



# Non-epistemic values and scientific assessment: an adequacy-for-purpose view

Greg Lusk<sup>1,2</sup>  · Kevin C. Elliott<sup>3</sup>

Received: 3 September 2021 / Accepted: 8 March 2022 / Published online: 1 June 2022  
© The Author(s) 2022

## Abstract

The literature on values in science struggles with questions about how to describe and manage the role of values in scientific research. We argue that progress can be made by shifting this literature's current emphasis. Rather than arguing about how non-epistemic values can or should figure into scientific assessment, we suggest analyzing how scientific assessment can accommodate non-epistemic values. For scientific assessment to do so, it arguably needs to incorporate goals that have been traditionally characterized as non-epistemic. Building on this insight, we show how the adequacy-for-purpose framework recently developed for assessing scientific models can provide a general framework for describing scientific assessment so that it goes beyond purely epistemic considerations. Adopting this framework has significant advantages and opens the possibility of effecting a partial rapprochement between critics and proponents of the value-free ideal.

**Keywords** Science and values · Hypothesis testing · Assessment · Modeling · Adequacy-for-purpose

## 1 Introduction

The contemporary literature on values in science has been deeply influenced by a pair of essays written by Thomas Kuhn and Ernan McMullin. In his famous piece,

---

✉ Greg Lusk  
gregory.lusk@durham.ac.uk

Kevin C. Elliott  
kce@msu.edu

<sup>1</sup> Lyman Briggs College and the Department of Philosophy, Michigan State University, East Lansing, MI, USA

<sup>2</sup> The Department of Philosophy, Durham University, Durham, UK

<sup>3</sup> Lyman Briggs College, the Department of Fisheries and Wildlife, and the Department of Philosophy, Michigan State University, East Lansing, MI, USA

“Objectivity, Value Judgment, and Theory Choice” (1977), Kuhn conceptualized theory assessment as a matter of weighing multiple “values,” such as accuracy, scope, and simplicity. Thus, he emphasized that science is guided by value judgments both in determining when these values are met and in weighing their relative importance. McMullin (1982) built on Kuhn’s work by distinguishing different sorts of values that could influence theory assessment. He focused on “epistemic values,” which are “presumed to promote the truth-like character of science” (1982, 19). He acknowledged that it could be difficult to delineate these epistemic values from other considerations that influence scientists. For example, scientists often consider pragmatic factors, such as limitations of time or resources, as well as a wide variety of political, ethical, and social considerations. McMullin promoted a distinction between epistemic and non-epistemic values. He frankly admitted, “The list [of potential non-epistemic values] is as long as the list of possible human goals. I shall lump these values together under the single blanket term ‘non-epistemic’” (McMullin, 1982, 20).

In contrast to the emphasis placed on epistemic values by Kuhn and McMullin, much of the recent literature on science and values has explored how non-epistemic values can play legitimate roles in scientific reasoning despite purportedly being irrelevant to advancing the “truth-like character of science.” Some authors have focused on the underdetermination of theory by evidence, arguing that non-epistemic values can help to break “ties” between theories that are equally well supported from an epistemic point of view (see e.g., Biddle, 2013; Longino, 1990). Others have appealed to inductive risk, arguing that scientists need to consider non-epistemic values when deciding how to handle the potential for error when accepting or rejecting hypotheses (see e.g., Douglas, 2009; Elliott & Richards, 2017). Nevertheless, although this literature has proliferated widely, it has been plagued by nagging questions. Some authors continue to question whether non-epistemic values can justifiably influence the assessment of scientific theories or hypotheses (e.g., Betz, 2013; Hudson, 2016). Even those who accept roles for non-epistemic values in scientific assessment struggle to characterize their position. For example, scholars have found it difficult to categorize the multitude of factors that can influence scientific reasoning (Biddle, 2013), specify how these factors relate to each other (e.g., Douglas, 2009; Brown, 2013; Elliott & McKaughan, 2014), or whether the epistemic/non-epistemic distinction is even tenable (see e.g., Longino 1996; Rooney, 2017).

In this paper, we argue that progress can be made in addressing these challenges by developing an account of scientific assessment that can accommodate non-epistemic factors. This moves us beyond previous work in two important ways. First, we argue that if one rejects the value-free ideal, then one *must* develop a view of scientific assessment that incorporates additional goals beyond the pursuit of truth. This is significant because it establishes that practical conceptions of assessment are not just a unique feature of some work on values in science but rather an essential (albeit not always explicit) presupposition of those who challenge the value-free ideal. Second, we provide a generalized framework for engaging in scientific assessment in a manner that is not purely epistemic. Although we are not the first to call for the consideration of non-epistemic factors when assessing the products of science (see e.g., Elliott & McKaughan, 2014; Intemann, 2015; Brown, 2020), we have not seen any detailed, systematic frameworks for providing such an assessment in this

previous literature. Instead, there is typically a somewhat generic claim that value influences are acceptable to the extent that they accord with the goals of inquiry in a particular context. In the final section of the paper, we show how the development of a more systematic framework for scientific assessment from a non-epistemic perspective could strengthen the existing scholarship on values in science.

## 2 Social values and the goal of scientific assessment

In this section we advance a meta-philosophical claim: arguments that reject the value-free ideal, and thus permit non-epistemic values a role in scientific assessment, are committed to rejecting the notion that assessment focuses solely on truth. By “scientific assessment” we mean the appraisal of the products of science to determine their status or acceptability. Scientific assessment therefore includes traditional approaches to hypothesis testing and evaluation, as well as similar forms of analysis applied to other scientific products, for example, evaluations of fit or similarity between a model and a target. We first construct a formal argument to demonstrate that arguments against the value-free ideal that employ a traditional distinction between epistemic and non-epistemic values must presume that assessment does not focus solely on truth. We thus show that, in fact, those who reject the value-free ideal also reject the sole focus on truth in scientific assessment. We assert that the generality of this point, and its implications, have not yet been widely realized.

In order to construct our formal argument, we need to start with the distinction that defenders of the value-free ideal – that is, those who would seek to prohibit non-truth conducive considerations from entering the internal or core aspects of science – typically make between epistemic values and non-epistemic values (see e.g., McMullin, 1982):

Definition: Epistemic values are indicators of truth, whereas non-epistemic values are indicators of something else, and only something else.<sup>1</sup>

Given this distinction, one can formulate a crucial but typically implicit presupposition that drives arguments for and against the value-free ideal:

(1) If (A) the goal of scientific assessment is to evaluate truth only, then (C) epistemic values are the only values relevant in scientific assessment.

---

<sup>1</sup> For the purposes of this paper, we are defining epistemic values in a narrow way, such that they are indicators of truth. Similarly, we are defining non-epistemic values such that they are never indicators of truth, even though it is plausible that ethical and social values (which are typically regarded as non-epistemic) can sometimes be helpful for arriving at truths. There is an extensive literature that clarifies different ways of conceptualizing and characterizing epistemic and non-epistemic values (see e.g., Douglas, 2009, 2013; Lacey, 2017; Rooney, 2017; Steel, 2010). Our point in this paper is not to debate the nature of epistemic values but only to use that language to help provide an explicit argument that if scientific assessment is focused only on truth, then only indicators of truth are relevant to scientific assessment.

We have labeled the antecedent and the consequent of this premise as (A) and (C) for referential ease. This premise appears to be uncontroversial.<sup>2</sup> If one's only goal for assessment is to evaluate truth, then one would not appeal to indicators of properties other than truth. One can flesh out these concepts (e.g., what "truth" means or what qualifies as an "indicator") in different ways, but we take it that the plausibility of the premise remains unchanged. For example, one could say that epistemic values are indicators of truth insofar as they play an evidential role. This evidential role could be described in probabilistic terms: an epistemic value,  $E$ , raises or lowers the probability of a scientific hypothesis, such that  $P(H|E) > P(H)$  or  $P(H|E) < P(H)$ . Empirical adequacy, accuracy, and so on are examples of widely acknowledged epistemic values that plausibly operate as "evidence" in this way. If non-epistemic values did not raise or lower the probability of a hypotheses (or other object being assessed), and if the only job of assessment were to evaluate the probability of truth, then only epistemic values would be relevant to assessment.

In our view, proponents and opponents of the value-free ideal respond to this crucial premise differently. We contend that proponents of the value-free ideal almost invariably (though often implicitly) affirm the antecedent of (1):

(A) The goal of scientific assessment is to evaluate truth only.

Based on premises (1) and (A), by modus ponens, they draw the following conclusion:

(VFI) Epistemic values are the only values relevant in scientific assessment.

Opponents of the value-free ideal have resisted this conclusion. They hold the following:

(~VFI) It is not the case that epistemic values are the only values relevant in scientific assessment.

This conflict between (VFI) and (~VFI) has been the focus of previous debates regarding values in science. We contend that these debates can be clarified by shifting attention to the role of premise (1) in the background of the debate. One can see that (~VFI) is the denial of (C) of premise (1). Thus, unless they deny premise (1) altogether, opponents of the value-free ideal are forced by modus tollens to deny the antecedent of (1):

<sup>2</sup> Although the premise initially appears to be uncontroversial, one might question whether it applies to the adequacy-for-purpose (AFP) hypotheses that we discuss later in this paper. Because those hypotheses make reference to purposes that can be non-epistemic in character, non-epistemic values can play a role in setting the truth conditions for those hypotheses. For example, a model might be "adequate for purpose" if it displays the non-epistemic characteristic of requiring little computing power to use. Thus, in contrast to the case of most hypotheses that have been considered in the literature on values in science, non-epistemic values actually serve as indicators of truth and can play the role of evidence in assessing these AFP hypotheses. One might argue, however, that if these non-epistemic values are serving as indicators of truth, then they actually function as epistemic values in the context of AFP hypotheses. Depending on whether one classifies them as epistemic or not, this might pose an exception to premise (1). We believe this might be an interesting and fruitful consequence of our view, which we discuss in the conclusion.

( $\sim$ A) It is not the case that the goal of scientific assessment is to evaluate truth only.

This conclusion highlights a new narrative for describing debates about the value-free ideal: those who reject it not only argue for the permissibility of non-epistemic values in scientific assessment, but they also end up disputing the nature of scientific assessment itself. That is, they reject (VFI) because they disagree with the antecedent of premise (1); they hold that the goal of assessment is or should be broader than the examination of truth.

In light of the above argument, disputing the truth of (A) – disputing that scientific assessment examines truth alone – is a necessary step in arguing for the relevance of non-epistemic values to scientific assessment. Consider what would have to be the case if one affirmed both the relevance of non-epistemic values and asserted that the goal of hypothesis assessment was to examine only truth. In such a case, non-epistemic values would need to play the role of evidence or help interpret how evidence would support the hypothesis; put in terms of the probabilistic account of evidence provided earlier, non-epistemic values would raise or lower the probability of the hypothesis. However, those who accept a distinction between epistemic and non-epistemic values (whether they accept or reject the VFI) presumably agree that this is incoherent. Allowing the values which have been traditionally identified as non-epistemic to play this role essentially amounts to wishful thinking: a hypothesis is true because it helps one achieve one's non-epistemic goals. Often, proponents and opponents of the VFI indicate that the broad implications, social and scientific, of permitting this in scientific practice would be grave. In order to reject VFI then, one requires a notion of scientific assessment that can go beyond truth.

Although this analysis holds promise, one might raise at least two worries about it. The first is that our fundamental premise focuses on the concept of truth, and truth is not the only, or even the most appropriate, epistemic goal of scientific assessment. For example, one might think that empirical adequacy or understanding is a more appropriate epistemic goal for science (De Regt, 2017; Potochnik, 2017; Van Fraassen, 1980).<sup>3</sup> We do not think this is a particularly significant difficulty for our analysis. If one were a proponent of an alternative epistemic goal, one could just alter both our definition of epistemic values and our first premise accordingly. For example, one might formulate them as follows:

Definition': Epistemic values are indicators of understanding, while non-epistemic values are indicators of something else, and only something else.

(1') If (A') the goal of scientific assessment is to evaluate understanding only, then (C') only epistemic values are relevant in assessment.

While this would change the details of our analysis, the main point would remain: if one allows non-epistemic considerations to play a role in scientific assessment, one

<sup>3</sup> The extent to which understanding is a purely epistemic goal is itself subject to potential disagreement (see e.g., Khalifa, 2020; Potochnik, 2020). For the purposes of our example here, we assume that understanding is conceptually independent of non-epistemic values.

must acknowledge that scientific assessment involves more than solely the pursuit of epistemic goals.

The second worry is that (1) is formulated in terms of a distinction between epistemic and non-epistemic values, and this distinction is sometimes rejected (see e.g., Longino, 1990; Rooney, 2017). However, we contend that (1) should still be uncontroversial, even to those who reject this distinction. The most natural way for an opponent of the distinction to respond to premise (1) would be to deny both its consequent and its antecedent. If epistemic values cannot be distinguished from non-epistemic values, then it makes no sense to insist that only epistemic values should play a role in assessment. Thus, opponents of the distinction will presumably accept premise (1) but regard it as a somewhat quaint description of an imaginary situation. If scientific assessment were actually about truth alone, then it would make sense that only truth-oriented (i.e., epistemic) values would be relevant. But for those who deny that we can distinguish epistemic from non-epistemic values, it makes no sense to say that scientific assessment in the real world is solely about truth.

Admittedly, some opponents of the distinction between epistemic and non-epistemic values might claim that we have ascribed too radical a position to them. They could insist that in practice one can identify considerations that are likely to count for or against the truth of a scientific claim; they would just insist that one cannot generalize these considerations and universally label them as “epistemic” or “non-epistemic”. For example, they might claim that explanatory power would count as an indicator of truth in one context but not in another. Our response is that as long as one accepts a distinction between epistemic and non-epistemic considerations in specific instances of assessment, our argument still applies. One would simply need to specify that when one refers to “epistemic” and “non-epistemic” values, one is referring to factors that count as epistemic and non-epistemic in a particular case. Premise (1) would still apply in such cases: if scientific assessment were solely about truth, then only the factors regarded as epistemic in a particular case would be relevant to assessment in that case.

Thus far, we have provided a relatively abstract argument that premise (1) is lurking in the background of debates over the value-free ideal. To buttress our argument, here we show that prominent proponents and opponents of the value-free ideal in fact appear (implicitly or explicitly) to make the argumentative moves that we described above. In other words, proponents of the value-free ideal continue to maintain that scientific assessment is solely about truth, whereas opponents of the value-free ideal deny that supposition. For example, in his defense of the value-free ideal, Robert Hudson insists that “where the moral, social, or political implications of the decision [to accept a hypothesis] are irrelevant, or are simply ignored, the ancillary assumptions [involved in determining whether it is to be accepted] will be of a logical, empirical, or theoretical nature that compels a conclusion because of their assumed truth” (2016, 171). He then endeavors to show that scientific assessment can indeed be performed in a manner that focuses only on truth rather than on “moral, social, or political implications.” Betz (2013, 2017) similarly emphasizes that scientists face assessment choices that are arbitrary from an epistemic point of view, insofar as epistemic factors do not determine which choice is better (Betz 2017, 102). However, for the sake of protecting democratic institutions and personal autonomy, he insists that scientific assessment can and should

be performed in a manner that focuses solely on truth (see Betz, 2013, 207; cf. Lusk, 2020, 2021).<sup>4</sup>

Our position also receives support from those who clearly reject the value-free ideal. Consider, for example, three prominent views of this type: the aims approach, Helen Longino's contextual empiricism, and the argument from inductive risk. The aims approach to values in science justifies roles for non-epistemic values in science based on the ways they help to achieve the non-epistemic aims of research. Thus, it is relatively clear that proponents of this approach deny that scientific assessment is directed solely at truth. For example, according to Kristen Intemann, "The aims approach maintains that social, ethical, and political value judgments are *legitimate* ... insofar as they promote democratically endorsed epistemological and social aims of the research.... (2015, 219, italics in original). In other words, Intemann claims that nonepistemic (social, ethical, and political) values are relevant to scientific assessment because the assessment process incorporates policy-related interests and not solely epistemic ones. Elliott and McKaughan (2014) make the same point, and they apply it to scientific representations in general: "Given that scientific representations can legitimately be evaluated not only based on their fit with the world but also with respect to their fit with the needs of their users, ... nonepistemic values can play a legitimate role as factors that override epistemic considerations in assessing scientific representations for practical purposes" (2014, 1). Setting aside for the moment their language about nonepistemic values "overriding" epistemic ones, they are clearly justifying roles for nonepistemic values based on the fact that scientific assessment requires more than assessing "fit with the world." For Elliott and McKaughan, the assessment of scientific representations is partly a matter of assessing their "fit with the needs of their users"; therefore, they deny the consequent of (1), and in turn the antecedent, and take nonepistemic values to be relevant to the assessment of scientific representations.

Although Helen Longino's contextual empiricism is not typically regarded as a form of the aims approach, her view of scientific assessment is actually very similar because it focuses on the extent to which scientific representations meet the needs of their users. She justifies the role of values in science based on the presence of a logical gap between a state of affairs,  $x$ , and a hypothesis,  $h$ . She emphasizes that background beliefs or assumptions are necessary in order to fill the gap and establish evidential relationships between states of affairs and hypotheses (1990, 41), and values play a role in assessing those background beliefs. In her book *The Fate of Knowledge* (2002), she makes abundantly clear that this process of assessment is about more than truth alone. She argues that the concept of "conformation" is more appropriate than "truth" for describing the goal of scientific assessment, both because conformation can be described in terms of degrees and because it clearly expresses how scientific assessment is relative to people's interests. As Longino puts it, "We may, for example, think of propositions as conforming to their intended objects just in case there is sufficient alignment between the elements of the proposition and elements of the object that we can successfully carry out our projects with respect to the object" (2002, 118–119). Thus, Longino's focus on the usability of representations clearly goes beyond truth alone, but

<sup>4</sup> Some scholars, such as Hugh Lacey (2017), advance an intermediate position. Lacey holds that scientists can adopt different cognitive attitudes toward scientific hypotheses or theories. Some attitudes focus solely on epistemic considerations, whereas others include other non-epistemic considerations.



she does not provide a detailed account of how to assess scientific representations in terms of their conformation.

Turning finally to the argument from inductive risk, the most influential proponent of this approach in recent years has been Heather Douglas (see e.g., 2000, 2009). At first glance, her argument is less explicit about denying that scientific assessment is focused solely on truth, but her commitment to this presupposition becomes apparent when one realizes that she is focused not solely on determining whether a hypothesis is true but rather on whether scientists should “put forward” or propound a hypothesis for decision makers (Elliott, 2013).<sup>5</sup> She emphasizes that her argument is directed at scientists considering “the context of use and the potential consequences of error when deciding *what to say*” (2009, 66, italics added). She claims the determinant of whether one sides with the proponents or opponents of the value-free ideal is whether one thinks scientists should be considering the social context when deciding *what claims to make*:

If, with Rudner and Churchman, one thinks that scientists should consider the potential consequences of error when deciding which claims to make, then values have an unavoidable place in scientific reasoning. If, with Levi and McMullin, one thinks that scientists should not be considering the potential consequences of error, then scientists can safely exclude social and ethical values from the heart of scientific reasoning. (Douglas, 2009, 66)

Thus, Douglas’s argument is predicated on the notion that scientific assessment can be about more than truth alone; it incorporates the goal of making scientific claims responsibly in their context of use, given the ever-present potential for error.

In sum, we have argued that opponents of the value free ideal, who reject either (1) or (A), are committed to a view of scientific assessment that goes beyond truth. While many have pointed to the relevance of non-epistemic factors in scientific assessment, our argument goes further: a practical conception of scientific assessment is an essential presupposition of those who object to the value-free ideal. As such, a successful challenge to the value-free ideal requires the specification of a viable form of assessment that can incorporate non-epistemic values. There are few, if any, general systematic accounts of this sort. In the next section, we provide one.

### 3 The adequacy-for-purpose framework

In this section, we provide a general practical account of scientific assessment. To build such an account, we expand upon the adequacy-for-purpose view, which was originally

<sup>5</sup> There is a sense in which Douglas arguably regards scientific assessment as being focused solely on truth. She insists that (non-epistemic) values should not play a direct role in hypothesis assessment, meaning that they should not play the role of evidence or be used to override evidence (Douglas, 2009, 102). So if one regarded scientific assessment solely as a matter of assessing the evidence for a hypothesis, one could argue that Douglas accepts both (A) and (C). However, there is a more important sense in which Douglas regards scientific assessment as not being focused solely on truth. As discussed in the text, Douglas insists that (non-epistemic) values are relevant when scientists decide what claims to put forward in policy contexts. This sense of hypothesis assessment is arguably the central one in Douglas’s work, given that she rejects the value-free ideal, and this is the sense in which Douglas rejects (A).



developed as an approach to assessing scientific models. Such a view, we claim, can be generalized beyond its original context and demonstrate how non-epistemic factors can help guide scientific assessment.

The adequacy-for purpose (henceforth AFP) view of scientific assessment arose in response to a problem with the evaluation of models. Typical strategies for scientific assessment sought to establish the truth, or the probability of truth, of the object being assessed, which was typically a hypothesis or theory. Models – at least the kinds used by scientists in practice, like the scale model of the San Francisco Bay (Weisberg, 2012) or climate models (Parker, 2009) – were *known at the outset* to be less-than-true in some way. These models contained idealizations and simplifications introduced by scientists in attempts to increase the fidelity of a particular aspect of the model, make the model practically useful (e.g., by fitting it in a warehouse), or (more simply) save computational costs. These departures from truth indicated that model assessment must proceed differently from traditional assessment – since the models (in their entirety) are not candidates for truth.<sup>6</sup>

The AFP view grows out of a broader tradition of understanding models as tools (e.g., see Morgan & Morrison, 1999) that are designed or selected to meet a set of epistemic and non-epistemic purposes. Model evaluation proceeds by assessing a hypothesis regarding a model's ability to meet criteria that specify "the right tool for the job." Rather than assess the truth or representational fidelity of a model tout-court, the evaluation of models happens via adequacy-for-purpose hypotheses, which specify whether the model is apt for the particular task at hand. For example, an AFP hypothesis might take these forms: "Model M is adequate for the purpose of predicting the changing odds of extreme weather events due to climate change," or "Model M is adequate for the purpose of exploring the implications of a new theory of star formation," or "Model M is adequate for the purpose of determining if a nuclear waste storage plan will be safe and effective." By determining whether evidence supports the AFP hypothesis rather than hypotheses about the truth of the model itself, model evaluation can proceed even when there are aspects of a model that knowingly depart from reality.

Despite the popularity of early writing (Parker, 2009), the adequacy for purpose approach to model assessment has only recently been given a rigorous articulation in Parker (2020). Under her interpretation, asking if a model is adequate for purpose is nearly synonymous with asking "Can the model be used to do the job?" (Parker, 2020, 461). Purposes (or "jobs") can be epistemic, practical, social, or a combination thereof; goals for particular models might be to predict the number of droughts in a region, explore the implications of certain policy options, or help a student better understand a complex concept. Purposes might permit multiple interpretations, and they might be achieved in a variety of different ways.

In order to perform a successful assessment then, one needs conditions of adequacy. Parker develops several related notions of adequacy that cover different contexts. The

---

<sup>6</sup> There is a rich debate about how idealizations should be assessed and understood, as well as to what degree their epistemic value is connected to truth or not. For a broad overview of contemporary positions, see Potochnik (2020).

most useful for our purposes is adequacy with respect to reliability within a type of use (Parker, 2020, 462), which she specifies in the following way:

ADEQUACY: M is ADEQUATE-FOR-P iff, in C-type instances of use of M, purpose P is very likely to be achieved.

First, it should be noted that when assessing an AFP hypothesis using this schema for adequacy, the degree of likelihood of achieving P necessary to declare a model adequate is relative and will differ across cases depending on the implications that achieving P has. Following Parker, we take no stand on how these probabilities should be understood (e.g., as propensities or frequencies).

C-type instances include the characteristics of the users (U), who employ a methodology (W), in background contexts (B), to represent a target system with a particular degree of fidelity (T). The AFP view makes particularly salient that when assessing a model's adequacy, one must examine not just the model's fidelity to the target system of interest, but characteristics of the user, their ways of employing the model, and the background in which they operate. Parker talks of U, W, B, and T as constraints on a problem space constituted by P. Therefore, in order for a model to be adequate for purpose, "it must stand in a suitable relationship not just with a representational target T, but with a target T, user U, methodology W, circumstances B and goal P jointly" (Parker, 2020, 464). It is important to note that a model can fail to be adequate not just because it falls short along some dimension of fidelity (T), but also because it fails to meet any of the other constraints (e.g., U, W, and B) defined by P. Non-epistemic factors often matter to assessing an adequacy for purpose hypothesis just as much as epistemic factors, and there may be "interactive effects" between the epistemic and non-epistemic (Parker, 2020, 465).

While adequacy for purpose does not admit of degrees, one can sometimes assess the relative fitness for purpose of particular models, which may come in degrees. If it can be assumed that a purpose has a complex structure with a rank order, where  $P_{\min}$  is equivalent to being minimally adequate for the purpose and  $P_{\max}$  the maximal desired extent, then one can assess the relative fitness of a model. Parker (2020, 464) defines this as:

FITNESS: A model M's FITNESS-FOR-P is greater to the extent that M is ADEQUATE-FOR-P for a higher-ranking member of  $P = \{P_{\min}, \dots, P_{\max}\}$ .

For example, a pedagogical model that is likely to result in a significant increase in students' knowledge is more fit-for-purpose than a model that results in a moderate increase in students' knowledge, though both models may be adequate for the purpose of learning.

An example with more context is helpful to illustrate how the AFP view of model assessment works in practice. Haasnoot et al. (2014) deployed something very much like the AFP view when developing and assessing a model for evaluating climate adaptation pathways in the Rhine River delta for the Dutch government. Their purpose was to assist the Dutch government in making better adaptation decisions for the Rhine by screening and ranking potential policy actions, aware that the policy actions taken today may constrain subsequent actions in the future. First, to fulfill this purpose, the

modelling team chose a specific kind of model – an integrated assessment model – that was capable of integrating the social, economic, and environmental aspects of the problem. To evaluate their model, they worked with stakeholders to establish conditions of adequacy. They knew that the model would need to “produce credible outcomes with sufficient accuracy for the screening and ranking of promising actions,” and established metrics for evaluating accuracy in this way. Furthermore, they knew the model would need to be flexible; there would be a large number of simulation runs, and the ability to change the scenarios in the model would be advantageous. In response, they chose a model structure that allowed for fast calculation, which in turn allowed them to accept a certain degree of inaccuracy in the model provided it did not rise above the level of indifference for decision making. Because this model was able to provide the required information under the various constraints imposed by the purpose at hand, Haasnoot et al. judged it adequate for purpose.

One might begin to notice how the AFP view could ground an approach to scientific assessment that goes beyond truth. As we argued in Section 1, what opponents of the value-free ideal essentially argue is that scientific assessment incorporates non-epistemic aims or purposes. As we show in the following section, it has been difficult to express these aims or purposes precisely in the previous literature on values in science, but the AFP approach to assessment looks more promising in this respect.

However, in order for the AFP view to be broadly applicable to the cases under dispute in the values-in-science literature, the view needs to be applied beyond modeling. Fortunately, there is nothing inherent in the AFP view that restricts this kind of evaluation to models. It is an orientation to purposes that seems to be essential, not the use of models; it seems then that any purpose-directed scientific activity that attempts to represent the world could potentially be accommodated by the AFP view. Take, for example, purpose-directed uses of the Newtonian theory of mechanics. One could assess whether the theory is adequate for the purpose of building bridges given users with a certain acumen, in a particular geographical context. If the resulting bridges were very likely to fulfill their purposes of, for example, sustaining car or pedestrian travel, then such a theory could be adequate-for-purpose given the requisite constraints. Similarly, the AFP view could be applied to the selection of other objects, like measuring instruments, whose outputs are taken as representations of the objects measured.

As we will show, the AFP view can be a powerful tool for those analyzing the role of values in science: it provides a means of performing scientific assessment that can accommodate non-epistemic factors in a precise fashion. One need not completely give up on truth; one can still assess the truth of the adequacy-for-purpose hypothesis.<sup>7</sup> What the AFP view does permit, however, are additional constraints related to the user and

---

<sup>7</sup> Our proposal might be seen to vindicate the value-free ideal, in the sense that this kind of assessment is one focused on the truth of a hypothesis. Although social values play a role in constructing the adequacy-for-purpose hypothesis, most proponents of the value-free ideal have long held that social values can help decide which hypotheses to investigate (Douglas, 2009). While we discuss how our proposal could result in a partial rapprochement between proponents and opponents of the value-free ideal in the conclusion below, proponents of the ideal have traditionally tried to exclude values from the core of scientific assessment by locating values elsewhere in decision-making. Our proposal is precisely the opposite (i.e., it incorporates social values in assessing the adequacy of “plain” hypotheses) and thus does not vindicate the value-free ideal as normally understood.

the context to influence judgements about the adequacy of a hypothesis, theory, model, or tool. In doing so, the AFP view can be used to advance the literature on values in science by providing an account of how scientific assessment might accommodate contextual factors.

## 4 Recasting values in science using the adequacy-for-purpose framework

To illustrate how the AFP framework enhances the literature on values in science, we now examine three ways it can improve on previous work. We argue that it provides a better way of describing the purpose-relativity of many scientific hypotheses, a better way to examine the role of consequences in assessing hypotheses, and a better way to describe the relationships between epistemic and non-epistemic values in scientific assessment. In other words, we contend that the extant literature on values in science has not provided a systematic account of scientific assessment that grounds its appeal to non-epistemic considerations. We show that the AFP framework provides a promising way of structuring hypotheses for understanding the value-laden character of assessment.

### 4.1 The purpose-relativity of scientific hypotheses

The examination of inductive risk was largely responsible for the resurgence of interest in values and science. As such, it serves as an apt case to examine how AFP hypotheses work in this context and the benefits of using them. The AFP framework can, for example, better express the purpose-relativity of scientific hypotheses. In his original (1953) explication of the argument from inductive risk, Richard Rudner argued, “How sure we need to be before we accept a hypothesis will depend on how serious a mistake would be” (1953, 2). He insisted that judging the seriousness of a mistake – and thus whether a hypothesis is acceptable – relies on non-epistemic values. That is, if the consequences would be severe according to some set of non-epistemic values, there should be a higher (or lower) epistemic standard of “acceptance” than there would be if the consequences were trivial.

Despite arguing for context-of-use considerations, Rudner lacked an explicit way of representing such context in the hypotheses being assessed. For example, Rudner considered the following hypothesis:

(h<sub>R</sub>) “that, on the basis of a sample, a certain lot of machine stamped belt buckles was not defective” (1953, 2).

Call this hypothesis a “plain hypothesis,” since it is not purpose relative.<sup>8</sup> On Rudner’s analysis, the standards for evaluation of this hypothesis will differ based on whether these buckles are to be used for securing pants or are to be used for securing bodies in a

<sup>8</sup> It seems that the intention was for this to be purpose-relative; the fact that they are “stamped” is perhaps a way to indicate to the reader the quality and potential applications of such belt buckles. Such quality concerns then may help relativize “defective” to a particular context. If so, the substantive assumptions and shared knowledge required to simply understand this belt-buckle hypothesis helps show the difficulties in judging its truth, difficulties that the AFP approach can overcome.

car crash (i.e., seat belt buckles) because of the relative seriousness of mistakes in judging them defective. Yet, Rudner's articulation of the hypothesis is as if it stands apart from use-considerations; it seemingly appears that *the same plain hypothesis*  $h_R$  is being accepted or rejected in different instances. As Jeffrey (1956) points out in response, the same hypothesis does not seem to capture differences between cases. For example, when applying this reasoning to a plain hypothesis about the safety of the polio vaccine, Jeffrey notes: "One cannot, by accepting or rejecting the hypothesis about the polio vaccine, do justice both to the problem of the physician who is trying to decide whether to inoculate a child, and the veterinarian who has a similar problem about a monkey" (1956, 245). Jeffrey was pointing out that contextual assessments of plain hypotheses seem to lead to odd consequences, where a plain hypothesis could be accepted and rejected by a single scientist at the same moment in time. This is not inconsistent – the acceptance or rejection are based on different proposed uses and thus different epistemic standards – but it is still confusing that the same hypothesis  $h$  is both affirmed and denied.

One might argue that Rudner has a good response to this worry, namely, that he would call for scientists to decide whether to accept a hypothesis based on *all* its potential consequences in *all* the contexts in which it could be used. To do otherwise might seem to be irresponsible. However, we contend that Jeffrey's worry is still relevant insofar as Rudner does not have a framework for specifying precisely which consequences and contexts are considered by the scientists who assess a plain hypothesis. Thus, if other scientists come to different conclusions about whether to accept the plain hypothesis under consideration, or if circumstances change in the future and thus alter the consequences of accepting the hypothesis, there is no formal way to determine which consequences or contexts were considered when assessing it. Shifting to our framework of adequacy-for-purpose hypotheses provides a formal way to specify these considerations.

The adequacy for purpose approach can be used to disambiguate the various statements being assessed in a non-competitive way. There are plain hypotheses that do not make reference to particular purposes, and then there are adequacy-for-purpose-style hypotheses that do make reference to particular purposes. What can be gleaned from Rudner's discussion is that it is often more apt to consider AFP hypotheses rather than plain hypotheses.<sup>9</sup> In this case, the AFP hypotheses are different because the purposes under consideration are different:

- (1) AFPh<sub>R1</sub>: Hypothesis  $h_R$  is ADEQUATE-FOR-Pants (where  $h_R$  is ADEQUATE-FOR-Pants iff employing  $h_R$  within the production of waist-belt buckles makes it highly likely that one's pants will stay up when the waist-belt and buckle are used.)

OR

- (2) AFPh<sub>R2</sub>: Hypothesis  $h_R$  is ADEQUATE-FOR-Seat Belts (where  $h_R$  is ADEQUATE-FOR-Seat Belts iff employing  $h_R$  within the production of seat belt

<sup>9</sup> The extent to which scientific practice does (or should) involve AFP-style hypotheses as opposed to plain hypotheses is an important, and potentially contentious, question to consider. We will return to it briefly in the conclusion.

buckles makes it highly likely that the buckle will keep an adult body secure in a car crash when used with a particular type of restraint).

There are important differences between  $AFPh_{R1}$  and  $AFPh_{R2}$  and Rudner's original  $h_R$ . The most obvious difference is that the AFP hypotheses specify purposes that include particular contexts. The "high likelihood" demanded for pants security and that demanded for seat belt safety can and should vary according to the circumstances, and that likelihood helps determine if the sampling method referenced in either AFP hypotheses is indeed adequate. Affirming  $AFPh_1$  while denying  $AFPh_2$  no longer seems odd because they are different hypotheses that do not compete with or exclude each other. One can do justice to different contexts within scientific assessment.

## 4.2 Examining consequences in hypothesis acceptance

Rudner's early form of the argument from inductive risk (AIR) is premised on the notion that scientists can consider the potential consequences of accepting or rejecting a hypothesis. However, the myriad consequences of accepting a hypothesis are often underdetermined, opaque, or unknowable. Modern proponents of the AIR, such as Heather Douglas, acknowledge this and claim that scientists merely need to predict the consequences as best they can. Nevertheless, because proponents of the AIR have focused previously on plain hypotheses, they do not have a formal way to represent or keep track of the consequences to be considered. Again, the importance of this is apparent in Jeffrey's response to Rudner's version of the AIR: "But what determines these consequences? There is nothing in the hypothesis, 'This vaccine is free from active polio virus', to tell us what the vaccine is for, or what would happen if the statement were accepted when false" (1956, 242).

The AFP framework not only calls for the provision of purpose-specific hypotheses but requires the articulation of conditions of adequacy for judging them. These elements of AFP reasoning require considering possible consequences. To illustrate, consider another case used to argue for the inclusion of non-epistemic values in the internal aspects of science via inductive risk.

Douglas (2000) analyzes toxicology experiments, particularly animal model studies used to assess dioxins, to demonstrate how non-epistemic values can indirectly but permissibly inform the internal aspects of scientific reasoning. She argues that these values enter reasoning at several decision sites, including the choice of a level of statistical significance and the characterization of evidence. For example, choosing a level of statistical significance invites non-epistemic values into the consideration of this hypothesis via the errors that might result. Though often set conventionally, the choice of a level of statistical significance balances the chance of mistakenly accepting the hypothesis as true (i.e., a false positive or type I error) with that of mistakenly rejecting the hypothesis as false (i.e., a false negative or type II error). When these toxicology experiments on model organisms are used as part of a process to *establish regulations regarding human health*, an excess of false negatives will have the result of "causing dioxins to appear less harmful than they actually are, leading to underregulation of the chemicals" (Douglas, 2000, 567). Under-regulation could foreseeably result in costs to human and environmental health, whereas overregulation

could foreseeably result in greater costs to industries guided by such regulations. Douglas claims the statistical significance should be set depending on how one values these effects, which involves consideration of non-epistemic values.

Though Douglas provides significant context regarding the consequences that need to be considered when making these decisions, she does not have an account of scientific assessment that provides a systematic way of representing or keeping track of these consequences. The hypothesis that can be extracted from many of her examples seems to be the following:

( $h_D$ ) Model animals exposed to dioxins exhibit cancer at a significantly greater rate than control animals not exposed to dioxins.

Unfortunately, the structure of  $h_D$  does not clarify which consequences to consider; it is not the case that human health is always at risk. Suppose, for example, that the manufacturing processes that produce dioxins as byproducts were just beginning to be developed by a chemical company, and they were screening the byproducts to determine if they might cause harm. The consequences of errors in such a case would be very different from a situation in which the manufacturing processes were already in widespread use for important industrial purposes.

AFP hypotheses provide a structure for indicating potential consequences in a way that plain hypotheses do not. In contrast to a plain hypothesis, an AFP hypothesis could specify whether  $h_D$  was being considered for screening purposes or for regulatory purposes, thereby suggesting the consequences to be considered. Furthermore, one expects that the setting of adequacy conditions would involve the consideration of consequences and the specification of particular constraints such that foreseeable negative consequences are avoided. For example, one might formulate hypotheses like these:

AFP $h_{D1}$ : Hypothesis  $h_D$  is ADEQUATE-FOR-Private Sustainability Assessment (where  $h_D$  is ADEQUATE-FOR-Private Sustainability Assessment iff the employment of the hypothesis within a specified decision-making structure makes it highly likely that chemical companies of a certain sort will achieve their sustainability goals.)

OR

AFP $h_{D2}$ : Hypothesis  $h_D$  is ADEQUATE-FOR-Federal Regulation (where  $h_D$  is ADEQUATE-FOR-Federal Regulation iff the employment of the hypothesis within a regulatory regime like that used by the U.S. Environmental Protection Agency makes it highly likely that resulting regulations will protect human and environmental health).

Here we have used “within a regulatory regime...” and “within a specified decision-making structure...” as a place holder for a more fine-grained specification of fidelity (T), the users (U), methodology (W), and background assumptions (B) that need to be specified as conditions of adequacy. The need to specify conditions of adequacy at all, however, directs one to consider the consequences of using a hypothesis, model, or tool



in a particular context. In order for something to be adequate for purpose, the conditions must be such that obstacles that would prohibit the fulfillment of the stated purpose are avoided. That means, for example, that one could not judge  $h_D$  adequate for the purpose of federal regulation, if it were analyzed in a way that would knowingly jeopardize health when used within a certain regulatory regime. The AFP framework makes perspicuous how the conditions of adequacy used in assessing hypotheses relate to consequences and purposes. In short, Jeffrey's rhetorical objection offered in the opening to this section is overcome: the use of AFP hypotheses and their related conditions of adequacy *can* tell us what a hypothesis (or model or method) is for, and *can* help one realize what would happen if it were in fact inadequate.

### 4.3 Describing relationships between epistemic and nonepistemic values

The AFP framework can also help specify how non-epistemic values relate to epistemic ones within scientific assessment. Many ascribe to the "lexical priority of evidence" (as described in Brown, 2013), which specifies that non-epistemic considerations should always play a secondary role to epistemic ones (see e.g., Steel, 2017). Others, however, reject the lexical priority of evidence. For example, Elliott and McKaughan explicitly take the view that scientists can "prioritize" nonepistemic values over epistemic ones, and they sometimes express this in terms of the idea that nonepistemic values can "trump" epistemic ones (2014, 5–6; see also Brown, 2017). The AFP framework defended here can clarify what is at issue and potentially dissolve these debates.

The language currently used to describe the relationship between epistemic and nonepistemic values is highly metaphorical and imprecise, which can hide different notions of what it means for some values to take "priority" over others. For example, Elliott and McKaughan (2014) envision scenarios in which two hypotheses are being compared, with one hypothesis having been assessed more favorably from an epistemic perspective and the other hypothesis having been assessed more favorably from a nonepistemic perspective. They claim that it is appropriate in some cases for scientists to accept the hypothesis that has been assessed more favorably from the nonepistemic perspective. But it appears that others have different scenarios in mind when picturing nonepistemic values "trumping" or being "prioritized" over epistemic ones. For example, when Steel (2017) criticizes Elliott and McKaughan's view, he seems to consider scenarios in which typical standards of adequate science, such as proscriptions against fudging data or rigging experiments, are abandoned for nonepistemic reasons. One could also view nonepistemic values to be trumping epistemic values if they affected scientific practice in ways that hindered the attainment of truths, such as by blocking particular research projects (Hicks, 2014). It is not clear that these scenarios should all be handled in the same way, but without providing a more precise way of characterizing what it means for nonepistemic values to be prioritized over epistemic ones, the potential to conflate these scenarios and talk past one another is significant (see e.g., Brown, 2017; Steel, 2016, 2017).

The AFP framework is ideally suited to provide a more precise characterization of the way epistemic and nonepistemic values relate to each other. Consider, for example, a case that Elliott and McKaughan (2014) discuss as part of their argument that nonepistemic values can be prioritized. They describe an expedited approach to chemical risk assessment that was tested by the state of California. The approach was

slightly more prone to false positive errors (i.e., declaring chemicals to be harmful when they were not) and false negative errors (i.e., declaring chemicals to be safe when they were not) than the standard risk assessment approach already in place. However, it was vastly faster than the standard approach. According to Elliott and McKaughan, this made it more appropriate for achieving the aim of protecting society against exposure to toxic substances, because the standard approach was slow enough that it left numerous toxic chemicals on the market until they could be assessed.

What becomes apparent when this case is analyzed in terms of adequacy-for-purpose is that rather than one kind of value “trumping” the other, there are epistemic and non-epistemic requirements that are satisfied *jointly* based on the purposes to be achieved.<sup>10</sup> To see this point, consider how the chemical risk assessment case described above might be handled within the adequacy for purpose framework:

(AFPh<sub>EM</sub>) Standard Risk Assessment (SRA)/Expedited Risk Assessment (ERA) is ADEQUATE-FOR-Chemical Risk Assessment (where a method is ADEQUATE-FOR-Chemical Risk Assessment iff in instances where the government employs it to assess toxic chemicals, effective regulation is very likely to be achieved).

Now “effective regulation” might be a contentious expression admitting of many different interpretations. What is clear is that “effective regulation” (like many scientific purposes or goals that might be analyzed using the AFP framework) is a multi-dimensional notion; it is not as if effective regulation can be clearly determined simply by assessing, for example, the accuracy of the two methods. The AFP framework suggests conceiving of this purpose as constituted by a multi-dimensional problem space with the set of constraints discussed previously (T, U, W, B). For example, effective regulation might require a number of usage constraints: assessment methods might need to be reasonably standardized so that chemical companies can predict how their products will be handled by regulators, they might need to be conducted in a timely fashion, the results would need to be interpretable by regulators, etc. They would specify fidelity constraints, such as that assessments would need to be accurate enough to ensure a specified degree of human safety. Risk assessment methods that fail to meet these constraints would be deemed *inadequate*-for-purpose; adequacy-for-purpose hypotheses can be rejected if *any* of the various criteria for adequacy are not met. Furthermore, a fast method that is only moderately accurate and easy to use could be adequate, just as a slow method that is very accurate and moderately difficult to use could be adequate. This is a consequence of the constraints requiring joint satisfaction: there may be more than one way to achieve a given purpose.

In sum, the AFP framework provides a way of characterizing epistemic and nonepistemic values as being satisfied jointly: both epistemic and nonepistemic properties are evaluated *in light of* an established set of purposes. In other words, the purposes specify the extent to which particular properties need to be satisfied. Within this framework, it would be awkward to claim that some properties are prioritized over

<sup>10</sup> Janet Kourany (2010) also holds a view of scientific assessment that calls for the joint satisfaction of epistemic and non-epistemic values, but she lacks a detailed way of representing the nature of this joint satisfaction. The AFP framework provides this missing detail.

others; rather, each epistemic and nonepistemic property just needs to be satisfied to the extent specified by the purposes that have been set. Consider an example where an accurate risk assessment method is unfit for purpose because it fails to meet the constraints placed on it by its adequacy conditions. One could propose a highly tailored, and very accurate, method of risk assessment that essentially analyzed each chemical differently on a case-by-case basis. Despite having the necessary fidelity characteristics, such an assessment would fail to be adequate for purpose because it would not have other necessary (nonepistemic) characteristics: manufacturers could not predict how chemicals would be handled, it would be slow, and so on. What this demonstrates is that purposes – when they are agreed upon – demand that epistemic criteria and non-epistemic criteria both be satisfied. When assessing adequacy for purpose there is no need to talk about prioritizing or trumping.

Nevertheless, one might still think that different kinds of values need to be prioritized or weighed against each other when two claims or methods are being compared to one another. For example, the debate surrounding Elliott and McKaughan’s example is not about whether SRA or ERA are each adequate, but rather which is preferable to the other. In the process of choosing one method, it might seem like one is prioritizing the characteristics (i.e., values) of that method over those of the other one. In a sense this might be true, but this “weighing” of values can be precisely expressed within the AFP framework using fitness-for-purpose.

The fitness for purpose of a risk assessment method can be described as follows:

**FITNESS:** A risk assessment method’s FITNESS-FOR-PURPOSE is greater to the extent that the risk assessment is ADEQUATE-FOR-P for a higher-ranking member of  $P = \{P_{\min}, \dots, P_{\max}\}$ .

Fitness is a relative measure of the degree to which something is purpose apt. In this case, assuming that SRA and ERA both meet the minimum requirements for the particular purpose they will serve, they both meet  $P_{\min}$  for all constraints. But since fitness admits of degrees, one may be more fit than the other for the purpose at hand. Since purposes are constituted by a set of constraints (T, U, W, B), one cannot choose between methods by focusing on a single constraint, such as fidelity with the target. Rather, multiple constraints determine the extent to which a method is fit for purpose. In the case of chemical risk assessment, for example, fitness also depends on the timeliness of the method and the extent to which it meets the needs of the users under particular background conditions. Although it might be difficult in some cases to create a fitness-for-purpose ranking from  $P_{\min}$  to  $P_{\max}$ , that ranking provides a precise way to determine what combination of characteristics is most important given a particular purpose. Under the fitness-for-purpose framework, two characteristics or values (e.g., speed and fidelity) could be weighed against each other in the sense that a risk-assessment method with a particular combination of speed and fidelity could be judged to be more fit for purpose than a risk-assessment method with a different combination of speed and fidelity under particular background conditions.

Thus, the AFP framework provides a precise way of representing several ways in which values can relate to each other in scientific assessment. First, both epistemic and nonepistemic values can play a role in choosing which purposes are most appropriate to achieve in a given context. Second, when assessing whether something is adequate for

purpose, epistemic and nonepistemic values must be jointly satisfied in a way that does not involve particular kinds of values “trumping” others. And lastly, when assessing fitness for purpose, combinations of characteristics (i.e., values) are weighed against each other in the sense that one must determine which combination results in greater fitness. It is only in this last, limited sense that particular values (which will typically be assessed in combination) might be said to “trump” others in scientific assessment. Thus, debates about whether nonepistemic values can be prioritized over epistemic values in scientific assessment can be cashed out as disagreements over what combination of values render a hypothesis, model, or method most fit for purpose and what kinds of purposes are appropriate to begin with.

## 5 Conclusion: Future applications of the AFP framework

We have argued for a conceptual shift in the literature on values in science. Rather than focusing on how non-epistemic values can figure into scientific assessment, we suggest analyzing how scientific assessment can accommodate non-epistemic values. We have argued that the AFP framework can provide a general framework for this kind of scientific assessment. Moreover, we have shown that adopting this framework can move the literature on values and science forward. We have argued that it provides a better way of describing the purpose-relativity of scientific hypotheses, a better way to examine the consequences that play a role in assessing hypotheses, and a better way to describe the relationships between epistemic and non-epistemic values in scientific assessment.

Moving forward, there are both more modest and more ambitious ways one might use the AFP framework to further illuminate the roles of values in science. More modestly, one might explore additional avenues for using the AFP framework to clarify key concepts used in the values-in-science literature. For example, a number of authors have worried that the word “value” has been used to describe an overly broad array of both individual and institutional factors that can influence scientific research (see e.g., Biddle, 2013; Ward, 2021). The AFP framework could potentially help to address this difficulty by providing a more systematic way of characterizing all these factors. Perhaps they could be categorized based on the ways they contribute to the purposes described in the AFP hypotheses or the different kinds of constraints that determine whether the purposes have been met (e.g., the users, the way a tool is used, various kinds of background conditions, and so on). Thus, whether or not all these factors continue to be labeled “values,” the AFP framework might help to organize them more systematically.

One might also explore a more ambitious way of applying the AFP framework. In Section 1, we argued that premise (1) plays a crucial implicit role in debates over values in science. We showed how proponents of the value-free ideal accept both the antecedent and the consequent of the premise (thereby insisting that scientific assessment is about truth or other epistemic goals alone), while critics of the value-free ideal reject both the antecedent and the consequent of the premise (thereby denying that scientific assessment is about truth or other epistemic goals alone). The AFP framework suggests that perhaps a rapprochement can be generated between these seemingly opposed positions. In Section 1, which pitted the proponents and the opponents of

the value-free ideal against each other, we focused on what we have called “plain hypotheses,” which are not purpose-apt and do not require conditions of adequacy. But one might return to premise (1) and focus instead on AFP hypotheses. Is the antecedent of premise (1) true when AFP hypotheses are under consideration? Is it possible to focus solely on truth (or whatever alternative epistemic goal one might have) with *this* kind of hypothesis? An affirmative answer would mean that proponents of the value-free ideal could maintain that science is ultimately about the pursuit of truths or other epistemic goals (insofar as AFP hypotheses are assessed solely based on their truth). However, critics of the value-free ideal could still insist that scientific reasoning incorporates non-epistemic values (insofar as non-epistemic values are relevant to assessing AFP hypotheses).<sup>11</sup> The key to this rapprochement would be to shift the entire conversation about values in science so that it focuses on AFP hypotheses; on this view, the reason the proponents and opponents of the value-free ideal have seemed so far apart in their views is that they have been focusing on the wrong (plain) kind of hypotheses.

We recognize that this application of the AFP framework would be controversial. Critics of the value-free ideal would presumably argue that the same kinds of reasoning that established roles for non-epistemic values in the assessment of “plain” hypotheses will also apply to AFP hypotheses. For example, they might insist that the argument from inductive risk (Douglas, 2009) or the underdetermined relationship between data and theory (Longino, 1990) are just as relevant to AFP hypotheses as they are to any other hypotheses. But this response merits further discussion. It might be the case that some of these considerations could be incorporated into the formulation of the AFP hypotheses (e.g., into the standard for concluding that a purpose P is “very likely” to be achieved) so that the hypotheses as a whole could be evaluated solely in terms of their truth or other epistemic qualities. These are questions that go beyond the scope of this paper. For present purposes, we have tried to show that the AFP framework can strengthen the values-in-science literature by providing better ways of describing scientific hypotheses and the roles that epistemic and non-epistemic values play in assessing them.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

<sup>11</sup> One might, however, wonder how non-epistemic values can play a role in assessing the truth of AFP hypotheses without violating premise (1). The answer is that these non-epistemic values can play the role of evidence (and thus functionally operate as epistemic values) in the assessment of AFP hypotheses because those hypotheses make reference to goals or purposes. For example, the fact that a model can be used quickly might be a non-epistemic value under most circumstances (see Steel, 2016 for a potential objection). However, when assessing the AFP hypothesis that a model is fit for the purpose of generating results that can be used effectively within a particular regulatory regime, the fact that the model can be used quickly might operate as an epistemic value that provides evidence for the truth of the hypothesis.

## References

- Betz, G. (2013). In defence of the value free ideal. *European Journal for Philosophy of Science*, 3(2), 207–220.
- Biddle, J. (2013). State of the field: Transient underdetermination and values in science. *Studies in History and Philosophy of Science Part A*, 44(1), 124–133.
- Brown, M. J. (2013). Values in science beyond underdetermination and inductive risk. *Philosophy of Science*, 80(5), 829–839.
- Brown, M. J. (2017). Values in science: Against epistemic priority. In *Current controversies in values and science*. Routledge.
- Brown, M. J. (2020). *Science and moral imagination: A new ideal for values in science*. University of Pittsburgh Press.
- De Regt, H. W. (2017). *Understanding scientific understanding*. Oxford University Press.
- Douglas, H. (2000). Inductive risk and values in science. *Philosophy of Science*, 67(4), 559–579.
- Douglas, H. (2009). *Science, policy, and the value-free ideal*. University of Pittsburgh Press.
- Douglas, H. (2013). The value of cognitive values. *Philosophy of Science*, 80, 796–806.
- Elliott, K. C. (2013). Douglas on values: From indirect roles to multiple goals. *Studies in History and Philosophy of Science Part A*, 44(3), 375–383.
- Elliott, K. C., & McKaughan, D. J. (2014). Nonepistemic values and the multiple goals of science. *Philosophy of Science*, 81(1), 1–21.
- Elliott, K., & Richards, T. (Eds.). (2017). *Exploring inductive risk: Case studies of values in science*. Oxford University Press.
- Hicks, D. J. (2014). A new direction for science and values. *Synthese*, 191(14), 3271–3295.
- Hudson, R. (2016). Why we should not reject the value-free ideal of science. *Perspectives on Science*, 24(2), 167–191.
- Intemann, K. (2015). Distinguishing between legitimate and illegitimate values in climate modeling. *European Journal for Philosophy of Science*, 5(2), 217–232.
- Jeffrey, R. C. (1956). Valuation and acceptance of scientific hypotheses. *Philosophy of Science*, 23(3), 237–246.
- Khalifa, K. (2020). Understanding, truth, and epistemic goals. *Philosophy of Science*, 87(5), 944–956.
- Kourany, J. (2010). *Philosophy of science after feminism*. Oxford University Press.
- Kuhn, T. S. (1977). Objectivity, value judgment, and theory choice. *Arguing about Science*, 74–86.
- Lacey, H. (2017). Distinguishing between cognitive and social values. In *Current controversies in values and science* (pp. 15–30). Routledge.
- Longino, H. E. (1990). *Science as social knowledge: Values and objectivity in scientific inquiry*. Princeton University Press.
- Longino, H. E. (2002). *The fate of knowledge*. Princeton University Press.
- Lusk, G. (2020). Political legitimacy in the democratic view: The case of climate services. *Philosophy of Science*, 87, 991–1002.
- Lusk, G. (2021). Does democracy require value-neutral science? Analyzing the legitimacy of scientific information in the political sphere. *Studies in History and Philosophy of Science Part A*, 90, 102–110. <https://doi.org/10.1016/j.shpsa.2021.08.009>
- McMullin, E. (1982). Values in science. In *PSA: Proceedings of the biennial meeting of the philosophy of science association* (Vol. 1982, pp. 3–28). Philosophy of Science Association.
- Morgan, M. S., & Morrison, M. (1999). *Models as mediators: Perspectives on natural and social science*. Cambridge University Press.
- Parker, W. S. (2009). Confirmation and adequacy-for-purpose in climate modelling. *Aristotelian Society Supplementary Volume*, 83(1), 233–249. <https://doi.org/10.1111/j.1467-8349.2009.00180.x>
- Parker, W. S. (2020). Model evaluation: An adequacy-for-purpose view. *Philosophy of Science*, 87(3), 457–477. <https://doi.org/10.1086/708691>
- Potochnik, A. (2017). *Idealization and the aims of science*. University of Chicago Press.
- Potochnik, A. (2020). Idealization and many aims. *Philosophy of Science*, 87(5), 933–943. <https://doi.org/10.1086/710622>
- Rooney, P. (2017). The borderlands between epistemic and non-epistemic values. In *Current controversies in values and science* (pp. 31–45). Routledge.
- Rudner, R. (1953). The scientist qua scientist makes value judgments. *Philosophy of Science*, 20(1), 1–6.
- Steel, D. (2010). Epistemic values and the argument from inductive risk. *Philosophy of Science*, 77(1), 14–34.

- Steel, D. (2016). Accepting an epistemically inferior alternative? A comment on Elliott and McKaughan. *Philosophy of Science*, 83(4), 606–612.
- Steel, D. (2017). Qualified epistemic priority: comparing two approaches to values in science. In *Current controversies in values and science* (pp. 49–63). Routledge.
- Van Fraassen, B. C. (1980). *The scientific image*. Oxford University Press.
- Ward, Z. (2021). On value-laden science. *Studies in History and Philosophy of Science Part A*, 85, 54–62.
- Weisberg, M. (2012). *Simulation and similarity: Using models to understand the world*. Oxford University Press.

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.