The Difference-to-Inference Model for Values in Science

Stegenga Jacob & Tarun Menon. 2023. "The Difference-to-Inference Model for Values in Science" *Res Philosophica* 100(4): 423-447

For a pdf version of the published article, please email me: jms303@cam.ac.uk.

Abstract
The value-free ideal for science holds that values should not influence the core features of scientific reasoning. We defend the difference-to-inference model of value-permeation, which holds that value-permeation in science is problematic when values make a difference to the inferences made about a hypothesis. This view of value-permeation is superior to existing views, and suggests a corresponding maxim, namely that scientists should strive to eliminate differences to inference. This maxim is the basis of a novel value-free ideal for science.

1. Introduction

The value-free ideal for science has been deemed untenable by nearly every philosopher writing on the subject. Yet thus far there exists no general, precise characterization of how exactly values permeate scientific reasoning and what is epistemically problematic about value-permeation. Our aim in this paper is to offer such a characterization, in what we call the difference-to-inference model of value-permeation. Value-permeation is the state in which non-epistemic values influence the core of scientific reasoning, and the difference-to-inference characterization of this holds that value-permeation is a problem in need of resolution when the permeating values make a difference to the inferences made about a hypothesis. There are a variety of possible responses to the threat of value-permeation of scientific reasoning. A clear view of the nature of the threat of value-permeation suggests a very particular prescription about how to respond to the threat, and this prescription in turn forms the basis of a new value-free ideal for science.

Value-permeation is particularly prevalent in scientific work which informs policy. The most compelling argument for value-permeation—both the descriptive argument for the necessity of value-permeation and the normative argument for the appropriateness of value-permeation—is the inductive risk argument, and this argument is especially persuasive in policy-oriented scientific contexts in which decisions must be made based on the results of scientific work.

When values make a difference to inference about a hypothesis, values impede the achievement of scientific consensus about that hypothesis, and this, we argue in §2, is a threat to the status of that hypothesis as scientifically warranted. In §3 we further characterise the notion of consensus central to our account. We then describe some of the more prominent theses about the nature of value-permeation in science (§4). We argue that while the extant theses have some intuitive plausibility, they each face problems, which motivates what we call the difference-to-inference view of value-permeation (§4.4). We proceed to describe some of the most prominent proposals to mitigate the threat of value-permeation (§5). Here, too, we find a range of extant views of varying plausibility. We argue that the best way to mitigate the threat of value-permeation is the prescription that scientists should strive to eliminate values when they are difference-makers to inferences (§5.5). This, in turn, provides the basis of a novel version of the value-free ideal for science.

2. Consensus and Scientific Knowledge

Much of the literature on the epistemic role of consensus in science has focused on the question of whether and when one can use consensus as a reliable indicator of truth. For instance, can we regard the consensus among climate scientists about anthropogenic global warming as a good reason to believe that the Earth is in fact being warmed by human activity (Oreskes 2004)? Some argue that the consensus should not be regarded as a good reason for belief because the existence of the consensus is best explained not by epistemically relevant factors, but by sociological, political, or economic factors (such as the unavailability of grant money for dissenting views, or hierarchically maintained groupthink within the scientific community). As a consequence philosophers have attempted to delineate the conditions under which a consensus is legitimately "knowledge-based" (the term comes from Miller 2013; see also Tucker 2003).

However, this sort of evaluation of the epistemic significance of consensus is often premised on the assumption that consensus is merely a mark of scientific success, and not a constituent of scientific success. On this view, science should not aim for consensus, it should aim for some criterion of epistemic success, which may be characterised in a variety of ways—as truth or truthlikeness (Niinluoto 2014), knowledge (Bird 2022), understanding (Dellsen 2016), or empirical adequacy (van Fraassen 1980), among others—and any resulting consensus is an indicator that epistemic success has been achieved (see also Miller 2013; Stegenga 2016; Vickers 2023). (In what follows, we will consider knowledge to be the criterion of epistemic success, but much of what we say will also fit with alternative conceptions.) A consensus that emerges as a consequence of a search for knowledge will be appropriately knowledge-based, on this view, but if we independently aim for consensus then we sever the link between consensus and knowledge, and the consensus we arrive at could not be regarded as a reliable indicator of the truth. Knowledge is the aim of science, on this view, and consensus is a mere by-product of pursuing that aim.

We argue that this conservative position on the value of consensus is misguided. Consensus should be construed as partially constitutive of scientific success. It is not an *independent* aim of science; we agree that the only consensus worth having is (in a sense to be explained) knowledge-based. However, this does not entail that the goal of scientific consensus is entirely parasitic on the goal of scientific knowledge, because scientific knowledge is (again, in a sense to be explained) consensus-based. Neither knowledge nor consensus is a mere by-product of the other aim, but they are not separate aims either. Rather, in science, knowledge and consensus are deeply interwoven, so that we cannot sharply distinguish between the pursuit of knowledge and the pursuit of consensus. Epistemically legitimate consensus-forming procedures are important precisely because they expand our capacity to acquire scientific knowledge.

2.1 Strong consensus

In the domain of public reason, consensus theorists such as Habermas argue that consensus must be achieved not merely by a unanimously held conclusion but by arguments for that conclusion that are endorsed by all parties to the consensus (1996).

We argue, in what follows, that scientific knowledge requires a similar type of *strong* consensus.

Strong consensus involves more than mutual knowledge. If every member of a scientific community knows C, it does not follow that there is a strong consensus regarding C. Suppose two scientists, Aisha and Bheem, rely on the same evidence and arrive at the same conclusion, C. However, the inferential routes they take to get from evidence to conclusion are very different. They use different epistemic toolboxes – different background theories, different techniques of analysis, different computational and mathematical tools, and so on. Both toolboxes are sufficiently reliable to license the claim that both Aisha and Bheem know C to be true. However, Aisha does not accept Bheem's toolbox as valid, and *vice versa*, so Aisha does not believe that Bheem knows C, and Bheem does not believe that Aisha knows C. This is a case of mutual knowledge, where both parties know C, but neither party knows that the other party knows C. This is *weak* consensus.

If other members of my scientific community agree with me on the truth of some claim, but I believe that their reasons for believing that claim are (unlike mine) wrong, then I am not part of a strong consensus with them. For there to be a knowledge-based strong consensus, there must be *common* knowledge. In other words, not only must it be the case that both Aisha and Bheem know C, it must also be the case that Aisha knows that Bheem knows C (and *mutatis mutandis* for Bheem), Bheem knows that Aisha knows that Bheem knows C (and *mutatis mutandis* for Aisha), and so on. Standardly, this condition will be met when Aisha and Bheem have common knowledge that each party's epistemic toolbox is a source of knowledge. Strong consensus does not require that they actually use the same inferential tools, but it does require that each of them has the in-principle capacity to evaluate, endorse and use the other's tools. In other words, strong consensus about scientific claims is only possible if there is consensus about the validity of the epistemic tools used to justify the claims (a similar point is made in Lehrer and Wagner 1981). We refer to this as the condition of *shared epistemic toolboxes*.

2.2 Shared epistemic toolboxes

Shared epistemic toolboxes are not just a requirement for genuine knowledge-based consensus. They are also, we claim, a requirement for scientific knowledge. We are distinguishing here between knowledge *simpliciter* and scientific knowledge. The latter is a subset of the former. Only certain kinds of propositions—roughly, descriptions, explanations, and predictions of the behaviour and properties of concrete objects and processes—can qualify as objects of scientific knowledge. (So, on a plausible conception of values, they would not qualify as the kinds of things we can scientifically know.)

But scientific knowledge is not just delineated by its subject matter. The manner in which a proposition is known matters as well. Suppose Sara has an uncannily reliable instinct for discovering physical law. She can look at a physical system and just intuit the laws that apply. When she has this intuition, it comes with a deep sense of epistemic conviction, and it is always right. On certain conceptions of knowledge, Sara's intuitions count as knowledge. However, they cannot plausibly count as *scientific* knowledge because the inferential route Sara uses to get from evidence to conclusion cannot be shared with other members of the scientific community. They may be able to use Sara's inferences as a guide to physical law, but they cannot replicate her inferences themselves. Her epistemic toolbox is not shared with them. There is a publicity requirement on properly scientific knowledge (Merton 1942, Longino 1990), and particularly on practices of scientific justification (Gerken 2022). In our terms, scientific knowledge must be arrived at using epistemic toolboxes which can, in principle at least, be shared.

Weak consensus in science implies that some scientists doubt the justification or reliability of their peers, despite the fact that they share the same conclusion. That alone is a concern that motivates a resolution. Another way of putting this point is that weak consensus implies dissent about the epistemic toolbox of peers, and thus anything less than strong consensus implies some form of dissent. This means that for a piece of knowledge to count as scientific knowledge, it must be justified using tools about whose validity there is a consensus in the relevant community. This is the sense in which scientific consensus and scientific knowledge are intertwined as aims of science. Proper scientific consensus must be knowledge-based—specifically, it should be best explained by common knowledge acquired by a high enough proportion of the members of a particular community. Simultaneously, scientific knowledge must be consensus-based—for it to count as properly *scientific*, it must be arrived at using shared epistemic toolboxes.

The presentation so far has assumed a somewhat artificial binary: Either epistemic toolboxes are completely shared, in which case we have strong consensus, or there is some portion of an epistemic toolbox that is not shared, in which case consensus fails. Consensus in science, however, is better thought of as a matter of degree. Suppose Aisha and Bheem both rely on very similar models to calculate some physical quantity. Their models are identical to each other except for the value of one parameter, on which they disagree somewhat. Furthermore, neither of them can endorse the parameter value chosen by the other. Fortunately, the output of both their models are nearly insensitive to the value of the contentious parameter, and so prediction for the physical quantity, as calculated by both models, is effectively the same up to the desired level of approximation. Let us assume that the predictions of both models are accurate and reliable, and this counts as knowledge. On the binary view, this is a mere weak consensus, so the calculated quantity cannot be considered to be *scientific* knowledge. However, the degree of disagreement between the epistemic toolboxes here is small enough that we might say we have approximate strong consensus. Good enough, perhaps, to license a claim to scientific knowledge. On the other hand, if Aisha was using a completely different model from Bheem, and they could not endorse even the basic structure of each others' models, agreement on the predictions would be a clear case of weak consensus, and thus not scientific knowledge.

Scientific knowledge claims rest on (perhaps approximate) strong consensus. So, we now have a minimal sense in which consensus must be an aim of science. Without consensus about epistemic toolboxes, there cannot be scientific knowledge. In so far as knowledge is an aim of science, scientists need to also aim for consensus about epistemic toolboxes. But so far, it may seem that consensus should be an explicit aim of science only in this restricted circumstance. Aisha and Bheem need to aim at consensus about their toolboxes, but once they have arrived at shared toolboxes, any further scientific progress just requires them to use these shared toolboxes to gather knowledge. Knowledge gathered in this manner will be scientifically warranted, since the toolboxes are shared. It might seem that there would be no further need to explicitly aim for consensus, once this base level of consensus about toolboxes is acquired. Any further knowledge-based consensus will simply be a by-product of aiming at knowledge. We now have a picture in which consensus-formation needs to be an aim of science at an early stage, when epistemic tools are being developed, but then drops out in favour of traditional epistemic aims.

But this ignores a crucial point about science. Scientific progress is not just about acquiring more scientific knowledge with existing epistemic toolboxes. It is also about expanding the scope of what could be potential scientific knowledge. The epistemic toolboxes used in science are not static. A crucial part of scientific inquiry is developing new tools which allow us to acquire hitherto inaccessible knowledge, convert previously non-scientific knowledge into scientific knowledge, or figure out the precise limits of our scientific knowledge. Aiming at consensus formation often allows scientists to redraw the lines separating scientific from non-scientific knowledge. This suggests that consensus must be regarded as a constitutive aim of science, as we now argue.

2.3 Consensus-based knowledge

Aisha and Bheem are back, but now they disagree about the conclusion they have reached based on the evidence E. Aisha infers from E to C1, Bheem from E to C2, and C1 and C2 are incompatible. Neither scientist has made a mistake in the use of a shared epistemic toolbox, or failed to employ a shared epistemic toolbox that may be relevant to settling the disagreement. The current set of shared epistemic toolboxes is insufficient to license a unique inference from E to either C1 or C2. The resolution of this disagreement is a goal the scientific community is committed to *qua* scientific community.

There are a number of ways this resolution could take place. Aisha and Bheem might collect more evidence, and this new evidence might settle the dispute between C1 and C2 without any change in epistemic tools, because now the evidence is sufficient to license a common inference using only shared epistemic tools. They might come to a consensus about a new shared epistemic tool that, when used, settles the dispute to the satisfaction of both scientists in favour of either C1 or C2. Or—and this is the relevant alternative for our argument—they might come to a consensus about a new shared epistemic tool that does not settle the issue in favour of either C1 or C2, but allows them to identify a new conclusion, C3, which is the most informative conclusion that is entailed by both C1 and C2, but can be inferred from E using only shared epistemic tools.

In arriving at the conclusion C3, Aisha and Bheem do not acquire new information about the world. C3 was already a part of both of their belief systems, at least implicitly, since both C1 and C2 entail C3. However, neither C1 nor C2 was a candidate for scientific knowledge, since neither conclusion was inferred using only shared epistemic tools. C3, on the other hand, is a candidate for scientific knowledge. Aisha and Bheem have developed a tool that does not tell either of them anything new about the world, but does tell them which aspects of their understanding of the world could be justified using their shared epistemic toolbox. What we have here is not knowledge-driven consensus, consensus that is a by-product of acquiring new knowledge (or potential knowledge). It is consensus-driven knowledge, new methods of consensus formation that tell us what counts as potential scientific knowledge.

An example of this type of scientific advance might be a case where Aisha and Bheem both infer distinct causal models from a given set of evidence, with the difference in models due to the fact that some of their epistemic tools are not shared. To resolve the disagreement, they come up with a statistical technique they both endorse that allows for an inference from the evidence to a more coarse-grained causal model, one for which both their earlier models are different fine-grained specifications. The coarse-grained model specifies just those aspects of each of their fine-grained models that could be regarded as scientific knowledge, given the current epistemic state of the art. The details from each fine-grained model that are not included in the coarse-grained model are excluded from consideration as present scientific knowledge, at least until some future change in the available epistemic tools.

A real-world example of this type of innovation is the set of tools that have been developed by the Intergovernmental Panel on Climate Change to characterize the nature and extent of consensus in climate science. These tools include new ways to measure and express uncertainty that take into account both quantitative and qualitative variation in beliefs across the community (Betz 2013). They also include multi-model ensemble techniques which involve looking at distributions of projections emerging from multiple different and incompatible climate models, or looking at outcomes that are robust across multiple models, rather than focusing on the individual projections of each model (Winsberg 2018).

Dissent can be productive. Feyerabend famously urged what he called a principle of proliferation which held that scientists should invent and elaborate theories which are inconsistent with consensus views (for discussion see Shaw (2020)). Given the potential for dissent to contribute to scientific progress, one might think that our emphasis on consensus is misguided. To this we have three responses. First, dissent should be 'normatively appropriate' (de Melo-Martín and Intemann 2018; Miller 2021), so not all instances of dissent proliferation are warranted on normative grounds, let alone epistemic grounds. Second, as de Melo-Martín and Intemann (2018) argue, dissent can contribute to public confusion about science, the spread of false beliefs, and public doubt about established scientific findings; properly justified consent, on the other hand, can have the opposite effects. Third, dissent should not be seen as an end in itself, but rather as a means for exposing problems in existing scientific work, the resolution of which would amount to scientific progress and could result in consensus. We agree with philosophers such as Longino who emphasise the importance of criticism in science, and such criticism can promote at least transient dissent. But transient dissent should be seen as a means to the end of justification of scientific claims, and that in turn has a long-run consensus-forming effect. For example, while achieving justified consent about a sun-centred model of the solar system took many decades, once it was achieved it became wildly implausible to dispute; today one needs to hold very extreme and implausible views to deny a sun-centred model of the solar system.

In summary, once we understand that the contours of scientific knowledge are delineated by shared epistemic toolboxes, and that a central aim of science is to develop new epistemic tools to further expand the limits of potential scientific knowledge, it becomes plausible that consensus is a constitutive goal of science, not merely an indicator of epistemic success. The scientific enterprise is not only committed to the resolution of disagreement through the discovery of new evidence, new methods, and new theories, it is also committed to the resolution of disagreement through the identification of shared beliefs within competing perspectives. In the next section, we use this conception of the goals of science to argue for a novel characterization of value-permeation, as well as the thesis that resisting value-permeation is a norm for science.

3. Consensus and Value-Permeation

Strong consensus implies shared epistemic toolboxes. But there is a potential regress lurking here. To have strong consensus about scientific claims, a scientific community must have a consensus about the validity of the epistemic tools used to warrant those claims. But what kind of consensus do we need regarding the tools? If that must be a strong consensus as well, then we will need another shared epistemic toolbox undergirding that consensus, and so on apparently *ad infinitum*.

To avoid either infinite regress or circularity, it cannot be the case that every element of a shared epistemic toolbox must be the object of a strong consensus.

The key here is that strong consensus is required to qualify scientific assertions as fully warranted. However, epistemic toolboxes are not simply collections of declarative claims about what is or is not the case. Epistemic tools are meant to guide scientific inquiry and action, and so the content of the toolbox must be at least partly imperative rather than indicative: "Do this" rather than "This is how things are." Declarative claims may be part of an epistemic toolbox. For instance, part of accepting the validity of the microscope as a tool for acquiring visual evidence of minute objects is accepting the veracity of various optical claims. There needs to be a strong consensus about these declarative claims for there to be a consensus about the validity of the tool. However, accepting the validity of the microscope also involves accepting imperative claims about how to use the microscope to form beliefs, such as claims about how to recognize and prevent artifacts during specimen preparation.

These action-guiding claims do not constitute potential scientific knowledge in the declarative sense, and so do not have to be the subjects of a strong consensus. All we need for such claims is a weak consensus. For instance, in order for me to endorse your use of a microscope as an appropriate tool for acquiring reliable evidence, I need to have good reason to think that the way you identify and eliminate artifacts is truth-conducive. But importantly, I do not have to endorse your reasons for carrying out those procedures, as long as I believe the procedures are reliable. You may have been taught the procedure as a mechanical ritual, and have no real understanding of how it works, but I can still endorse your use of the microscope. There is an asymmetry here. A scientific consensus about what to believe requires members to endorse each other's beliefs as well as each other's reasons for believing. A scientific consensus about how to use tools requires members to endorse each other's use of the tools, but it does not require them to endorse each other's reasons for using the tools as they do.

This is how values can enter the domain of scientific knowledge. As discussed in the previous section, value claims cannot be part of the corpus of declarative scientific knowledge. They are simply not the right sorts of propositions. However, value claims can enter into scientists' reasons for preferring one epistemic toolbox over another. I might, for instance, choose an evidence threshold for acceptance of hypotheses in a particular domain based on value considerations. If our epistemic toolboxes had to be subject to a strong scientific consensus to be shared, this would disallow these value considerations. If there is strong scientific consensus about a claim, then it is the sort of claim that can be scientific knowledge. Value claims are not the sorts of claims that can be scientific knowledge, *ergo* there cannot be strong scientific knowledge about them. But the imperative claims about use of epistemic tools only require a weak consensus. There does not have to be a consensus regarding the reasons for accepting those tools, so values may play a role in those reasons.

Thus we have a preliminary account of value-permeation in science. Values become relevant to scientific knowledge if they are involved in the choice of epistemic toolbox. However, value-permeation of this sort is not a threat to scientific knowledge *per se*. If Aisha and Bheem choose the same epistemic toolbox based on different value considerations, they have a weak consensus on the epistemic toolbox, and this can provide the basis for a strong consensus on some claim inferred using the toolbox. The scientists' decisions are influenced by their values, but the difference in values in this case does not hinder the discovery of scientific knowledge.

However, value differences can become a hindrance to scientific knowledge if they hinder the formation of a strong consensus. If Aisha and Bheem, based on differing value considerations, endorse different toolboxes that lead to different inferences from the same evidence, then their value disagreement is standing in the way of strong consensus, and therefore standing in the way of scientific knowledge formation. It is this kind of case, in which the values make a difference to the inference drawn from the evidence, that is a case of values consequentially permeating science. This permeation of values can be blunted by consensus-forming activity in science, like the development of new shared toolboxes that both Aisha and Bheem could endorse despite their divergent values, thus allowing them to reach a shared conclusion from the evidence that might qualify as scientific knowledge.

Conceptualizing consensus-formation as a goal of science gives us an important new way to think about values in science. This is based on our intertwined understanding of consensus and scientific knowledge. We have an account of how value disagreements influence science by preventing the consensus necessary for knowledge formation. And we have a new motivation for moving towards value-freedom, understood as the quest for consensus formation in the face of value disagreement. In the following sections, we develop this "difference-to-inference" view of value permeation and contrast it with other accounts of value permeation.

4. Characterizing the value-permeation thesis

That the value-free ideal has merit is motivated by the thought that value-free science is more truth-conducive, more trust-enhancing, and more democratic than value-permeated science (see, for example, Bright 2018; Boulicault and Schroeder 2021). Value-permeation threatens those aspects of science that motivate the value-free ideal. However, many philosophers have argued that certain forms of value-permeation are inevitable, and perhaps even beneficial.

Many aspects of science are influenced by value-based considerations, such as the choice of research topics, constraints on research methods, and technological applications of the results of science, but these forms of value-permeation are not philosophically controversial. The most pressing concern about value-permeation is about the influence of values on scientific reasoning, and specifically the inferences drawn from evidence to scientific conclusions.

Yet, how exactly the influence of values on scientific reasoning should be characterized remains somewhat opaque in the literature. Ward notes the diverse ways that philosophers have written about value-permeation (2021): values "lead" scientists to make particular choices (Schroeder 2017) or are "embedded" in those choices (Elliott 2011) or "factor into" those choices (Carrier 2013). This heterogeneity suggests to Ward that contributors to the values-in-science literature are talking past each other. In turn Ward argues that values can serve as either motivating reasons or justifying reasons for scientific choices, and values can be either causal effectors of choices or be causally impacted by choices. We argue that vaue-permeation often involves values as causal effectors in which the choice being impacted is whether accepting or rejecting a particular hypothesis is justified; however, the first view of value-permeation we discuss, the reasoning-to-a-foregone-conclusion view, is best characterized as holding that values are motivating reasons for scientific choices. We construe value-permeation as significant (and, as we will argue, epistemically threatening) when values are difference-makers to inferences. This characterization of value-permeation is more precise and more general than existing characterizations, to which we now turn.

4.1 Reasoning to a foregone conclusion

There is a form of value-permeation that is almost universally acknowledged to be epistemically illegitimate. This is when values lead us to determine the bottom line –

the conclusion—prior to any consideration of the evidence, or to seek or generate evidence which is likely to support the conclusion. The evidence is then used simply to justify the pre-determined conclusion—by, for example, ignoring evidence that would challenge the conclusion and highlighting evidence that would support it. We call this type of value permeation 'reasoning to a foregone conclusion'. To put this in Ward's terms, this view characterizes values as motivating reasons for scientific choices (2021). The literature survey in Peters (2021) shows that this characterisation of value-permeation is widely promulgated.

For example, in her well-known article on the role of values in the study of divorce, Anderson claims that "We need to ensure that value judgments do not operate to drive inquiry to a predetermined conclusion. This is our fundamental criterion for distinguishing legitimate from illegitimate uses of values in science." (2004, p. 11). Similarly, Brown claims that the main concern about value-permeation is that it will "drive inquiry to a predetermined conclusion" (2013, p. 838). Peters further quotes Elliott (2017), Douglas (2009), Douglas (2016), and Longino (1990) as all asserting versions of the reasoning-to-a-foregone-conclusion view as the primary or sole threat associated with value-permeation. Reasoning to a foregone conclusion is certainly not the only way in which values permeate science, but there is a widely held view that it is the only form of value-permeation that is inherently illegitimate.

Consider a historical example of a scientist who plausibly engaged in reasoning to a forgone conclusion. The statistician Ronald Fisher famously disputed the hypothesis that smoking causes lung cancer, and this seems to have been a value-motivated conclusion. Fisher argued that the correlations between smoking and lung cancer observed in case-control studies did not justify the conclusion that smoking causes lung cancer, on the grounds that the correlation was probably a result of a common cause, likely genetic, between smoking and lung cancer (1958). Stolley suggests that Fisher's position was deeply value-permeated (1991): Fisher's political conservatism contributed to ideological objections to public health campaigns; Fisher felt that genetics had not been sufficiently considered in medicine in general; the importance of Fisher's own methodological innovation, the randomized trial, was threatened by causal inference from a non-randomized study; Fisher was a smoker; and Fisher had financial ties to tobacco companies. Fisher's claim that smoking does not cause lung cancer appears to have been for him a foregone conclusion.

We do not dispute that reasoning to foregone conclusions is generally unreliable. However, we disagree with the view that this form of value-permeation is the only variety that is illegitimate *per se*. Moreover, we argue that there are cases where reasoning to a foregone conclusion is plausibly not an epistemic error.

Value-permeated reasoning to foregone conclusions is generally not a truth-conducive or trust-inspiring way to do science. Yet, this concern cannot fully account for the threat of value-permeation. A scientist could have a special talent at discovering true facts and could have a value-laden commitment to providing justification to claims about those facts, and thus engage in value-permeated reasoning to a foregone conclusion about those claims, yet thanks to her special talent this would generally be reliable. From an epistemic perspective there appears to be no problem here, despite the gut reaction that one might have about scientists engaged in value-permeated reasoning to foregone conclusions.

Peters (2021) argues that value-permeation which causes scientists to infer to foregone conclusions can in fact contribute to reliable science at the group level, via Mandevillian effects (epistemically imperfect processes at the individual level which have group-level epistemic benefits). If this is correct, then the reasoning-to-a-foregone-conclusion characterization of the threat of value-permeation is at best incomplete.

Moreover, a scientist could engage in reasoning to a value-laden foregone conclusion which is a universally accepted scientific claim, in which case the fact that this pattern of reasoning was value-laden would not make a difference to the scientific community or to decision-makers depending on the results of those sciences, particularly if that reasoning involved a shared epistemic toolbox. Similarly, a scientist could engage in reasoning to a value-laden foregone conclusion in which the permeating values are universally shared, and thus any other scientist who shared this scientist's evidence and epistemic toolbox would reach the same conclusion—this scientist would be in a strong consensus with her peers. Both considerations suggest that, in principle at least, value-motivated reasoning to a foregone conclusion does not always make a difference to the warranted inference, and thus does not hamper scientific consensus nor threaten the status of a conclusion as scientific knowledge.

Most importantly, value-permeation is not limited to instances of reasoning to foregone conclusions. Value-permeation which operates in more subtle ways, such as those described in the following sections, are pervasive in scientific practice. For example, consider the phenomenon of publication bias in medical research, in which evidence that supports a company's pharmaceutical is published while evidence that suggests this pharmaceutical is not safe or effective does not get published. It is difficult to see how publication bias necessarily involves reasoning to a foregone conclusion, since it could result from a variety of other factors, though the net effect of the various decisions that contribute to publication bias amount to value-permeation influencing the inferences of relevant stakeholders. This mechanism of value-permeation is widespread, yet it cannot be generally accommodated by the reasoning-to-a-foregone-conclusion view of value-permeation.

4.2 The underdetermination view

The underdetermination view of the threat of value-permeation is probably the most prevalent in the literature on values in science. The general idea is that evidence necessarily underdetermines the truth of hypotheses, yet often we must hold a dichotomous doxastic attitude towards hypotheses, particularly when we are required to make decisions based on a hypothesis; our assessment of the practical consequences of a mistaken belief about a hypothesis is based on our values, and that in turn modulates our doxastic attitude. Thus, values have an influence on scientific inference, though this influence is not as direct as is the influence of values on the reasoning-to-a-foregone-conclusion view (§2.1).

A specific and important version of this argument is the inductive risk argument, which was articulated by Rudner (1953) and resuscitated by Douglas (2000). Empirical methods, modelling, and statistical inferences involve trading off the probability of committing false positive errors (accepting a hypothesis as true when it is false) with the probability of committing false negative errors (rejecting a hypothesis as false when it is true). False positives and false negatives have different practical consequences, and our evaluation of those consequences is of course based on values, and—goes this argument—such evaluation must and should influence choices about the empirical methods, modelling, and statistical inferences. Those choices, in turn, influence our inference about the hypothesis in question. (For a recent characterization of this argument, see Havstad 2022). Thus, scientific reasoning itself is value-permeated.

The inductive risk argument is compelling and has been influential, but it is also incomplete. For some scientific hypothesis, values could modulate the acceptable false positive and false negative probabilities for all relevant scientists, decision-makers, or people in general in the same way, such that all people either believe or disbelieve that hypothesis. That, in turn, would entail that values would not threaten the potential of a strong consensus about the hypothesis and thus the status of that hypothesis as scientific knowledge. For many real scientific scenarios that may sound fanciful, but the mere possibility of such a scenario suggests that something is missing in the underdetermination characterization of value-permeation. What is missing is the role of values making a difference to inference (§4.4).

When values make a difference to inference, they block the ability of science to achieve consensus. As argued in §2 and §3, consensus is important to science, and this difference-making consequence of value-permeation threatens the potential for strong consensus about a hypothesis. We thus disagree with Douglas, who claims that when values influence science indirectly, in the sense described in this section, they can "completely saturate science, without threat to the integrity of science" (2009, p. 96). On the contrary, we argue that the integrity of science is threatened when values function as difference-makers for inference, because they block the potential for consensus and thus the status of a hypothesis as scientific knowledge.

4.3 The mixed claims view

Alexandrova (2018) introduces the notion of 'mixed claims' and argues that mixed claim hypotheses constitute a form of value-permeation. A mixed claim is a scientific hypothesis about a causal or statistical relation in which at least one of the variables in the hypothesis "is defined in a way that presupposes a moral, prudential, political, or aesthetic value judgement about the nature of this variable" (Alexandrova 2018, p. 424). She claims that they "present a special case of value-ladenness" (p. 421), and so reasoning about mixed claims should be understand as a form of value-permeation. Yet, Alexandrova argues convincingly that mixed claims do not generate value-permeation of the same kind as that usually discussed in the literature (such as in §4.1 and §4.2). Value-permeation, on this account, is a result of mixed claims 'presupposing' value judgements, in which case a conclusion about the truth or falsity of a mixed claim somehow inherits this value judgement.

Yet, one might respond to Alexandrova by maintaining that scientists need not presuppose value judgements about the relata in a mixed claim. Of course, scientists often have such presuppositions; the issue is whether those presuppositions are necessary for scientific work and particularly for scientific inference, and whether those presuppositions thereby necessarily permeate scientific reasoning itself. We think not. You could be an empirical scientist studying the rates of poverty in various communities, for example, but have no presupposition whatsoever about whether an absolute or a relative conception of poverty is superior. Or you might believe that in developed countries a relative measure of poverty is generally more appropriate than an absolute measure of poverty—a value-laden judgement—but you might anyway think it worthwhile to estimate the frequency of absolute poverty in Saint Louis—and while the pursuit of this question is of course value-motivated, the empirical work and subsequent inferences would not be any further value-permeated than they otherwise would have been due to inductive risk considerations (see Peters 2020 for a similar argument).

Consider a case about a controversial disease in medicine: female sexual dysfunction ('female sexual interest/arousal disorder' in the Diagnostic and Statistical Manual of Medical Disorders). Maria might think, contrary to the status quo, that female sexual dysfunction is not a real disease but rather is a result of fatigue or stress or bad relationships or capitalism or boredom due to monogamy, and its inclusion in disease manuals is a value-motivated ploy by pharmaceutical companies to sell more pills, but if Maria operationalizes the condition in a particular way (say, with a standardized sexual satisfaction questionnaire), then Maria can perform a randomized trial to test if some intervention can increase the average score on that questionnaire in the drug group of the trial compared with a control group—and though a hypothesis about the intervention for the condition is a mixed claim, any conclusion about the capacity of the intervention need not be permeated by Maria's misgivings about the genuine nature of the condition or the value-ladenness of the relata of the hypothesis about the intervention.

Mixed claims make a difference to inference if the following conditions hold:

(i) Multiple scientists make an inference about the same mixed-claim hypothesis
(ii) The value-presuppositions in that mixed claim vary among the scientists
(iii) The varying value-presuppositions entail that some scientists justifiably infer the truth of the hypothesis while others justifiably infer its falsity

Note that (ii) and (iii), while plausible in some cases, could be false in many other cases. For example, the poverty rate in Sweden is less than the poverty rate in Russia on both an absolute and a relative conception of poverty, and thus (iii) would not hold when comparing the poverty rates of Sweden and Russia. More importantly, though, is (i). Suppose two economists are comparing poverty rates between England and France, and one, who makes her measurements according to her normatively-preferred absolute conception of poverty, concludes that France has a higher poverty rate than England, but the other, who makes her measurements according to her normatively-preferred relative conception of poverty, concludes that England has a higher poverty rate than France. Now (ii) and (iii) are satisfied. Yet, (i) is not satisfied, since the two scientists are making inferences about different hypotheses—one about absolute poverty rates between the countries and the other about relative poverty rates between the countries. It would not be enough for the scientists to emphatically declare that they are both measuring the same thing—the poverty of England and France simpliciter—since they would be wrong, and we are interested here not in the idiosyncratic declarations of scientists but rather in their justified inferences. So, in general, the phenomenon of mixed claims does not entail a difference to inferences resulting from the value-permeation of those claims.

4.4 The difference-to-inference view

We have been alluding to the centrality of value-permeation making a difference to inferences. It is when values make a difference to the inferences made based on the available evidence that value-permeation creates problems. This is the *difference-to-inference* model of value-permeation. When values make a difference to inference, consensus about the hypothesis in question is threatened, and in turn the potential for that hypothesis to be scientific knowledge.

Recall that in §2 we argued that scientific consensus and scientific knowledge are intricately interwoven. The difference-to-inference view of value-permeation can be given both a consensus-first argument and a knowledge-first argument. Let us consider them in turn. In an as-yet unpublished manuscript we offer a specific characterization of values making a difference to inference. There we introduce the notion of a 'bifurcation point' to explicate the difference-to-inference view of value-permeation. A bifurcation point exists when two individuals share some evidence E, hold different value sets V_1 and V_2, and:

(i) E & V_1 justify conclusion C_1
(ii) E & V_2 justify conclusion C_2
(iii) C_1 and C_2 are inconsistent

It is in this precise sense that values influence scientific reasoning. These three conditions represent value-permeation, because when conditions (i) and (ii) are met, values make a difference to the inferences pertaining to C_1 and C_2, and when condition (iii) is met, the possibility of consensus is blocked. In §2 we argued that scientific knowledge is consensus-based, and thus, when conditions (i) – (iii) are met for some hypothesis, that hypothesis cannot be deemed scientific knowledge. This is the consensus-first argument.

The knowledge-first argument is straightforward. Evidence is truth-relevant, but values are not—values pertain to how one wishes the world were and clearly do not indicate how the world is. A plausible desideratum for scientific inference is the *difference-to-inference* principle, which says: only truth-relevant properties should be difference-makers for inference. Value-permeation implies that our inferences are sensitive to values, despite the fact that values are not truth-relevant. If values make a difference to inference, and values are not truth-relevant, then the difference-to-inference principle is violated. That is precisely what is wrong, epistemically, with value-permeation. In turn, if a set of scientists and stakeholders do not share the relevant values permeating inferences about a hypothesis, they will not form a consensus about that hypothesis, and that hypothesis would not get the status of scientific knowledge.

To be clear, we are not suggesting that the difference-to-inference principle should never be violated. As with many principles, there are times when the difference-to-inference principle *must* be violated: all compelling instances of the inductive risk argument involve violations of the difference-to-inference principle. Particularly when science must be relied on for guiding decisions, the inductive risk argument is sound, and so violations of the difference-to-inference principle are warranted. The point is only to emphasise the threat of value-permeation when values make a difference for inference.

The difference-to-inference characterization of value-permeation accommodates compelling examples of value-permeation in the values in science literature. For instance, the case of health harms of dioxins presented in Douglas (2000) exemplifies our account, as the case shows that “two opposing and plausible interpretations” (p. 576) of the evidence were available, depending on background assumptions and different evaluations of the consequences of error. More generally, the fact that different conclusions are justified in (i) and (ii) need not be due to the value sets V_1 and V_2 playing a direct confirmatory role in the sense discussed in §4.1, but rather is best thought of in terms of the modulation of credential thresholds to accept or reject hypotheses in the sense discussed in §4.2.

5. Responding to value-permeation

Philosophers have offered a variety of strategies to respond to the threat of value-permeation. In this section we articulate some of the most prominent responses and argue that the response that corresponds to our difference-to-inference model of value-permeation is superior. In short, we argue that scientists should strive to eliminate all values which are difference-makers for inference. While there is substantial overlap in the following approaches, we argue that what we take to be the most compelling response to the threat of value-permeation—the difference-to-inference elimination approach—provides the basis for a novel version of the value-free ideal.

5.1 Forbid reasoning to a foregone conclusion

We saw that many philosophers characterize the threat of value-permeation in terms of reasoning-to-a-foregone-conclusion. The idea is that values threaten the legitimacy of scientific reasoning when they directly influence inferences. This characterization of the threat of value-permeation in turn suggests a natural prescription: “One cannot use values to direct the selection of a problem and a formulation of a methodology that in combination predetermines (or substantially restricts) the outcome of a study.” (Douglas 2009, p. 100). That is, values should not be permitted to serve as “reasons in themselves” for or against a hypothesis (Douglas 2009 p. 96). Importantly, this response is articulated by those who appeal to the argument from inductive risk to conclude that values can and should have some influence over scientific reasoning. For that to be tenable, such philosophers distinguish between the direct and indirect influence of values on scientific inference.

We saw in §4.1 reasons to doubt the completeness of the reasoning-to-a-foregone-conclusion view of the threat of value permeation. Similar reasons entail that the corresponding prescription for science is incomplete. For instance, reasoning to a *true* foregone conclusion, though not generally a meritorious procedure, has a meritorious outcome.

Moreover, we saw in §4.1 that the threat of value permeation is not limited to instances of reasoning to foregone conclusions and indeed, cases of reasoning to foregone conclusions very likely make up a tiny proportion of instances in which value-permeation threatens the integrity of science. Thus, the prohibition on such reasoning is at best incomplete. Relying only on this prescription as a means of managing value-permeation would permit widespread value-permeation.

5.2 Promote transparency about value-permeation

A common response to value-permeation is that scientists ought to be transparent about the role of values in scientific reasoning (Elliott 2022). The rationale is straightforward: transparency can help to inform other scientists, decision-makers, and the public about how a conclusion was reached based on the available evidence. Transparency can be achieved in a variety of ways; one strategy that we think is especially valuable is the inclusion of a diverse range of scientists and stakeholders in the scientific process to help uncover the influence of existing values and biases (Longino 1990).

Transparency is fine as far as it goes, but it is limited in several ways. Transparency about value-permeation does not necessarily (arguably not typically) lead to consensus. Moreover, transparency about value-permeation may contribute to a decline in public trust in science, particularly when the public holds a naïve conception about the inner-workings of science (John 2018).

A compelling criticism of the feasibility of the transparency response to value-permeation is offered by Winsberg, who argues that the influence of values does not have to be so coarse-grained and evident as the views of value-permeation described in §4.1 and §4.2 seem to suppose, but rather, value-permeation can be subtle, occurring in the 'nooks and crannies' of scientific research (2012). Winsberg's example was complex climate models. When climate scientists make predictions about future climate change, those predictions are necessarily value-permeated, because the predictions are based on models which are developed on the basis of numerous unforced choices, and which are liable to the inductive risk argument. Given the complexity of the models and the extent of distributed epistemic labour, and the sheer number of free choices in the modelling work, transparency about the role of values in the scientific work is simply impossible. Kukla (2019) similarly argues that values can be "buried too deep" for transparency to afford critical engagement with value-permeation. Thus, even if one agreed that transparency about value-permeation is good practice, it may not always be achievable.

5.3 Permit only the right values

If scientific reasoning is value-permeated, then a natural response is to require that the permeating values be justified on some basis, such that only the legitimate values end up permeating scientific reasoning. Holman and Wilholt call this the 'axiological' strategy for managing value-permeation (2022). There have been a variety of such proposals.

For example, Kourany argues that value-permeation should involve only values that are "socially responsible" and that promote "human flourishing" (2010 p. 68). Similarly, Intemann argues that value-permeation is legitimate if it promotes "democratically endorsed epistemological and social aims of the research" (2015). However, the executive of a company that manufactures a consumer product which includes a potential carcinogen may have a different view about what is required for human flourishing or what the social aims of research on that potential carcinogen should be compared with the public health official tasked with evaluating the putative health harms of that product. On the other hand, the right-values approach need not precisely articulate from the outset whose values or what values should be permitted in value-permeation, but rather, such an approach could be procedural, specifying a method which states how values are to be permitted. Kitcher's notion of 'well-ordered science', based on an ideal deliberative mechanism, is such a view (2001). Yet Holman and Wilholt (2022) argue that in a society which holds a range of contradictory values about some domain, the well-ordered science procedure would fail to achieve consensus (a point convincingly made by Larroulet Philippi 2020).

Boulicault and Schroeder (2021) argue for what they call the 'idiosyncrasy-free ideal', which holds that scientific decisions should not be dependent on the influence of features of individual scientists, including those scientists' values; the implication is that non-idiosyncratic values could be permitted. They cite Wilholt (2013) as offering a detailed version of this: Wilholt argues that shared conventions, such as 'only accept a hypothesis with a p value of less than 0.05', can block the influence of values on the inferences of individual scientists. Yet, there cannot be conventions for all the unforced decisions that scientists must make, especially in the nooks and crannies of scientific work (Winsberg 2012). Moreover, scientific conventions can be epistemically impoverished; the p-value convention for hypothesis testing is arguably a good example of a flawed scientific convention, and for another example, Stegenga (2017) argues that the conventions for pharmaceutical regulation are unreliable guides to the benefits and harms of pharmaceuticals.

5.4 Constrain value-permeation

Given the threat of value-permeation, in conjunction with the arguments that conclude that value-permeation is necessary and in specific ways good for science, a natural response is to maintain that science should strive to constrain or decrease value-permeation while not aiming to altogether eliminate it. We have seen that Douglas seems to have such a view. In the context of articulating an "ideal for scientific reasoning" (2009 p. 87), Douglas argues that values have an ineliminable role in scientific reasoning, yet she argues that there are limits to the ways in which values should be allowed to permeate science: her thesis is that while values ought not play a 'direct' role in hypothesis choice, in the sense of the reasoning-to-a-foregone-conclusion view (§4.1), values can and should play an 'indirect' role in the sense of influencing scientists' acceptable risks of error when accepting hypotheses (§4.2), and in that capacity values can and should "saturate science" (2009 p. 96). This is a prominent articulation of a constraint response to value-permeation yet allows for and indeed encourages rampant value-permeation.

Elliott defends another version of the constraint response to value-permeation (2017). He suggests that one way to minimize the influence of value-permeation on the reasoning of individual scientists is to develop community guidelines that specify how scientists should act when faced with unforced choices and uncertainty. Community-centred strategies can be effective in responding to value-permeation, and thus we have some sympathy for the broad contours of community-based responses to value-permeation. Yet Elliott's specific proposal has problems. Elliott's proposal looks very similar to that of Wilholt (2013), mentioned in §5.3, and so inherits the same problems: community guidelines can be unreliable, and guidelines cannot be made for the vast number of fine-grained choices that scientists are faced with in the nooks and crannies of their work. Moreover, as Elliott himself notes, this proposal "would merely shift the value judgements from the individual scientists to the community that creates the guidelines" (p. 98). Finally, this shift of value-permeation from an individual scientist to a community of scientists would correspondingly shift the threat of value-permeated science not being trust-inducing or democratic with respect to the broader public.

There is a bigger problem lurking for this view. Suppose we grant the lesson from §4.2 that, based on the inductive risk argument, values can have a positive influence on scientific reasoning. Any response to value-permeation which holds that we should constrain or minimize value-permeation to the point that the threat of value-permeation is minimized, and the positive influence of values is optimal, assumes that there is an ideal degree of value-permeation which scientists should aim for. Yet, discerning that one has achieved that ideal degree of value-permeation is virtually impossible. Moreover, what exactly that ideal range consists of is sensitive to contingent changes in science, technology, and methodology.

To illustrate this last point, consider again the example of Fisher's value-motivated denial that smoking causes lung cancer. This may have been a tenable position in the 1950s when Fisher published a letter in *Nature* arguing that due to the possibility of a confounding factor in the case-control studies that showed a correlation between smoking and lung cancer, we ought not conclude that smoking causes lung cancer (1958). Fisher alluded to twin studies which showed that smoking habits of monozygotic twins were more similar than that of dizygotic twins, suggesting a genetic basis of smoking, and he concluded that his disputants who believed smoking causes lung cancer were themselves value-motivated, calling their proclamations "propaganda" ("there is nothing to stop those who greatly desire it from believing that lung cancer is caused by smoking cigarettes.") In response, Bradford Hill argued that there is a variety of features of evidence which can provide warrant to causal conclusions, including the observed strength of the effect size, the reproducibility of the findings, the biological plausibility of the causal relationship, and several others (1965)—these are sometimes called 'Hill's criteria'. By appealing to these considerations, it became much less plausible to deny that smoking causes lung cancer, and this in turn led to a stronger and better-justified consensus about the matter. Bradford Hill's methodological innovation effectively blocked one way in which values could influence reasoning about this hypothesis, and in so doing he shifted the ideal degree of acceptable or optimal value-permeation (for a discussion of Hill's criteria, see Bird 2011).

5.5 Eliminate differences to inference

The Fisher-Bradford Hill debate suggests an alternative response to value-permeation. Imagine a contemporary of these disputants—let us call her Ronalda Pescadora—who, like Fisher, has made a value-permeated inference that smoking does not cause lung cancer, but unlike Fisher, holds her permeating values less sternly, and has a nagging doubt that she might be wrong. Bradford Hill then develops his novel way of assessing causal relationships, and, drawing on a range of relevant empirical findings, builds an argument that smoking causes lung cancer, which subsequently convinces Ronalda. Her values did not change, but the empirical context and available methodological apparatus did change. Prior to Bradford Hill's innovation, Bradford Hill and Ronalda disagreed about the hypothesis in question, but after Bradford Hill's innovation they agreed. The difference between Ronalda's values and Bradford Hill's values made a difference to their inference prior to Bradford Hill's innovation, but that difference-maker was eliminated by Bradford Hill's innovation.

In the terms introduced in §4.4, Bradford Hill's innovation eliminated a bifurcation point between Bradford Hill and Ronalda. This, in turn, created a strong consensus between the two scientists. This suggests a general response to value-permeation. When values permeate scientific reasoning, scientists should strive to eliminate that value-permeation. In the precise sense characterized above, if a bifurcation point exists between two individuals, scientists should strive to eliminate that bifurcation point.

In our presently-unpublished companion paper on the value-free ideal, we describe various strategies for how scientists can eliminate bifurcation points. These strategies involve routine aspects of science, such as strengthening the quality of available evidence and hedging one's conclusions. Parker (2014) offers an example of hedging. Parker notes that if coarser rather than finer estimates of future climate states are being considered, then the influence of values "will be reduced insofar as choices in model development will less often *make a difference* to the uncertainty estimates produced" (28). A coarse-graining of a hypothesis can eliminate the difference that values make to inferences about that hypothesis. John discusses another strategy for eliminating bifurcation points which is based on the strengthening of evidence (2015).

Our view has several advantages over those sketched in §5.1 to §5.4. First, many of the above strategies in the other responses to value-permeation are also strategies to eliminate differences to inference, yet our view is immune to the problems of the other responses. Consider the prohibition against reasoning-to-a-foregone-conclusion: insofar as such reasoning introduces bifurcation points with those who do not share the same values motivating such reasoning, the imperative to not reason to this foregone conclusion involves an elimination of the corresponding bifurcation point; yet if such reasoning did not introduce a bifurcation point (say, because the hypothesis being reasoned to was believed by everyone), then an imperative to not reason that way would seem excessively strict. Second, a constitutive aim of science is the discovery of truth (notwithstanding antirealist cautions), and if an inference is modulated by features that are not truth-relevant, that inference is less truth-conducive than it otherwise could have been. So, the capacity of those features that are not truth-relevant to be difference-makers for inference should be eliminated. And that is precisely what our view states. Third, the elimination of bifurcation points contributes to consensus-formation, which, as we argued in §2, is interwoven with the development of scientific knowledge. Fourth, truth aside, the results of science should contribute to social coordination, particularly in sciences that are meant to provide guidance for policy. The existence of a bifurcation point entails that two individuals differ regarding their view about some putative fact, so any decision based on that putative fact will seem right or just to only one of those individuals. Were that bifurcation point eliminated, then both individuals would have the same view regarding the resulting decision or policy. Value-permeation of science drives controversy in important domains such as the study of climate change or vaccine safety (Elliott 2017). Our prescription can help to alleviate such controversies.

This view about how to respond to value permeation constitutes a novel version of the value-free ideal for science, one which we describe and defend in the as-yet unpublished manuscript mentioned earlier. That value-free ideal says: scientists should act as if science should be value-free. Our value-free ideal departs significantly from the existing ideal defended by Betz (2013) and others and allows us to respect the arguments challenging the classical value-free ideal, particularly from the inductive risk argument. When value-permeation is characterized in the way we have done so here, our new value-free ideal says: scientists should strive to eliminate bifurcation points, or differences to inference.

References

Alexandrova, Anna. 2018. "Can the Science of Well-Being Be Objective?" *British Journal for the Philosophy of Science* 69 (2): 421-445.

Anderson, Elizabeth. 2004. "Uses of Value Judgments in Science: A General Argument, with Lessons from a Case Study of Feminist Research on Divorce." *Hypatia* 19 (1): 1–24.

Betz, Gregor. 2013. "In Defence of the Value Free Ideal." *European Journal for Philosophy of Science* 3: 207-220.

Bird, Alexander. 2022. *Knowing Science*. Oxford University Press.

Bird, Alexander. 2011. "The Epistemological Function of Hill's Criteria." *Preventive Medicine* 53(4-5): 242-245.

Bradford Hill, Austin. 1965. "The Environment and Disease: Association or Causation?" *Proceedings of the Royal Society of* Medicine 58 (5): 295–300.

Brown, Matthew. 2013. "Values in Science beyond Underdetermination and Inductive Risk." *Philosophy of Science* 80: 829-839.

Boulicault, Marion and Schroeder, Andrew. 2021. "Public Trust in Science: Exploring the Idiosyncrasy-Free Ideal." In *Social Trust* edited by Kevin Vallier and Michael Weber, Routledge.

Carrier, Martin. 2013. "Values and Objectivity in Science: Value-Ladenness, Pluralism and the Epistemic Attitude." *Science & Education* 22 (10): 2547–68.

Dellsén, Finnur. 2016. "Scientific Progress: Knowledge versus Understanding." *Studies in History and Philosophy of Science* 56: 72-83.

de Melo-Martín, Inmaculada and Kristen Intemann. 2018. *The Fight against Doubt: How to Bridge the Gap between Scientists and the Public.* Oxford: OUP.

Douglas, Heather. 2000. "Inductive Risk and Values in Science." *Philosophy of Science* 67 (4): 559-579.

Douglas, Heather. 2009. *Science, Policy, and the Value-Free Ideal*. Pittsburgh: University of Pittsburgh Press.

Douglas, Heather. 2016. "Values in Science." In *Oxford Handbook of Philosophy of Science* edited by Paul Humphreys. 609-31 New York: Oxford University Press,.

Elliott, Kevin. 2011. *Is a Little Pollution Good for You?: Incorporating Societal Values in Environmental Research*. New York: Oxford University Press.

Elliott, Kevin. 2017. *A Tapestry of Values: An Introduction to Values in Science*. New York: Oxford University Press.

Elliott, Kevin. 2022. "A Taxonomy of Transparency in Science." *The Canadian Journal of Philosophy* 52(3): 342-355.

Feyerabend, Paul. 1965. "Reply to Criticism: Comments on Smart, Sellars and Putnam." *Proceedings of the Boston Colloquium for the Philosophy of Science*, 223-61.

Fisher, Ronald. 1958. "Cancer and Smoking" *Nature* 182: 596.

Gerken, Mikkel. 2022. *Scientific Testimony: Its Roles in Science and Society.* Oxford University Press.

Habermas, Jürgen. 1996. *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy*, W. Rehg (trans.), Cambridge, MA: MIT Press.

Havstad, Joyce. 2022. "Sensational Science, Archaic Hominin Genetics, and Amplified Inductive Risk" *Canadian Journal of Philosophy* 52(3): 295-320.

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

Intemann, Kristen. 2015. "Distinguishing Between Legitimate and Illegitimate Values in Climate Modeling." *European Journal for Philosophy of Science* 5: 217-232.

John, Stephen. 2015. "The Example of the IPCC Does Not Vindicate the Value Free Ideal: A Response to Gregor Betz." *European Journal for Philosophy of Science* 5:1-13.

John, Stephen. 2018. "Epistemic Trust and the Ethics of Science Communication: Against Transparency, Openness, Sincerity and Honesty." *Social Epistemology* 32(2): 75-87.

Kitcher, Philip. 2001. *Science, Truth, and Democracy.* Oxford: Oxford University Press.

Kourany, Janet. 2010. *Philosophy of Science After Feminism.* New York: Oxford University Press.

Kukla, Rebecca. 2019. "Infertility, Epistemic Risk, and Disease Definitions." *Synthese* 196: 4409-4428.

Larroulet Philippi, Cristian. 2020. "Well-Ordered Science's Basic Problem." *Philosophy of Science* 87(2): 365-375.

Lehrer, Keith and Carl Wagner. 1981. *Rational Consensus in Science and Society*. D. Reidel Publishig Company.

Longino, Helen. 1990. *Science as Social Knowledge.* Princeton University Press.

Merton, Robert K. 1942 [1973]. "The Normative Structure of Science." In *The Sociology of Science: Theoretical and Empirical Investigations* Chicago: University of Chicago Press, pp. 267–278.

Miller, Boaz. 2013. "When is Consensus Knowledge Based? Distinguishing Shared Knowledge from Mere Agreement." *Synthese* 190: 1293-1316.

Miller, Boaz. 2021. "When is Scientific Dissent Epistemically Inappropriate?" *Philosophy of Science* 88(5): 918-928.

Niiniluoto, Ilkka. 2014. "Scientific Progress as Increasing Verisimilitude." *Studies in History and Philosophy of Science Part A* 46: 73–77.

Oreskes, Naomi. 2004. "The Scientific Consensus on Climate Change." *Science* 306: 1686.

Parker, Wendy. 2014. "Values and Uncertainties in Climate Prediction, Revisited." *Studies in History and Philosophy of Science Part A* 46: 24–30.

Peters, Uwe. 2020. "Values in Science: Assessing the Case for Mixed Claims" *Inquiry* DOI: 10.1080/0020174X.2020.1712235

Peters, Uwe. 2021. "Illegitimate Values, Confirmation Bias, and Mandevillian Cognition in Science." *The British Journal for the Philosophy of Science* 72(4): 1061-1081.

Rudner, Richard. 1953. "The Scientist qua Scientist Makes Value Judgments." *Philosophy of Science* 20(1): 1–6.

Schroeder, Andrew. 2017. "Using Democratic Values in Science: An Objection and (Partial) Response." *Philosophy of Science* 84 (5): 1044–54.

Shaw, Jamie. 2020. "Feyerabend and Manufactured Disagreement: Reflections on Expertise, Consensus, and Science Policy." *Synthese* 198(25): 6053-6084.

Stegenga, Jacob. 2016. "Three Criteria for Consensus Conferences." *Foundations of Science* 21: 35-49.

Stegenga, Jacob. 2017. "Drug Regulation and the Inductive Risk Calculus." In *Exploring Inductive Risk,* edited by Kevin Elliott and Ted Richards). Oxford University Press

Stolley, Paul. 1991. "When Genius Errs: R. A. Fisher and the Lung Cancer Controversy." *American Journal of Epidemiology* 133(5):416–425.

van Fraassen, Bas. 1980. *The Scientific Image*. Oxford University Press.

Vickers, Peter. 2023. *Identifying Future-Proof Science.* Oxford University Press.

Ward, Zina. 2021. "On Value-Laden Science." *Studies in History and Philosophy of Science Part A* 85: 54-62.

Wilholt, Torsten. 2013. "Epistemic Trust in Science." *The British Journal for the Philosophy of Science* 64: 233-53.

Winsberg, Eric. 2012. "Values and Uncertainties in the Predictions of Global Climate Models." *Kennedy Institute of Ethics Journal* 22: 111-137.

Winsberg, Eric. 2018. *Philosophy and Climate Science*. Cambridge University Press.