

WELL-ORDERED SCIENCE AND PUBLIC TRUST IN SCIENCE

Gürol Irzık and Faik Kurtulmus

Forthcoming in *Synthese*

Abstract Building, restoring and maintaining well-placed trust between scientists and the public is a difficult yet crucial social task requiring the successful cooperation of various social actors and institutions. Philip Kitcher's (2011) takes up this challenge in the context of liberal democratic societies by extending his ideal model of "well-ordered science" that he had originally formulated in his (2001). However, Kitcher nowhere offers an explicit account of what it means for the public to invest epistemic trust in science. Yet in order to understand how his extended model and its implementation in the actual world address the problem of trust as well as to evaluate it critically, an explicit account of epistemic public trust in science needs to be given first. In this article we first present such an account and then scrutinize his project of building public trust in science in light of it. We argue that even though Kitcher's ideal model and his proposals for its implementation in the real world face a number of problems, they can be addressed with the resources of our account.

Key Words Public trust . well-ordered science . inductive risk . diversity

Scientific expertise has a crucial role to play in dealing with many of the challenges our societies face from health policy to addressing climate change. Yet, liberal democratic societies are thought to experience significant distrust of expertise (Drezner 2017; Gauchat 2012; Nichols 2017). This has placed questions of trust in science firmly on the agenda not only of scientists (ALLEA 2018; Resnick, Sawyer, and Huddleston 2015), but also philosophers of science (Douglas 2009; Grasswick 2010; Kitcher 2011; Wilholt 2013). The latter emphasize that science is in the pursuit not of any truths but significant truths and that the acceptance of hypotheses involves trade-offs regarding inductive risk, both of which are influenced by social and moral values. Granting the role of such values in science raises the question of "whose values"? If the values guiding science are not shared by the public,

the public's trust may not be forthcoming. How can we then incorporate people's values in science to build trust? What exactly does public's investing trust in science involve? The aim of this article is to pursue these questions largely, though not exclusively, within the framework of "well-ordered science" developed by Philip Kitcher.

Kitcher has long been arguing that if science is to promote the collective good, it should pursue those ends (from agenda setting to their applications) that would be agreed upon through a process of hypothetical democratic deliberation by ideal deliberators tutored by scientists. The hypothetical nature of Kitcher's model of well-ordered science has been the subject of sustained criticism. A number of philosophers have forcefully argued that since we are not in a position to know its outcome, it cannot generate well-grounded substantive recommendations (Brown 2013; Douglas 2013). Others have rightly pointed out that arguments for actual democratic deliberation do not carry over to ideal democratic deliberation (Keren 2013). Finally, Kitcher has also been criticized for not offering any suggestions as to how his model can be applied in practice (Hausman 2003). While we concede the force of these criticisms, we see merit in further engagement with his work for two reasons. First, in (Kitcher 2011) he does offer some practical proposals for approximating his ideal through Fishkin style deliberative polling. Although this goes some way towards addressing the criticisms above, critics have for the most part ignored his proposals.¹ Second, Kitcher (2011) contains a number of significant revisions to his original model developed in (Kitcher 2001), and they are foremost aimed at rebuilding public trust. To our knowledge, how his revised model, when combined with his practical proposals, tackles the problem of public trust in science has not received the attention it deserves. Hence, engaging with Kitcher's views for building and maintaining public trust in science is, we believe, instructive in its own right.

Even though epistemic public trust in science occupies a central place in his (2011), Kitcher's conception of it remains implicit throughout. Yet in order to understand how his revised model and its implementation in the actual world address the problem of trust as well as to evaluate it critically, an explicit account of epistemic public trust in science needs to be given first. Accordingly, we present such an account in Sect. 1. In Sect. 2 we provide our interpretation of Kitcher's suggestions for building public trust in science in light of our account. In Sect. 3 we evaluate Kitcher's model and his proposals

¹ Notable exceptions are (Brown 2013 and Wilholt 2014).

for its implementation. We argue that they face a number of problems and suggest how they can be addressed with the resources of our account.

1 An Account of Epistemic Public Trust in Science

“Public trust in science” is an ambiguous expression as “public” and “science” can refer to different things. Moreover, the object of trust is usually left implicit. First, “public” can refer to collective bodies acting on behalf of the public, such as the courts or ministries, or it can refer to individuals making up the public. Second, “science” can refer to scientific methods, to individual scientists or groups of scientists, or to their social organization. Third, the object of trust can be the information provided by scientists, or it can be the consequences of scientific activities more broadly conceived.²

In this paper we shall assume that members of the public have a certain level of trust in science as a whole as a way of obtaining knowledge about the world. Drawing on our previous work (Irzik and Kurtulmus, forthcoming) we will present in this section some necessary conditions for the public’s placing warranted epistemic trust in a scientist or a group of scientists as providers of information. For this reason, our account does not apply to those who reject *tout court* scientific methods and practices.³ Even if trust in scientific methods and the organization of science is assumed, the question of trust in individual scientists or groups of scientists remains important. We cannot assume that scientists have always employed scientific methods and procedures correctly. While the incentive structure of science makes scientists more likely to be careful in their research and truthful, this is far from a guarantee. Moreover, our trust in the maintenance of this incentive structure assumes that there are scientists who are trustworthy as individuals and act in light of a set of ethical values. Finally, when a novel claim has been made or there is disagreement among scientists about a particular issue, our trust in scientific methods and the organization of science will not settle the matter. We will need to make decisions about the trustworthiness of individual scientists.

² This last difference is picked out by survey questions asking members of the public whether they have confidence that scientists “act in the best interests of the public” and trust scientists “to give full and accurate information” about certain topics. For an example see The PEW Research Center’s report (Funk and Kennedy 2016).

³ This restriction in the scope of our account is in line with Kitcher’s approach. See Kitcher (2011, pp. 155-168).

We also underline at the outset that we are concerned with *warranted* epistemic trust, not with a socio-psychological account of why public trust or distrust scientists. By warranted epistemic trust we mean well-placed epistemic trust: trust with good grounds invested in those scientists with the required qualities for being trustworthy. The idea of warranted trust also resonates well with Kitcher's well-ordered science as a normative ideal.

We find it useful to distinguish between qualitatively two different types of trust: basic and enhanced.

Basic trust. Let M be a member of the public, S stand for a scientist or a body of scientists, and P be a proposition or a group of propositions. Then in order for M to place warranted *basic* epistemic trust in S as a provider of P, the following conditions are necessary:

(C1) S believes that P and communicates it to M honestly,

(C2) M takes the fact that S believes and has communicated that P to be a (strong but defeasible) reason to believe that P,

(C3) P is the output of reliable scientific research carried out by S, and

(C4) M relies on S because she has good reasons to believe that P is the output of such research and that S has communicated P honestly.

Several clarifications are needed. First, a piece of scientific research that yields P is reliable if it is carried out according to methodological standards that make the objective probability of P's being true sufficiently high (Goldman 1986). Second, we can move from M's trust to the public's as follows: the relevant public will be said to invest warranted epistemic trust in S as a provider of P when a sufficient number of its members invest warranted epistemic trust in S as a provider of P. Our formulation refers to the relevant public because the boundaries of "the public" are determined contextually and usually left implicit. Finally, while (C1) and (C3) spell out the necessary conditions for S to be trustworthy, (C2) and (C4) specify the necessary conditions for M to trust S and do so with good reasons. What is essential for the existence of trust as opposed to mere reliance is both the requirement of honesty in C1 and the requirement that it is S's communication (and not just

S's belief) which is a reason for M's belief that P. Trust differs from reliance in two crucial respects. First, it involves dependence on the good will of another person (Baier 1986, pp. 234-5). Second, when the person trusted fails to live up to the trust entrusted in them, this gives rise to reactive attitudes such as feelings of betrayal and resentment (Holton 1994, pp. 66-7). When we take what a scientist has told us as a reason to believe the claim they are making, we take them to be honest, thus displaying some level of good will towards us. It is not the case that we only assume the incentives to tell the truth is what guides their action. If that were the case, we would not feel betrayed if we found out they had lied to us.⁴

Recall that S can be an individual scientist or a group of scientists such as a research team or a body like the IPCC. The requirement of honest communication implies that S has beliefs, which in turn implies that groups can have beliefs. This point requires two clarifications regarding our account. First, our account is intended to be neutral between summative accounts of group beliefs, which require that a group has a belief that P only if all or most of its members believe that P, and non-summative accounts that reject this requirement (Gilbert 1987, p. 186). That is, we set aside metaphysical questions regarding the reducibility or irreducibility of collective intentional attitudes to individual ones. For our account to be applicable to groups, only the following conditions need to obtain: Groups can have beliefs, aims and desires and can make assertions, and furthermore we can evaluate their competence and sincerity. Second, successful epistemic collaboration within groups depends on group members placing warranted trust in other group members. Individuals who are deciding whether to place trust in the claims made by such a group do not need to decide whether to trust each individual member. Instead, they can rely on the following principle:

Transmission of warranted trust principle: If (a) S_1 has placed warranted basic epistemic trust in S_2 as a provider of P and (b) M has a strong reason to believe that (a) obtains in light of facts about S_1 and S_1 's conduct, then M has warranted basic epistemic trust in S_2 as a provider of P.

This principle holds not only for members of research groups but generally. Thus, for instance, if a journalist has well-placed trust in a scientist and if

⁴ For further discussion, see Irzik and Kurtulmus (forthcoming).

their readers know that the journalist would have taken due care in placing her trust, then they can also trust the scientist.

Enhanced trust. Since there is always an inductive gap between a hypothesis and evidence, every scientific hypothesis is vulnerable to two types of error known as inductive risks (see Rudner 1953; Douglas 2000 and 2009). One may accept a hypothesis as true even though it is false, or one may reject it as false even though it is true. Both Rudner and Douglas argue that social and moral values can legitimately influence scientists' methodological standards that determine the threshold of evidence needed for the acceptance or rejection of the hypothesis in view of the consequences of its inductive risks. Suppose now that either risk has serious consequences for the public. If scientists (S) make their methodological decisions in line with the public's assessments of inductive risks, and moreover if the public has good reasons to believe that S have indeed done so, then they would certainly invest further (over and above basic) trust in S. Hence, two further conditions can be added to our earlier account:

(C5) When public welfare is at stake, S make their methodological decisions regarding the distribution of inductive risks with respect to P in agreement with M's assessments of those risks, and

(C6) M has good reasons to believe that Condition 5 is satisfied.

Basic epistemic trust plus these two conditions yield public's (warranted) *enhanced* epistemic trust in science.

Condition 5 in effect says that for there to be enhanced trust the social and moral values of the scientists and the public should more or less match with respect to the distribution of inductive risks regarding P, but only when public welfare is at stake.⁵ However, it is not enough that there is a value match.

⁵ This qualification is necessary because when inductive risks are not relevant to the public, they need not play a role in the public's trust. However, this introduces a complication: research that does not appear to be relevant to the public's welfare may later turn out to be relevant. Thus, there may be cases when we, as individuals, mistakenly assume that considerations of enhanced trust are irrelevant. Such errors on our part are unavoidable. Nevertheless, judgments regarding trust are not made once and for all. We can revise our judgments about trust in the light of new information and correct our errors.

Members of the public should also have reason to think that this is indeed the case, and that's what Condition 6 says.

The difference between basic and enhanced trust can be highlighted with a hypothetical example from a slightly different context. Suppose that I have some money I am considering to invest. I hear about a new venture and wonder whether it is a good investment. I ask three people for advice. All three say that I should invest all of my money in the company. However, their credentials and risk profiles are different. The first person has no knowledge about such matters. This person does not merit basic trust and their advice should have no bearing on my decision. The second person has proper credentials, but is unlike me in that they are not risk averse. In this case, it would be a good idea to invest some but not all of my money in this venture. The third person also has proper credentials. Moreover, they are, like me, risk averse. In this case it makes sense for me to invest all of my money in this venture. The first person exemplifies the absence of trust. The second case exemplifies a case with basic trust, and the third a case with enhanced trust. This hypothetical example shows that there is indeed a conceptual and practical difference between basic trust and enhanced trust – a difference that does translate into different behaviors.

In cases where human welfare is at stake and enhanced trust is absent, basic trust is second-best and preferable to the total absence of trust. The existence of basic trust means that scientists who issue a certain claim have carried out reliable research. Thus, this gives the public reason to believe their claim and act on it. But as this example illustrates, the public's actions will not be solely based on the information offered. They will also include precautions against risks that the research has not fully taken into account.

At this point, it is useful to note the following feature of our account: it includes internalist as well as externalist elements. Both basic and enhanced trust are warranted only if they are placed in people worthy of such trust and in light of good reasons. It entails that warrant for trust depends on both the truster and the trustee. This does not point to a problem with our account; rather, it underscores the difficulty of successful relationships of trust. Our trust would not be well-placed if the person we trusted was unreliable. Similarly, if we chose whom to trust randomly, our trust would not be well-placed even if the person we trusted turned out to be trustworthy.⁶

⁶ We thank one of the anonymous referees for urging us to clarify this aspect of our account.

2 Well-ordered science: From promoting collective good to restoring public trust

Kitcher (2001) finds the current social organization of science in many affluent democracies unsatisfactory because research priorities and their applications are determined jointly by scientists and a privileged group of ‘paymasters’ with the funds to support them in a way in which the values and the interests of less privileged groups of people tend to be neglected (Kitcher 2001, 127ff; Flory and Kitcher 2004; and Reiss and Kitcher 2009). The ideal of well-ordered science aims to overcome this shortcoming of the existing social organization of science by taking into account the perspectives and values of all groups in scientific agenda setting and the allocation of resources to them. In this way, Kitcher argues, collective good will be promoted.

By contrast, Kitcher (2011) focuses on the loss of public trust in science. “My primary thesis” is that, he writes, “scientific authority has been eroded in recent decades” (Kitcher 2011, p. 20). He has in mind the American people's resistance to the Darwinian account of evolution and their skepticism about the safety of genetically modified foods, the reality of human induced global warming and the like despite the existence of wide scientific consensus with respect to these issues. He takes such cases as evidence for decline in public trust in science. He cites two major reasons for this: (a) the fact that many people perceive certain areas of science (like evolutionary biology and climate science) to conflict with their deeply held values and (b) the overambitious and often unsubstantiated claim that all questions of major importance can eventually be answered by science, a view he calls “scientism”, that “stimulate[s] reactions that detract from credibility of more sober scientific judgments” (*ibid*, pp. 16-17). Scientism is seen by ordinary citizens as the illegitimate intrusion of scientists' own interests under the guise of a value-free science, and, according to Kitcher, that's “the deepest source of the current erosion of the scientific authority” (*ibid*, p. 40). Accordingly, one of the main motivations behind extending the scope of his model is to restore public trust in science (Kitcher 2011, pp. 30-31, 40, 152, 178, 185-187, and 222-225). Let us see how.

The central element of Kitcher's new model is the notion of “the system of public knowledge”, science being its most successful part. It incorporates a public depository of knowledge and consists of the contexts of inquiry,

submission, certification, and transmission. Whereas the earlier model is solely concerned with the context of inquiry, the revised model additionally includes all other contexts of the public depository. As in Kitcher (2001), the context of inquiry is about determining research agendas and the allocation of resources to them. The context of submission is concerned with the questions of who is entitled to submit reports to the public depository of knowledge and what standards (mostly, ethical) they must meet in their investigations. The context of certification deals with the methods, procedures and standards for accepting and rejecting submitted reports as part of public knowledge. Finally, the context of transmission involves people's having access to the public depository and specifying which applications of public knowledge are to be pursued.

Kitcher now argues that the contexts of inquiry, submission, certification, and transmission must all be well-ordered and that the system of certification must be ideally transparent. Science would be well-ordered when these four contexts are governed by an ideal hypothetical deliberative reasoning process carried out by ideal deliberators who are representative of all segments of society and who are tutored by scientists.⁷ Ideal transparency, on the other hand, means that people see how knowledge claims are accepted into the pool of public knowledge. Opening up the context of certification to the scrutiny of ideal deliberators and requiring the system of public knowledge to be ideally transparent is a radical revision to Kitcher's original model and the key to his argument for restoring epistemic public trust in science. Let us then look into them carefully.

Kitcher begins his discussion of the context of certification by noting that in order for a piece of research to be certified we need some methodological standards that would ensure its reliability (Kitcher 2011, p. 148). He writes: "The system of public knowledge is expected to accord with methodological guidelines that are reliable, in the sense that following them would tend to generate correct conclusions... Reliable standards and processes are those giving rise to conclusions that are true *enough*, at a frequency that is high

⁷To be exact, in contrast to the original formulation, the ideal deliberators now represent the entire world population and future generations, but for our purposes we will focus on a single society.

enough” (*ibid.*, 148; emphasis original).⁸ This is similar to how we defined reliability in Sect. 1.

How close to truth is close enough and how high is high enough? According to Kitcher’s model, the levels of sufficiency will be determined by ideal deliberators and scientists collectively in the context of certification. The context of certification will be well-ordered “just in case an ideal deliberation would endorse levels of proximity to the truth and of probability of generating truth so that both the general methodological standards enunciated and the particular judgments extending those explicit standards fall within the range of reasonableness determined by those levels” (*ibid.*, p. 149). In setting the levels of sufficiency, Kitcher tells us, the deliberation should take into account “values of consequences” of research (*ibid.*, p. 148). We interpret this to mean taking into account the associated inductive risks. It is for this reason that values pervade science not just at the agenda setting stage, but also at the very stage of determining the threshold for what counts as sufficient evidence for a given hypothesis. Turning now to the requirement of transparency, Kitcher writes that “a system of public knowledge is ideally transparent just in case *all people, outsiders* as well as researchers, can recognize the methods, procedures and judgments used in certification ... and can accept those methods, procedures and judgments” (*ibid.*, p. 151; our emphasis).

How exactly is all this supposed to address the problem of trust? The relationship between public trust in science and a well-ordered and ideally transparent system of certification can be seen more clearly in light of our analysis. While Kitcher’s model does not include our C1, C2 and the transmission of warranted trust principle, it does contain requirements that correspond to our conditions C3 and C4. A well-ordered system of certification implies that research results are based on methodological standards and procedures that make their probability of being true sufficiently high. An ideally transparent system of certification means that the public understands how those standards and procedures produce reliable research and accepts them. Hence, both reliability and its recognition by the public appear as necessary conditions of well-placed trust. Kitcher’s conditions can be formulated more explicitly as follows:

⁸ Kitcher speaks of the system of public knowledge here, but recall that science is its most central and successful component.

(K3) Scientific research should be based on reliable methodological standards and procedures, and

(K4) The public must be able to recognize and accept all methodological standards and procedures by which conclusions of research are included into the pool of public knowledge.⁹

Moreover, since the methodological standards and procedures are decided by taking into account the inductive risks involved, the public will see the role played by social and moral values in methodological decision making and recognize that the associated criteria adopted through ideal deliberation reflect their values as well as those of the scientists. As Kitcher puts it, “harmony not only obtains, but is recognized as obtaining.” (*ibid.*, p. 151) This can be formulated more explicitly in terms of two separate requirements as follows:

(K5) Methodological decisions about the choice of standards and procedures should reflect a harmony between the values of the scientists and those of the public regarding the distribution of associated inductive risks, and

(K6) The public should be able to see that the value harmony mentioned in K5 indeed obtains in light of their understanding of the methodological standards and procedures employed.

When these two conditions are satisfied, any suspicion to the effect that some claims are accepted simply because they serve scientists’ own interests and values will wither away. This will boost trust in science, and, by contrast, any lack of transparency in the system of certification will erode it.¹⁰

⁹ This way of numbering is intended to capture the parallel to the conditions in our account.

¹⁰ In a recent article Stephen John has offered a critique of the requirement of transparency. He uses the negative impact of the leaked emails from Climate Research Unit at the University of East Anglia to illustrate his reservations about transparency. As John observes, many of the practices that the scientists whose emails were hacked engaged in were “normal and respectable” (John 2018, p. 81). However, given a set of false background beliefs about how science works, these practices were seen as outrageous and undermined trust. In light of these, John maintains that non-experts’ epistemic interests would be better served in the absence of transparency. Note that he is making a conditional claim: if people have mistaken

The similarities between K3-K6 and conditions in our account should be obvious, but so should the striking differences. First, while K3 and C3 are the same, K5 and C5 are similar but not identical. K5 requires that in making their methodological decisions regarding a particular issue, scientists should *always* take into account the consequences of error in a way to reflect the values of the citizens. By contrast, C5 requires that value harmony should obtain only when inductive risks bear on public welfare. Indeed, it seems to us that K5 is unnecessarily too broad for the purpose of trust building because there is no reason to require value harmony in, for example, fundamental research in star formation or unified field theory, where it does not matter for the public how the inductive risks are distributed.

It may be thought that K4 cannot be a condition of trust because when one is in a position to understand and evaluate the methods behind a certain claim there is no role left for trust (Douglas 2013, p. 904). Even though there is less dependence on trust when K4 obtains, trust still has an important role to play. Scientists are often in this position vis-à-vis other scientists, yet they need to be able to trust them. For instance, when a laboratory team reports the results of a certain experiment, other scientists who are in a position to evaluate the team's methods will still need to trust that they carried out the experiment with competence and reported their results truthfully. The same point applies to ordinary citizens as well when K4 is satisfied.

Nevertheless, K4 and K6, drawn from the requirement of transparency, are much more demanding than our C4 and C6 and indeed impossible to meet in practice, for, as Kitcher himself notes (2011, p. 151), no one can recognize all the methods and standards by which scientific claims are certified. For the same reason, nor is it possible to see that value harmony obtains. Thus, if we insist on K4 and K6 as necessary conditions for warranted trust, both basic and enhanced trust become unattainable. By contrast, as we shall show in Sects. 3.1 and 3.4, C4 and C6 are realistic and much easier to satisfy since they require from the public only to have good reasons for thinking that research is reliable and that value harmony obtains. Finally, Kitcher nowhere

views about how science operates, then transparency is dangerous. Taken as a whole, Kitcher's proposal undermines the antecedent of this conditional. When Kitcher's ideal obtains, people have an accurate understanding of how science works. Nevertheless, John is also making a broader point, which is important. Under non-ideal conditions transparency can have costs. Thus, when we seek to bring about conditions necessary for warranted trust, we should not myopically focus on improving only one set of conditions but pursue a more holistic approach.

distinguishes between different types of trust such as basic and enhanced trust as we have done. As we shall see in the next section, these differences create difficulties for him.

Kitcher thinks that there are several institutions and practices, such as citizens' panels and juries, that could give substance to his ideal and approximate it to some degree. He favors deliberative polling developed by James Fishkin and modifies it into a three-stage process. At the first stage, the experts should present to a representative group of citizens the existing consensus among the scientific community about a particular scientific issue along with relevant central concepts, methodological procedures and standards of evidence, in short, all the major elements of the research that led to the consensus. At the second stage, challenges to the consensus should be presented. At the final stage, the scientific community should offer their counterarguments to the challenges, and the citizens should be "encouraged to press until they had resolved their doubts. Only if – and when – full resolution is achieved should citizens be prepared to report a unanimous verdict on behalf of the scientific consensus" (*ibid.*, p. 225).

Suppose now that a sufficient number of representative citizens can be tutored in the way envisioned above. Kitcher argues that they can then serve as credible witnesses, mediating between the scientific community and the larger public (*ibid.*, p. 129). They could share their tutored understanding of science with the broader public and then report back their experiences. This reiterable mediation would provide continued flow of information and dialogue between the scientists and the larger public and thus greatly contribute to restoring trust in science.

In short, for the purpose of building or restoring epistemic public trust in science, Kitcher proposes a two-step process for approximating his ideal in a liberal democratic society. The first step consists of an actual deliberative consensus building among a representative group of citizens tutored by experts with regard to a given scientific issue. The second step consists of the tutored citizens informing the larger public about the consensus they have reached. The whole process can be reiterated as many times as necessary.

3 Evaluating Kitcher's proposals for trust building

Kitcher's proposals regarding the building of public trust face a number of problems. Below we introduce them and suggest ways of addressing them.

3.1 The Problem of Recognizing Reliability

In Fishkin-type deliberative polling, the representative citizens are to be tutored by experts in such a way that they recognize the evidence, methodological standards and procedures that led to the consensus about a particular scientific claim. To put it differently, they are to be convinced of the reliability of research that underlies the consensus in question by acquiring first-order reasons. This is how, according to Kitcher, the condition K4 for trust is to be satisfied. We believe this is too demanding. While we do not deny that some representative citizens can be tutored in this way for some relatively simple scientific claims in a few days or weeks, we doubt that this can be achieved for all of them for all scientific issues. Take, for example, the scientific consensus that global warming is anthropogenic. Understanding the evidence itself and the methods by which they are obtained require sophisticated scientific knowledge, which may well be beyond the capacity of representative citizens who are not necessarily well-versed in climatology. Moreover, even if they can be somehow taught the first-order reasons for believing that global warming is human-induced, it is impossible for the tutored citizens to inform the public at large in the same way. Short of all citizens going through exactly the same type of deliberation that the representative citizens go through, which is practically impossible, the larger public are unlikely to understand and assimilate the first-order reasons that the tutored citizens have been introduced.

In view of these facts we propose that instead of the reliability of research itself the representative citizens and the larger public can judge the trustworthiness of the researchers themselves. To be able to do that, they need some criteria which could function as second-order reasons. These may include assessing expertise, honesty, and epistemic responsibility (see Anderson 2011 and Goldman 2001). Criteria for judging expertise reflect a hierarchy of expertise ranging from holding a PhD in the relevant field of inquiry to being an expert in that field whose work is received well by her peers, all the way up to being a leader in the field as indicated by being honored by prestigious awards such as the Nobel prize. Those who are higher in the ladder of hierarchy should be taken as more trustworthy epistemically than those who are lower. As for criteria for judging honesty, they include presence or absence of conflicts of interest and previous misconduct such as plagiarism, fabricating or suppressing crucial data and so on. While their absence is not a reason to trust a particular scientist issuing a particular scientific statement, their presence is a reason not to trust him with respect to

that statement. In the case of groups of scientists working together, we can evaluate the members of the group individually. A group can also be evaluated as a whole, especially in cases where it has formed an organization. For instance, research labs can also have conflicts of interest, be recognized by peers for doing good research or have cases of past academic dishonesty. Finally, epistemic responsibility can be judged by checking whether the researcher publishes in peer-reviewed scientific journals, whether she refuses to share data for no good reason, and whether she fails to acknowledge or engage with refutations of her own claims, and the like.

Therefore, we suggest that in addition to being informed about the first-order reasons, the representative citizens should also be tutored about the trustworthiness of scientists who have carried the research in question to judge its reliability in an indirect way and then share both types of reasons with the rest of the public. Arguably, untutored citizens can understand and employ second-order reasons even when they fail to master the first-order reasons. Moreover, if the tutored citizens come to accept a certain conclusion, that also is a second-order reason for the rest of the public to accept it as well. Even though they may not appreciate the first-order reasons for the conclusion in question, they can assume that had they also been tutored, they would come to accept it. Our proposal therefore has a greater chance of success in building public trust than Kitcher's own.

3.2 The Problem of Trust in Tutors and the Idea of Tutoring

The success of the Fishkin-type deliberative process crucially depends on the existence of a certain degree of initial trust between the representative citizens and their tutors. Kitcher simply assumes that such a trust base between them already exists. But if there is a serious erosion of public trust in science as Kitcher worries, then such an assumption appears to be unjustified.¹¹

Yet trust is not an all or nothing affair, and if the encounter between the experts and the representative citizens is structured appropriately, then a healthy level of trust between the two groups can be established during the process. To that end, first, the tutors must be chosen from trustworthy experts – a point Kitcher surprisingly neglects. Moreover, the tutors should be experts whose trustworthiness is relatively easy to observe by the representative citizens. Thus, for instance, they must not suffer from conflicts of interest.

¹¹ We thank one of the anonymous referees for raising the challenge of establishing trust between the tutors and the representative citizens.

Furthermore, the representative citizens should be provided with second-order reasons regarding their trustworthiness. This can be done by providing them with the credentials, the institutional affiliations of the experts, and so on.

Second, Kitcher's emphasis on one-way tutoring of the representative citizens by the experts smacks of indoctrination, which can be an obstacle for successful deliberation. Indeed, as some philosophers have noted, his entire project of well-ordered science is based on "the deficit model" (Brown 2013, p. 394), according to which the public's waning interest in science, their skepticism or outright hostility toward it is due to their lack of knowledge and understanding (see Kitcher 2011, pp. 128-129 and 187 ff).¹² What is needed is a genuinely dialogical communication between the participants of the deliberative process, where the experts should not only provide the representative citizens with up-to-date relevant knowledge, and indeed with the scientific consensus if any, about the issue in question, but also try to understand their beliefs and values, uptake their valid criticisms and address their worries (Goldenberg 2017; Grasswick 2010).

Finally, Kitcher and the tutors must be reminded that expertise always comes in degrees, even lay people can have a certain degree of expertise in some domains (Collins and Evans 2007), and at least on occasion it is the tutors that need to be tutored by the lay experts rather than the other way around (Wynne 1996).

We believe that these points and proposals, when taken into account, can establish sufficient grounds for trust between the tutors and the representative citizens and contribute to the success of the Fishkin-type deliberative polling.

3.3 The Problem of Value Harmony

Kitcher's entire discussion (both of his ideal model and its implementation) proceeds as if condition K5 can always be satisfied, that is, as if it were always possible to establish a harmony between the values of the public and those of the scientists regarding the inductive risk

¹² For an overview and the evaluation of the deficit model, see Sturgis and Allum (2004).

assessment.¹³ We believe that there are situations in which this may not be the case.

Consider, for example, the controversy regarding the (alleged) causal link between the measles-mumps-rubella (MMR) vaccination and autism.¹⁴ The causal claim was due to Andrew Wakefield's article published in the prestigious medical journal the *Lancet* in 1998 and publicized at a press conference. Later on Wakefield was found guilty of serious scientific misconduct, struck off the medical register, and his paper was retracted from the *Lancet*. The current scientific consensus is that there is no evidence for the existence of a causal relationship between MMR vaccination and autism. Let us then consider the hypothesis H: "The MMR vaccine does not cause autism". The inductive risks associated with H consist of accepting H as true when it is in fact false (false positive) and rejecting it as false when it is in fact true (false negative).

While we lack data about how parents and scientists informing the public in the UK ranked these risks, we surmise that scientists would be primarily concerned with public health, whereas parents would be concerned with the welfare of their children. As a result, their respective assessments of inductive risks would diverge. Consider the consequences of rejecting H when it is in fact true from the perspective of scientists. They have to consider not only the harms of mumps, measles, and rubella, but also the costs of allocating resources to treating the increased incidence of these diseases – resources that could otherwise be used to treat other diseases. Parents who are thinking foremost about their children's welfare need not worry about this.

Herd immunity provides a further reason for the two groups to judge risks differently. When the rate of vaccination within a population reaches a certain level, those who are not vaccinated are also protected. Suppose herd immunity obtains. If scientists falsely reject H and announce it to the public, this could lead to a decline in the vaccination rate which could in turn undermine herd immunity. They would thereby be putting the whole population, including people who cannot be vaccinated, at risk of MMR. Under the same circumstances, when a parent falsely rejects H and acts on this false belief,

¹³ This assumption is made also by others who argue for the relevance of moral and social values in scientific research and science policy. See, for example, Wilholt (2013).

¹⁴ We discuss this case in more detail in Irzik and Kurtulmus (forthcoming). For informative discussions of this controversy as it played out in the UK, see Boyce (2007) and Fitzpatrick (2004).

they are not taking a huge risk. They can rely on herd immunity. On the other hand, if they accept H when H is in fact false and vaccinate their child, they put their child at risk of autism.

In view of these considerations, the public and the relevant community of scientists could disagree about the assessment of inductive risks with respect to H. Therefore, condition K5 (or our C5 for that matter) is not satisfied. Since it is a necessary condition of warranted trust from Kitcher's perspective, the public trust in science can never be built in the MMR case even if all other conditions are satisfied. This conclusion is too strong. Fortunately, it can be avoided. Recall the distinction we made between basic and enhanced epistemic trust in Sect. 1. While the public may not be in a position to invest warranted *enhanced* epistemic trust, they can still have warranted *basic* epistemic trust in scientists provided that research that yields H is reliable and that people have good reasons to believe that H is based on such research.¹⁵

In the MMR case, a parent who has no trust in scientists and a parent who has only basic trust in scientists appear behaviorally identical: neither will vaccinate their child. Thus, it may be asked whether basic epistemic trust makes any difference at all. The following scenario illustrates the difference between complete lack of trust and having basic epistemic trust in the MMR case. Suppose that these people are not making a decision for their children but are responsible for determining public health policy. The parent who has no trust in scientists would not implement a public policy that would increase the vaccination rate, whereas the parent who has only basic epistemic trust in scientists would implement it. Thus, having basic epistemic trust differs behaviorally from both the case of no trust and that of enhanced trust.

This is a case where the reasons one has *qua* parent and *qua* public official entrusted with public welfare differs. It is not surprising that the reasons one has as someone concerned with an individual's welfare and as someone concerned with collective welfare should be in conflict. It is a familiar fact from collective action problems such as the Prisoner's Dilemma: when each individual does what is individually rational, the outcome is worse for everyone. If we were choosing the course of action for all, we would want

¹⁵ When members of the public do not have enhanced epistemic trust in scientists, this does not entail that scientists are at fault. It is possible that in certain cases the public is at fault for withholding enhanced trust. In fact, John argues that parents who demanded higher epistemic standards than the ones employed by scientists in the MMR case were morally wrong to do so. See (John 2011).

everyone to cooperate even though when we are acting individually defecting is the rational course of action.

The MMR case demonstrates why it is important to have an explicit and nuanced analysis of the notion of public trust in science. It also shows that value harmony between scientists and the public regarding the distribution of inductive risks may not always come about, in which case enhanced trust does not obtain. But this does not mean that the public cannot invest any epistemic trust in science at all; they can still enjoy a basic level of trust. It seems to us that these considerations can be fruitfully incorporated into Kitcher's account.

3.4 The Problem of Observing Value Harmony

As we saw in Sect. 2, Kitcher requires not only that there is value harmony (K5), but also that the public is able to see that it obtains in light of their understanding of the methodological standards and procedures employed (K6). We pointed out that this is an unrealistic requirement as Kitcher himself admits, but he offers no alternative. By contrast, our C6 requires only that there be good reasons for believing that there is value harmony. What could they be? We suggest that facts about individual scientists and diversity among them provide such reasons.¹⁶ General facts about individuals such as their sex, social class, racial and ethnic background influence the values they hold and offer them different perspectives regarding the goodness or badness of certain outcomes. Thus, if one narrow social group is dominant among scientists, the judgments regarding inductive risk is likely to reflect their perspective and values and suffer from the blind spots they may have due to their social position. However, if scientists are drawn from different social groups, then their value judgments would be more likely to represent the values of the whole of society. The same point applies to the selection of tutors in deliberative polling. A more diverse group of tutors has a greater chance of gaining the trust of the representative citizens during the Fishkin-type deliberative process, which in turn facilitates healthy consensus building.

Note that this line of reasoning is available to members of the public as well. They can also reason that if scientists represent all groups in society, the values of their group are more likely to be taken into consideration than would be the case if the group they belong to were excluded. Moreover, the diversity or

¹⁶ Diversity can also provide grounds for basic trust. Diversity with respect to gender, class, race and ethnicity can boost the reliability of research by exposing, eliminating or balancing biases and prejudices (Keller 1985; Longino 1990; Nielsen et al. 2017).

its lack within the scientific community is something that can be observed with relative ease. Thus, the diversity of the scientific community can act as a second-order reason that facilitates the building of enhanced trust.

Diversity of scientists can also help address a second challenge to building trust. The public is not a monolithic entity and does not share the same history of interactions with the scientific community. Thus, certain groups may need more evidence of the good will and honesty of scientists or the reliability of their research than others who do not share the same history.

Consider, for instance, the notorious Tuskegee syphilis study and its legacy. Four hundred African-American men were recruited into the study to observe the effects of syphilis in African-American males when it goes untreated. The subjects were not informed that they had syphilis. Instead, they were told that they were being treated for “bad blood”, an ill-defined condition referring to various health problems. During the course of the study, which lasted from 1932 to 1972, the subjects were not treated for syphilis even though penicillin’s effectiveness in curing the disease was established in the 1940’s. Condition C1 for public trust, which requires that scientists communicate the relevant information honestly, was obviously violated. When the news about the Tuskegee horror story finally broke, would it be surprising if the African-Americans became suspicious of the entire medical community? Indeed, it has been argued that the Tuskegee study has caused African-Americans to mistrust medical institutions (Epstein 2007, p. 193) and that this mistrust has contributed to health disparities between them and other social groups.¹⁷ Against a background of past injustices that have bred mistrust, the diversity of scientists and the composition of tutors become significant. If members of disadvantaged communities are well represented in both, there will be greater chance of restoring their trust in scientists. A recent study lends support to this suggestion. Alsan and her colleagues found that African American patients were more likely to accept and act on medical advice when it came from African American doctors and this effect was more pronounced in subjects who had mistrust of the medical system (Alsan, Garrick and Graziani: unpublished).

Nevertheless, if diversity is to have the desired effect, it should go beyond merely having members of different social groups within the scientific

¹⁷ For a remarkable attempt to document this quantitatively, see Alsan and Wanamaker (2018). For related philosophical discussions of this case see Grasswick (2010) and Scheman (2001).

community; they should be in a position to influence their discipline. Although, as some commentators have noted, a number of African American medical professionals were involved in the Tuskegee syphilis study (Jones 2008), they had very limited power. Indeed, African American doctors were entirely excluded from institutions such as the American Medical Association that have an influential role in guiding the profession of medicine.¹⁸ Furthermore, it is naïve to expect that having a diverse body of scientists in a single generation can undo a long history of bias that has undermined trust in science. In such contexts diversity can deliver its expected benefits only in the long run. Restoring lost or damaged trust is often more difficult than building trust.

4 Conclusion

Ensuring that people have well-placed trust in scientists is a crucial goal for democratic societies that seek to treat people as free and equal citizens. In authoritarian societies public opinion does not have the same impact on social outcomes. Moreover, such societies do not have the same concern with respecting people's autonomy and rationality: it does not matter if citizens trust scientists in light of good reasons or due to propaganda. Thus, if democratic societies are to have sound policies informed by science and at the same time maintain democratic legitimacy, conditions for well-placed trust need to be in place. Therein lies the significance of Kitcher's proposals for building and maintaining trust.

While a normative account, such as Kitcher's, can find faults with existing societies and promote goals that are not easily attainable, the ideal it sets out should not be impossible to realize. We have argued that Kitcher's proposal is unrealistic for a variety of reasons. Thus, if Kitcher's account were correct, warranted public trust in scientists would be unattainable. This would mean that democratic societies would have to make a very difficult choice. Either they would have to give up on democratic values or their policies would not be guided by science. Democratic societies do not have to make this choice because Kitcher's (implicit) account of warranted trust is more demanding than is necessary. Let us briefly summarize how our account offers a more attainable goal while respecting some of the guiding ideas in Kitcher's account.

¹⁸ For the view that the roles of African Americans in the study have been exaggerated see Washington (2008, pp. 175-177).

First, according to Kitcher value harmony is a necessary condition for warranted trust. Our distinction between basic trust and enhanced trust allows that there can be some warranted trust even in the absence of value harmony. Moreover, unlike Kitcher we maintain that value harmony is needed in a narrower set of cases, i.e., only cases where public welfare is at stake.

Second, Kitcher's account requires that people understand and accept the reasons scientists have for accepting scientific claims. Our account also maintains that people should trust scientists in light of reasons, but we expand the set of reasons that citizens can employ to include second-order reasons, thereby making warranted trust more attainable.

Whether scientists are trustworthy and whether people have access to reasons they need to recognize the trustworthiness of scientists depends on the way we organize our society, science, and deliberative fora. There is, thus, a division of responsibility for ensuring that members of the public have warranted trust in scientists. The burden that members of the public need to bear in this context can be lightened by altering the reasons available to them through different organizations of science and society. For instance, having a more diverse body of scientists that represent different social viewpoints would make value harmony more likely. Such diversity can be much more easily observed and understood than methodological choices.

The point about altering social conditions to make the realization of warranted trust easier carries over to the design of deliberative fora like Fishkin-type polling as well. By structuring the encounter between citizens and scientists in a dialogical way and by electing tutors whose trustworthiness is easier to observe, it is possible to make deliberation more conducive to the building of well-placed trust in science.

References

ALLEA (2018). All European Academies (ALLEA) Newsletter #18, June 2018.

http://www.allea.org/wpcontent/uploads/2018/06/ALLEA_Newsletter_Number_15_June_2018-1.pdf

Alsan, M. and Wannamaker, M. (2018). Tuskegee and the health of black men. *Quarterly Journal of Economics*, 133, 407-455.

Alsan, M., Garrick, O. and Graziani, G. (Unpublished). Does Diversity Matter for Health? Experimental Evidence from Oakland. NBER Working Paper No. 24787. <https://www.nber.org/papers/w24787>

Anderson, E. (2011). Democracy, public policy, and lay assessments of scientific testimony. *Episteme*, 8, 144-164.

Baier, A. (1986). Trust and Anti-trust. *Ethics*, 96, 231-260.

Boyce, T. (2007). *Health, risk and news: The MMR vaccine and the media*, New York: Peter Lang Publishing.

Brown, M. B. (2013). [Review of *Science in a Democratic Society* by Philip Kitcher]. *Minerva*, 51, 389-397.

Collins, H. and Evans, R. (2007). *Rethinking Expertise*. Chicago: The University of Chicago Press.

Douglas, H. (2000). Inductive risk and values in science. *Philosophy of Science*, 67, 559-579.

Douglas, H. (2009). *Science, policy, and the value-free ideal*, Pittsburgh: University of Pittsburgh Press.

Douglas, H. (2013).). [Review of *Science in a Democratic Society* by Philip Kitcher], *British Journal of Philosophy of Science*, 64, 901-905.

Drezner, D. W. (2017). *The Ideas Industry: How Pessimists, Partisans, and Plutocrats are Transforming the Marketplace of Ideas*. Oxford: Oxford University Press.

- Epstein, S. (2007). *Inclusion: The Politics of Difference in Medical Research*, Chicago: University of Chicago Press.
- Fitzpatrick, M. (2004). *MMR and autism: What parents need to know*, London: Routledge.
- Flory, J. and Kitcher, P. (2004). Global health and the scientific research agenda. *Philosophy and Public Affairs*, 32, 36-65.
- Funk, C., and Kennedy, B. (2016). *The Politics of Climate*. Washington, DC: Pew Research Center.
- Gauchat, G. (2012). Politicization of Science in the Public Sphere: A Study of Public Trust in the United States, 1974 to 2010. *American Sociological Review*, 77, 167–187.
- Gilbert, M. (1987). Modelling Collective Belief. *Synthese*, 73, 185–204.
- Goldenberg, M. J. (2017). Public Misunderstanding of Science? Reframing the Problem of Vaccine Hesitancy. *Perspectives on Science*, 24, 552-581.
- Goldman, A. (1986) *Epistemology and cognition*, Cambridge: Harvard University Press.
- Goldman, A. (2001) Experts: which ones should you trust? *Philosophy and Phenomenological Research*, 63, 85–110.
- Grasswick, H. E. (2010). Scientific and lay communities: Earning epistemic trust through knowledge sharing. *Synthese*, 177 (3), 387–409.
- Hausman, D. M. (2003). [Review of *Science, Truth and Democracy* by Philip Kitcher]. *Ethics*, 113, 423-428.
- Holton, R. (1994). Deciding to Trust, Coming to Believe. *Australasian Journal of Philosophy*, 72, 63–76.
- Irzik, G. and Kurtulmus, F. (forthcoming) What Is Epistemic Public Trust in Science?. *The British Journal for the Philosophy of Science*.

John, S. (2011). Expert Testimony And Epistemological Free-Riding: The MMRC Controversy. *The Philosophical Quarterly*, 61, 496-517.

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

Jones, J. H. (2008). The Tuskegee syphilis experiment. In E. J. Emanuel, C. Grady, R. A. Crouch, R. K. Lie, F. G. Miller and D. Wendler (Eds.), *The Oxford textbook of clinical research ethics* (pp. 86-96). Oxford: Oxford University Press.

Keller, E. F. (1985). *Reflections on Gender and Science*. New Haven. Yale University Press.

Keren, A. (2013). Kitcher on Well-Ordered Science: Should Science Be Measured against the Outcomes of Ideal Democratic Deliberation?. *Theoria*, 77, 233-244.

Kitcher, P. (1993). *The advancement of science*. Oxford: Oxford University Press.

Kitcher, P. (2001). *Science, truth, and democracy*. Oxford: Oxford University Press.

Kitcher, P. (2011). *Science in a democratic society*. Amherst: Prometheus Books.

Longino, H. (1990). *Science as social knowledge*. Princeton: Princeton University Press.

Nichols, T. M. (2017). *The Death of Expertise: The Campaign against Established Knowledge and Why it Matters*. Oxford: Oxford University Press.

Nielsen, M. W., Alegria, S., Börjeson, L., Etzkowitz, H., Falk-Krzesinski, H. J., Joshi, A., Leahey, E., Smith-Doerr, L., Woolley, A. W., and Schiebinger, L. (2017). Gender Diversity Leads to Better Science. *Proceedings of the National Academy of Sciences*, 114, 1740-2.

Reiss, J. and Kitcher, P. (2009). Biomedical research, neglected diseases, and well-ordered science. *Theoria*, 66, 263-82.

Resnick, H. E., Keegan, S., Huddleston, N. (2015). *Trust and Confidence at the Interfaces of the Life Sciences and Society: Does the Public Trust Science? A Workshop Summary*. Washington, DC: The National Academies Press.

Rudner, R. (1953). The scientist *qua* scientist makes value judgments. *Philosophy of Science*, 20, 1-6.

Scheman, N. (2001). Epistemology resuscitated: Objectivity as Trustworthiness. In N. Tuana and S. Morgen (Eds), *Engendering Rationalities* (pp. 23–52). Albany, NY: SUNY Press.

Sturgis P. and Allum, N. (2004). Science in Society: Re-Evaluating the Deficit Model of Public Attitudes. *Public Understanding of Science*, 13, 55-74.

Washington, H. A. (2008). *Medical apartheid: The dark history of medical experimentation on black Americans from colonial times to the present*. New York: Anchor Books.

Wilholt, T. (2013). Epistemic trust in science. *British Journal for the Philosophy of Science*, 64, 233–53.

Wilholt, T. (2014). [Review of *Science in a Democratic Society* by Philip Kitcher]. *Philosophy of Science*, 81, 165-171.

Wynne, B. (1996). May the Sheep Safely Graze? A Reflexive View of the Expert-Lay Knowledge Divide. In S. Lash, B. Szerszynski and B. Wynne (eds) *Risk, Environment and Modernity: Towards a New Ecology*, (pp. 44–83). London: SAGE.