohibited either, and researchers might often choose to report this information as well.

4. Putting EPA into practice

Having introduced the foundations and basic elements of EPA, this section discusses more concretely how EPA can be put into practice, illustrating with an example involving air pollution. There are three main steps.

Step 1. *Determine whether it is appropriate to try to make the research at hand responsive to contextual values*.

As mentioned earlier, EPA is primarily an account of how to make research responsive to contextual values, once it has been determined that some responsiveness is appropriate. It is beyond the scope of this discussion, and probably a fool's errand, to try to provide an exhaustive list of circumstances in which responsiveness to contextual values is appropriate. But a partial list of representative situations is easy to provide.

First, if research is undertaken to *inform a high-stakes practical decision* where one type of error in conclusions can be expected to have serious negative consequences for the public, then typically it will be appropriate to include a preference for avoiding that type of error (when methodological uncertainty cannot be propagated through to conclusions). Second, if research is being performed *in a consulting mode*, then often it will be appropriate to make research responsive to the client's values, e.g., as reflected in their epistemic priorities and inductive risk preferences. A third type of situation is when a regulatory agency conducts scientific research in support of its activities but also has *a value-related mandate*, e.g., to protect human health. Yet another situation is when there is *pervasive value-related bias* in an existing body of research; in some cases, this could warrant making research responsive to an alternative set of values, whether as a corrective epistemic measure or to combat harms ensuing from the existing bias (e.g., if a body of research is sexist or racist).

Example. Suppose a county government is considering granting a permit to a company to build a chemical manufacturing facility in the county. Residents in a city near the proposed site are concerned that the air pollutants that this type of facility is known to emit will sometimes reach harmful levels in the city. The company has produced an impact analysis in a standard way, which finds that pollutants throughout the county will remain within acceptable levels. But the analysis is at a fairly coarse spatial scale, and the residents know that the orography of the region is such that some neighborhoods in the city – smaller in scale than could be resolved by the company's analysis – sometimes experience much worse air quality than others. The city council contracts a group of scientists at a nearby university to probe in a more fine-grained way what the pollution might be like in particular neighborhoods; their overriding concern is with protecting community health.

In this example, the research is being done in a consulting mode, and the findings will inform a relatively high-stakes decision where one type of error could have serious negative consequences for the public (namely, exposure to chemicals that harm their health). This is a situation in which there is good reason to try to conduct research that aligns with the community's broader aim, especially if it can be done without sacrificing much epistemic quality.

Step 2. *Perform epistemic projection. The research problem that is specified must be scientifically acceptable and, ideally, will leave room for research of excellent epistemic quality.*

EPA does not require that anyone in particular participates in the process of epistemic projection. However, when research is being done to inform the decisions and actions of others, as in the air pollution example, then often it will be best for epistemic projection to be performed collaboratively by both the researchers and those actors. The researchers will have a better sense of which epistemic goals are feasible and which methodologies are available, but the actors will have a better understanding of the aims and values to which the investigation is meant to be responsive.

A helpful way to approach epistemic projection in practice is by answering a series of questions that correspond to the different elements that an epistemic research problem can have. They are articulated here for the case where there is some broader aim (reflecting contextual values) motivating the investigation, as in the air pollution example:

1. Which plausibly obtainable epistemic goods are particularly relevant to pursuit of the broader aim? (*goals*)
2. If multiple epistemic goods will be pursued, are some of them more important than others to achieving the broader aim? If so, which one(s)? (*priorities*)
3. Would the use of some epistemic resources (concepts, data, models, methods, etc.) rather than others be advantageous to the pursuit of the broader aim? If so, which ones? (*resources*)
4. Would some errors in findings, more than others, impede pursuit of the broader aim? If so, which ones? (*inductive risk*)

5. Would it facilitate pursuit of the broader aim to express or represent findings in a particular way? If so, which way? (*form of conclusion*)

Answers to these questions can guide the specification of one or more epistemic goals, as well as preferences and/or constraints, which jointly constitute an epistemic research problem that aligns with the broader aim.

A complication is that, since preferences come in degrees, some indication of their strength also should be provided. The best way to do this will vary with the context, but manageable approaches are readily identified. For instance, preferences might be classified as "weak", "moderate", or "strong," where some qualitative description of what this means is provided. Or specific limits might be set, such as, prioritize goal A over goal B, until goal B is compromised to this extent. These sorts of approaches will not perfectly capture preferences, but EPA accepts some imperfection for the sake of tractability.

As noted earlier, when deciding whether to include a particular preference or constraint in a research problem, not only the benefits (in terms of aligning research with the broader aim), but also the costs – both epistemic and practical – should be considered. At a minimum, the research problem needs to leave room for research that respects basic constraints on epistemically adequate science. In cases where this condition is met but researchers judge that satisfying the constraints of the research problem would still compromise epistemic quality more than they are comfortable with, they should feel free to decline to perform the research. (Likewise, researchers should be able to decline to address a research problem if they judge it to have ethically objectionable motives or, more generally, if they judge that there are not actually compelling reasons to make the research responsive to contextual values in the ways reflected in a research problem.)

Example. In the air pollution example, we can imagine that the scientists and city council members together perform epistemic projection with the help of questions (1)-(5). Table 1 shows an epistemic research problem that might be specified as a result of this process, along with an indication of the rationale for each component of the problem. Notice that the research problem's specification makes no mention of the broader aim of protecting community health; it refers only to features of research qua epistemic activity. Moreover, it seems to be scientifically acceptable. In fact, carrying out research that meets the constraints and preferences in Table 1 might not require sacrificing any epistemic quality at all (relative to alternative routes of investigation that might be selected in pursuit of the epistemic goal).

Research Problem Component	Research Problem Specification	Rationale
Epistemic Goal(s)	Learn what the statistical distribution of pollution levels for chemicals C_1 and C_2 would be at locations around town during a typical year	Learning this is highly relevant to the aim of protecting citizens' health
Epistemic Priorities	*Preference*: Prioritize accuracy of the upper quartiles of the distributions, and for neighborhoods {N} *Strength*: moderate; the quality of estimates of the median levels and for other neighborhoods shouldn't be greatly reduced	High-pollution days are of greatest risk to health, and neighborhoods {N} are suspected to be at greater risk of high-pollution days; but other neighborhoods deserve quality information too, and median levels are also of interest
Inductive Risk	*Constraint*: Risk overestimating pollution levels if some inductive risk must be taken on, but avoid taking on inductive risk whenever feasible, e.g., by propagating uncertainty	Overestimating pollution is better than underestimating it, given the goal of protecting health, but the council wants the results to be easily defensible if challenged by the company
Epistemic Resource	*Constraint*: Use computer simulation model that can resolve pollution at neighborhood scales	That some neighborhoods may experience higher pollution than others is of particular concern
Form-of-Conclusion	*Constraint*: Present findings as maps of median and 98th percentile pollution levels, and provide some indication of the associated uncertainty	Maps of middling and high levels are easier for community members to understand than statistical distributions; uncertainty information will help defend against challenges from the company

Table 1. Example of an epistemic research problem formulated via epistemic projection.

Step 3. *Try to solve the research problem.*

The third step occurs in – and indeed constitutes – the research stage of investigation. Researchers attempt to solve the epistemic research problem that was specified in Step 2. Solving an epistemic research problem requires, at a minimum, achieving its epistemic goals while satisfying its epistemic constraints; the more epistemic preferences are also satisfied, the better.

Example. In the air pollution study, we can imagine the university researchers making a number of choices when addressing the research problem. We can expect that they will employ a pollution model with a spatial resolution fine enough to resolve neighborhood-level variation. They might invest some time testing and improving the model's representation of a physical process that is particularly relevant for capturing the movement of pollution in the county's orography, to try to ensure that pollution levels, especially levels in neighborhoods {N}, are accurately estimated. We can expect that they will present their findings in the map format specified by the research problem. And so on.

Notice, however, that these choices by the university researchers will be justified as a means of addressing an epistemic research problem, not by appeal to any broader aims of the investigation (which, as noted earlier, in principle needn't even be known by individuals conducting research). At the same time, quite different choices might be made if, in the problem selection stage, the broader aim of protecting community health had not been considered, and researchers were simply asked to investigate the potential pollution changes stemming from the proposed facility.

In that case, the researchers might have chosen a standard lower-resolution model like the one the chemical company used, made no adjustments to the model, and presented results differently, e.g., as average pollution levels inside and outside the city limits.

5. Strengths and Limitations of EPA

As an approach to value management, EPA has some attractive features.

First, there is EPA's potential for rapprochement. On the one hand, EPA recognizes that sometimes there are good reasons to make research responsive to contextual values. On the other hand, EPA permits contextual values to play a justificatory role only in the problem selection stage of investigation, and it calls for researchers to use epistemic quality as the default criterion when choosing among routes of investigation. Consequently, EPA could appeal *both* to those who reject the VFI but hold that epistemic considerations should still enjoy a qualified priority *and* to defenders of the VFI who are primarily concerned that value influence will result in shoddy science. (It will not, however, be acceptable to defenders of the VFI who hold that contextual values should *make no difference* to how research is pursued, nor to challengers of the VFI who reject even a qualified epistemic priority.)

Second, even if no such rapprochement is achieved, EPA is an approach to value management that is relatively easy to implement and may be more comfortable for practitioners than many alternatives. For practitioners who are eager to make some of their research responsive to contextual values but are unsure *how* to do it, EPA provides concrete guidance: they can consider a limited set of dimensions of research (goals, priorities, resources, inductive risk, form of conclusions) in the problem selection stage of investigation, and then proceed with research in the usual way. And for practitioners who only reluctantly accept that some of their research should be responsive to contextual values, EPA may be more comfortable than approaches to value management that place less emphasis on epistemic quality.

At the same time, EPA also has some limitations.

First, EPA is a pragmatic approach to value management, not an ideal one; it sacrifices perfection for the sake of tractability. It is plausible that researchers would be less likely to miss opportunities for making research responsive to contextual values if those values were kept in view throughout the research stage and considered each time a methodological choice was made. But this would be much more onerous for practitioners. Likewise, EPA accepts that the strength of preferences regarding priorities, inductive risk, etc., identified in the process of epistemic projection typically will only approximately reflect the extent to which the anticipated users of the findings hold particular contextual values.

A second limitation of EPA is that it is incomplete in a significant way. EPA doesn't provide a general account of when it is appropriate for research to be responsive to contextual values. There seem to be some easy cases, as in the air pollution example, where research was undertaken in a consulting mode and could be significantly aligned with an ethically

unobjectionable broader aim without sacrificing much, if any, epistemic quality. But matters can quickly become controversial. What if the city council only cared about pollution in the city's wealthy neighborhoods? Is it appropriate for pharmaceutical companies to implement EPA, with profitmaking as one of the contextual values to which their research is responsive? If so, how much epistemic quality may they legitimately sacrifice for the sake of profitmaking? EPA in its current form does not directly address difficult questions like these; it does not provide a complete answer to the new demarcation problem.

That said, as noted earlier, researchers employing EPA may always decline to address a research problem that has ethically objectionable motivations or that includes constraints that can only be met by conducting research that fails to meet the researcher's own standards for epistemic quality. Moreover, when research does proceed, the transparency requirement of EPA, if faithfully implemented, will facilitate its critical evaluation.

6. Conclusion

EPA is an intermediate approach to value management. It recognizes that there are sometimes good reasons to make research responsive to contextual values, but it achieves this responsiveness via the careful formulation of a research problem in the problem-selection stage of investigation. EPA is thus an approach that could be acceptable to some parties on both sides of the debate over the VFI. And even if EPA fails to effect any such rapprochement, it has other strengths. It is an approach to value management that provides concrete, tractable guidance for practitioners, and it achieves significant value responsiveness while maintaining a firm commitment to the epistemic quality of research.

References

Betz, Gregor. "In defence of the value free ideal." *European Journal for Philosophy of Science* 3, (2013): 207–220. https://doi.org/10.1007/s13194-012-0062-x

Betz, Gregor. "Why the Argument from Inductive Risk Doesn't Justify Incorporating Non-Epistemic Values in Scientific Reasoning." In *Current Controversies in Values in Science*, edited by K. C. Elliott and D. Steel. Routledge Press, 2017. https://doi.org/10.4324/9781315639420-6.

Brown, Matthew. "Values in Science: Against Epistemic Priority." In *Current Controversies in Values in Science*, edited by K. C. Elliott and D. Steel. Routledge Press, 2017.

Carrier, Martin. "What Does Good Science–Based Advice to Politics Look Like?" *Journal for General Philosophy of Science* 53 (2022): 5–21. https://doi.org/10.1007/s10838-021-09574-2.

DiMarco, Marina, and Kareem Khalifa. "Inquiry Tickets: Values, Pursuit, and Underdetermination." *Philosophy of Science* 86, no. 5 (2019): 1016-1028.

Douglas, Heather E. “Inductive Risk and Values in Science.” *Philosophy of Science* 67, no. 4 (2000): 559–79. https://doi.org/10.1086/392855.

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

Elliott, Kevin C., and Daniel J. McKaughan. “Nonepistemic Values and the Multiple Goals of Science.” *Philosophy of Science* 81, no. 1 (2014): 1–21. https://doi.org/10.1086/674345.

Holman, Bennett, and Torsten Wilholt. “The New Demarcation Problem.” *Studies in History and Philosophy of Science* 91, (2022): 211-220.

Intemann, Kristen. “Distinguishing between Legitimate and Illegitimate Values in Climate Modeling.” *European Journal for Philosophy of Science* 5, (2015): 217–32. https://doi.org/10.1007/s13194-014-0105-6.

Jeffrey, Richard C. “Valuation and acceptance of scientific hypotheses.” *Philosophy of Science* 22, (1956): 237–246. https://doi.org/10.1086/287489.

Lusk, Greg, and Kevin C. Elliott. “Non-epistemic Values and Scientific Assessment: An Adequacy-for-Purpose View.” *European Journal for Philosophy of Science* 12, (2022): 35. https://doi.org/10.1007/s13194-022-00458-w.

Menon, Tarun, and Jacob Stegenga. “Sisyphean science: why value freedom is worth pursuing.” *European Journal for Philosophy of Science* 13, (2023): 48. https://doi.org/10.1007/s13194-023-00552-7.

Parker, Wendy S. “The Epistemic Projection Approach to Values in Science.” *Philosophy of Science* 91, no. 1 (2024): 18-36. https://doi.org/10.1017/psa.2023.107.

Resnik, David B., and Kevin C. Elliott. “Science, Values, and the New Demarcation Problem.” *Journal for General Philosophy of Science* 54, (2023): 259–286. https://doi.org/10.1007/s10838-022-09633-2.

Schroeder, S. Andrew. “Democratic Values: A Better Foundation for Public Trust in Science.” *The British Journal for the Philosophy of Science* 72, (2021): 545-562. https://doi.org/10.1093/bjps/axz023.

Steel, Daniel. “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. Routledge Press, 2017. https://doi.org/10.4324/9781315639420-4.