

## Should I Do as I'm Told? Trust, Experts, and COVID-19

*Matthew Bennett*

**ABSTRACT.** The success of public health responses to the COVID-19 pandemic is sensitive to public trust in experts. Despite a great deal of attention to attitudes towards experts in the context of such crises, one significant feature of public trust remains underexamined. When public policy claims to follow the science, citizens are asked not just to believe what they are told by experts, but to follow expert recommendations. I argue that this requires a more demanding form of trust, which I call recommendation trust. I argue for three claims about recommendation trust: recommendation trust is different from both epistemic and practical trust; the conditions for well-placed recommendation trust are more demanding than the conditions for well-placed epistemic trust; and many measures that have been proposed to cultivate trust in experts do not give the public good reasons to trust in expert-led policy.

### 1. Introduction

In 2019 public trust in politicians was at an all-time low in many parts of the world. Then again, it had been low for quite some time.[1] Public trust in epistemic authority is a more complicated matter. The success of many right-wing politicians—Trump, Johnson, Bolsonaro, etc.—is partly due to their efforts to discredit traditional sources of information, including the “mainstream media” and whichever scientific institutions prove inconvenient to their political interests. But worries about the erosion of epistemic standards in public discourse are easily overstated, and some polling data have shown that publics are not as distrustful of experts as we might fear.[2]

Nonetheless, there are some political and social crises in which even a small degree of suspicion of scientific authorities can have disastrous consequences. One such crisis is the climate change emergency. It is often with a focus on the climate crisis that social epistemologists and philosophers of science have tackled a range of questions around trust in expertise, including how we can defend the rationality of trusting experts and what measures we can take to cultivate such trust (see e.g., Anderson 2011; John 2016). Public health crises are also highly sensitive to public trust in experts, and the COVID-19 pandemic is no exception to the principle that when epidemics pose serious threats to the health of whole nations, distrust of the relevant public health experts can have serious consequences.

Despite a great deal of attention to attitudes towards experts in the context of such crises, one significant factor in public trust in experts remains underexamined. The vast majority of philosophical work on public trust in experts has addressed questions about epistemic trust: forming beliefs on the basis of the testimony, broadly construed, of scientists and scientific communities. But as the COVID-19 pandemic shows us, a distinct kind of trust in experts is needed when experts not only provide information and data, but also play an instrumental role in the formation of science-led policy. When public policy claims to follow the science, citizens are asked not just to believe what they are told, but to follow expert

recommendations. I argue that this requires a distinct form of trust, which I call recommendation trust.

I will argue for three claims about recommendation trust. The first is that recommendation trust is different from both epistemic trust and a third form of trust, also popular among philosophers, which I call practical trust (Section 3). The second is that the conditions in which recommendation trust is well-placed are more demanding than the conditions for well-placed epistemic trust (also Section 3). This is because to have good reason to follow an expert recommendation I must not only believe that the expert is sincere and competent in their field, but also that the action they recommend is in my interest. This leads me to argue for my third claim (Sections 4 and 5): because the norm that governs the rationality of recommendation trust is more demanding than that which governs epistemic trust, many measures that have been proposed to cultivate trust in experts do not help to cultivate recommendation trust in science-led policy. In short, experts have a higher bar to clear when we are asked to follow their recommendations. If we are to ask the public to trust the recommendations of scientists, we must acknowledge that this is different from asking novices to accept facts. When science leads policy, it must work harder to merit public trust

## 2. THE SIGNIFICANCE OF TRUST

My account depends foremost on the assumption that public trust in experts is both necessary and desirable for an effective public health response to the coronavirus pandemic.[3] One reason to think trust in experts is particularly important in the context of the pandemic is that most current public health policies in response to the crisis depend on a significant degree of voluntary compliance. Many countries have taken social distancing measures to reduce infection rates to a manageable level. This usually involves radical lifestyle changes for citizens, including the closure of businesses, workplaces, and many public spaces, bans on travel for all but essential journeys, and complete self-isolation for those at greater risk of serious illness. It is fair to assume that it would be wholly unpractical to apply social distancing solely through comprehensive policing and direct enforcement. And I am confident that even if compliance with social distancing could be guaranteed wholly through coercion, this would be undesirable.

A state could coerce compliance with social distancing in ways that are less direct than totalitarian policing. The simple threat of such measures might be enough to persuade the public to comply. I am also taking for granted that we would prefer voluntary compliance with a public health policy to compliance secured solely through public fear of what the state might do to them if they do not comply. But note that this is not an all-or-nothing test of the desirability of a public health policy, as if such a policy either depends on states applying coercive power to all citizens, or it does not depend on such power at all. Any state-enforced policy, even where it is granted democratic legitimacy through public trust, will involve exercising coercive power over the few who refuse to comply. In this respect, trust in government secures not just a sufficient level of compliance with government policy, but also sufficient public support for the policing of that policy. It would be wrong to

think that a public health response secured through trust would involve no coercive state power at all.

Though I am taking for granted that public trust is a necessary part of a desirable response to the pandemic, I do not assume it is sufficient. There are many reasons for this. One is that citizens may be willing but unable to comply with government policy. Emerging research on public behavior during the pandemic in the UK indicates that although willingness to follow social distancing measures is generally high, the ability to self-isolate is not evenly distributed, with the most economically disadvantaged unable to comply (Atchison et al. 2020). Thus, some citizens who trust policymakers enough to be willing to follow policy might not be able to do so.

It is also possible for citizens to trust policy that does not merit trust, and I will be assuming that the kind of public trust that features in a desirable response to the pandemic is specifically trust that is well-placed. Here it helps to invoke a distinction sometimes made by social epistemologists between trustworthiness and credibility (Rolin 2002). An authority is credible when it is likely to be believed. Credibility is vulnerable to faulty reasoning by individuals when making judgements about whom to believe, but the more insidious threat to well-placed trust is posed by structural injustices that grant illegitimate authority to privileged social groups (Fricker 1998). By contrast, an authority is trustworthy when its claims should be given credence. Following Onora O'Neill's work on public trust (see e.g., O'Neill 2020), I will assume that desirable public trust must be trust in the trustworthy, and it is only accidentally related to credibility. However, one of the tasks of good science-led policy is to bridge any extant gaps between credibility and trustworthiness. I will return to this challenge in Sections 4 and 5.

My framing of the issues so far might be accused of confusing two very different matters: public trust in scientific experts and public willingness to comply with government policy. If some readers worry about this, then all the better for the purposes of this paper, for the distinction I wish to argue for is very similar to the distinction between trust in science per se and compliance with scientifically-informed policy. But before proceeding to put some conceptual distance between trust in science and trust in technocracy, it is worth noting some observations about the tandem pairing of government policy and scientific expertise that we see in many (though not all) government responses to the coronavirus crisis. What we find in the pandemic is that when science has a significant role to play in policy, trust in experts and compliance with government go hand in hand.

In the UK, where I write this paper, the government's response to the pandemic has been consistently presented to the public as "following the science." Daily press conferences began on March 12, 2020 with a briefing delivered by the prime minister flanked either side by the government's chief medical officer, Chris Whitty, and chief science officer, Patrick Vallance. Vallance and Whitty have since played a prominent role in communicating the government's policy on the coronavirus, standing alongside government ministers in many more press conferences and appearing on television, on radio, and in print to explain the UK's public health strategy. Between March and May, Whitty and Vallance have been joined

in this communications effort by inter alia Angela McLean, deputy chief scientific adviser, Jenny Harries, deputy chief medical officer, and NHS England's medical director, Stephen Powis.

Each of these government spokespersons are scientists—epidemiologists, clinical pharmacologists, mathematicians, public health physicians, and medical professors. Their most evident role in the current emergency appears to be to amplify the message that there is authoritative scientific reasoning that supports the social distancing measures the public are being asked to adopt. This “following the science” message depends on the public believing not just that these scientists are part of the communication of the policy, but also that they are involved in deliberations that result in that policy. That the UK government is in fact following the science is very much open to question, but it is evident that the UK's approach to communicating its message to its citizens is to heavily emphasize its (putative) technocratic nature.

Most other national government responses lie somewhere between two extremes: namely, policy that is directly contradicting the recommendations of most scientists and policy that is designed and implemented by scientists. These extremes are somewhat idealized insofar as the question of whether a given national policy denies or follows “the science” ignores the fact that there is no single authoritative scientific perspective on what should be done. Nonetheless, some notable examples approximate these extremes. In Brazil, Jair Bolsonaro has argued publicly with regional governors and the national health ministry about the wisdom of advice to all citizens to stay home. In the US, Donald Trump has made a series of dangerously ill-informed claims about possible treatments for COVID-19, from hydroxychloroquine to disinfectant, prompting multiple high-profile corrections from Anthony Fauci, director of the US National Institute of Allergy and Infectious Diseases. By stark contrast, the Swedish government has in large part handed over their public health response to state epidemiologist Anders Tegnell, who has become the public face of the Swedish government during the pandemic and leads a state science agency with a significant degree of autonomy from the government. Though it would be easy to overstate the level of political authority granted Swedish public health agencies (responsibility for policy still formally lies with government ministers), the Swedish case is perhaps the closest thing we have to a purely technocratic response to the pandemic.

For those governments that are not actively discrediting the recommendations of scientific experts, public trust in experts and public compliance with government policy go hand in hand. The significance of public trust in experts increases with the degree to which government policy is presented as technocratic. Note that this does not mean that a government's strategy for dealing with the pandemic must in fact be technocratic in order for public trust in experts to be a significant factor in securing public support. Public trust in experts can support policy implementation by ensuring that citizens are willing to comply with government public health responses that are perceived to be led by science, regardless of whether they are in fact led by science. Public trust in experts is also a significant factor in public support for genuinely technocratic policy, provided we wish to avoid measures to secure public endorsement that involve coercion or dishonesty.[4]

### 3. RECOMMENDATION TRUST

Limited though they are, there are lessons to be learned from the coronavirus pandemic about the nature of the public trust that supports acceptance of (at least putatively) science-led policy. Before outlining those lessons directly, it will help to be more precise about the kinds of trust that are relevant to the pandemic.

I will be treating trust in all of its forms as a tripartite relation: one person trusts another with regard to a particular object of trust.[5] What I call practical trust is to stake something of importance on a particular action or range of actions that I expect another to perform. Thus, what I mean when I say 'I trust the babysitter' is that I place the wellbeing of a child in the hands of the babysitter, and I am confident in doing this because I expect the babysitter will take good care of the child. Moreover, the domain of my trust in the babysitter is restricted, insofar as I do not trust them to do just anything. Epistemic trust is to believe something because another person has told us it is true. Epistemic trust requires that the truster ascribes sincerity to the trusted, but also competence in the matter about which the trusted is providing testimony. A similar domain restriction applies. I might trust a classicist to help me with Greek etymology but not trust them on matters of bicycle repair, even if I expect them to be sincere when answering questions in either area.[6] Both epistemic and practical trust differ from a third form of trust that I will call recommendation trust. I recommend-trust someone when I believe I should do something because they have told me I should. The domain restriction also applies to the act recommended by the person I recommend-trust: when I trust the recommendation of another it is because there is a particular area in which I take that person's recommendations to be trustworthy.

As I am primarily interested in differences between recommendation trust and epistemic trust, I will stop short of a full account of these three forms of trust. A comprehensive definition of recommendation trust would need to do more to distinguish it from other, deceptively similar ways in which we form beliefs about what we should do. Say I am taking a trip by car with a friend. The friend recommends I avoid the M25 this afternoon because of the likely traffic. But I have good reason to think that taking a detour route will take just as long, despite the heavy traffic on the M25. There are a number of ways I might do as my friend suggests, in response to their suggestion, without thereby following their recommendation. Perhaps I avoid the M25 to humour my friend (I know they tend to be insecure about others not taking their travel advice). Perhaps I avoid the M25 because I know that my friend has a phobia of traffic jams and is too proud to cite this as the reason for asking that we take a detour. Perhaps my friend is an intimidating person, and I fear the consequences of questioning their advice. Doing as my friend recommends for any of these reasons is not trusting their recommendation.

Note also that the same principle applies to public trust in the recommendations of politicians and experts. If members of the public accord with the recommendations of government because they fear what the government would do were they not to follow its recommendations, then the public is not acting on trust, but on fear. Policy that depends on this form of compliance is precisely the coercion-based policy that earlier I suggested we

would prefer to avoid. Moreover, compliance with public health recommendation is significantly different from compliance with law. I will restrict my focus to the question of when we have good reason to follow public health recommendations issued by experts.

One very common concern about epistemic trust is that to believe what we are told by experts we must believe without the warrant available to other kinds of belief (Burge 1995; Hardwig 1985, 1991; John 2018; Kukla 2007). Experts know about things I do not know about and sometimes could not know about. But many areas of expertise are areas in which it is important for non-experts to form true beliefs. Such beliefs must be formed without access to all of the evidence and without expert ability to understand the evidence. This threatens the possibility of both knowledge and rational belief about a wide array of basic facts. Do I know whether the Earth is flat without travelling? Should I believe that penicillin can be used to treat infection without first studying biochemistry? Knowledge of such matters depends on the rationality of beliefs grounded in the testimony of others, and the rationality of such beliefs is questionable in ways that the rationality of other beliefs is not.

John Hardwig's influential solution to this problem was to argue for the following principle:

(T) If person A has good reasons to believe that person B has good reasons to believe that P, then A has good reasons to believe that P. (1991, 697–98)

Hardwig suggests that if such principles can be demonstrated, we can rescue testimony-based beliefs from scepticism. If (T) holds, Hardwig argues, then we can be epistemically justified in believing what others tell us is the case without having to investigate the evidence ourselves. If Hardwig is right, the principle also gets us part of the way to understanding how epistemic trust in experts could be rational. Say that I believe my doctor to be competent, well-informed, and sincere. This same doctor tells me that if a person's heart stops outside of hospital there is a 20 percent chance that CPR will save their life. Provided I have good reason to believe the doctor to be competent, well-informed, and sincere, I have good reason to believe that if a person's heart stops outside of hospital there is a 20 percent chance that CPR will save their life.

With much greater complexity, a novice can adopt a comparable relation of epistemic trust in public experts communicating facts about COVID-19. The greater complexity comes partly from the fact that it is very unlikely that most non-experts would have direct access to the testimony of experts. Most of us will read about what experts have to say about the pandemic as it is reported by journalists. Thus, for example, the question of whether I should (epistemic) trust UK chief medical officer Chris Whitty is made more difficult if, rather than watching the press briefings led by Whitty, I instead read reports based on the briefings in a national newspaper. Most obviously the problem of the rationality of my trust in this situation is compounded by the additional question of whether I can trust the newspaper; I must have reason not just to believe that Whitty is competent and sincere, but also to believe that the journalists working at the newspaper are competent and sincere.

Hardwig's principle does not render epistemic trust easily won; there are many obstacles to rational epistemic trust.[7] It nonetheless does give us an account of how epistemic trust can be rational, and it is not implausible that its conditions for rational trust might sometimes be met. But even if the conditions set by Hardwig's principle for rational trust are met, we have not thereby met conditions for rational recommendation trust. This much is evident when we attempt to adjust Hardwig's principle to apply to recommendation trust.[8] When a person provides factual testimony, they communicate their belief that something is the case, and the question is whether this gives the receiver of testimony good reason to believe the same. When a person makes a recommendation, they communicate their belief that the receiver of the recommendation should do something. Accordingly, the analogous question is whether the recommendation gives the receiver good reason to agree that they should act as recommended. If we adapt Hardwig's principle to apply to recommendations, we arrive at the following:

(R) if A has good reasons to believe B has good reasons to believe that A should act a certain way, then A has good reasons to believe that they should act that way.

But (R) is not true.[9] Consider again the doctor who has told me about the likelihood of a person being revived through CPR. Imagine that this same doctor follows up with a recommendation: given the low likelihood of resuscitation and the potential for distress if CPR is attempted unsuccessfully, it would be best if I signed a Do Not Attempt CPR (DNACPR) form. Grant that the doctor has good reasons to believe I should sign the DNACPR and that I still think my doctor competent and sincere. And set aside for the sake of argument the significant ethics violations involved in a physician making a recommendation on the decision to sign a DNACPR. Do I have good reason to sign?

Not necessarily. If the doctor's good reason to think I should sign is that their employer has imposed perverse targets for DNACPR signatures, then this is not likely to be a good reason for me. Perhaps that is stretching what counts as a good reason for the doctor to recommend signing. But say instead that the doctor recommends signing because they have seen too many patients and families suffer the indignity of a failed CPR attempt at the end of life or that they think it is likely that patients whose lives are saved by CPR would only live for a few days more with broken ribs and severe pain. Though these are still questionable reasons for the doctor, they are not evidently bad reasons for them to believe I should sign the DNACPR. But even if they are good reasons for the doctor to think I should sign, they are not necessarily good reasons for me to think this. For I may, for example, have a deeply held conviction in the value of life or in the purpose of medicine to preserve life at all costs. The doctor's good reasons are not necessarily even reasons that I accept but decide are overridden by other reasons ("I agree, but I consider other things more important"), for I may altogether repudiate the value of a dignified end of life. Thus, even if I am in a position, given the sincerity and competence of the doctor, to believe that they have good reasons to think I should sign the DNACPR, I do not thereby have reason to sign.

An additional condition must hold for it to be rational to trust in the recommendation of an expert. That additional condition is that the person has good reason to believe that the expert bases their recommendation on values that are held by the recipient of the recommendation. To be more specific, a recommendation is based on the values of another when B recommends an action to A because B believes that action is in A's interests. The ambiguity of "interest" here helps to keep this condition flexible enough to cover actions that are valuable for a person in virtue of their desires ("Given you wish to lift heavier weights, I recommend you work on your shoulder mobility.") and actions that are valuable for a person despite their occurrent desires ("Smoking is bad for you; stop it."). This flexibility is important because without it we would be committed to thinking that one can never rationally trust a person who is trying to teach us to change our values. Sometimes we turn to experts not just because they can show us the means to our ends, but because we want their advice on the ends themselves.

We thus reach a more demanding version of a principle for recommendation trust:

(R\*) if A has good reasons to believe B has good reasons to believe that a certain action is in A's interest, then A has good reasons to believe that they should perform that action.

One way in which this is a more demanding principle than that which governs epistemic trust is that it requires that we have good reason to think that the expert issuing the recommendation understands what is in our interest. In some cases, this condition might be easily met. Most of the time when we ask a doctor for a recommendation we expect them to tell us what is good for our health, and we will be happy to defer on precisely what health means and what we should do to preserve it. But even in medical cases we will sometimes encounter situations in which we cannot take for granted that a competent doctor's recommendations will be in our interest. The potential divergence of values in a decision about a DNACPR is one example of this.

#### 4. HOW DO WE CULTIVATE TRUST?

The difference between epistemic trust and recommendation trust has ramifications for the measures appropriate to cultivating trust in experts. Much of the literature theorising ways in which we could help to build trust in science focuses on the sceptical problem posed by the decline of the value-free ideal of science (see, e.g., Anderson 2011; John 2018; Schroeder 2019). In doing so, this work focuses on how we can cultivate epistemic trust. There is no one answer to this challenge for epistemic trust in science, nor is this the only challenge.[10] To keep things manageable, I will focus only on the problem that values in science poses for trust, and to a selection of solutions to this problem. The question I raise is whether any proposed measures to cultivate epistemic trust in experts could also build recommendation trust. The purposes of this are twofold: to demonstrate that not all measures that help to cultivate epistemic trust help to build recommendation trust; and to arrive at a positive—though provisional—proposal for two measures to cultivate recommendation trust in experts.[11]



Consider first a popular argument against the value-free ideal of science, namely the argument from inductive risk (Douglas 2000; John 2016). Evidence gathered to test a scientific hypothesis typically underdetermines whether we should accept or reject the hypothesis because it fails to support deductive inferences either way. This introduces an inductive risk of either false positives or false negatives, depending on the level of caution we exercise in rejecting or accepting hypotheses. The level of caution we exercise cannot be grounded in epistemic reasons, and so we must appeal to non-epistemic reasons—primarily ethical values—to decide how cautious we will be.

Such judgements about the inductive risk we are willing to accept are particularly prevalent in our current emergency, where decisions need to be made quickly and data are scant. The parameters of our epidemiological modelling, for example, will partly be determined by non-epistemic reasons for whether we are more willing to risk overestimates or underestimates in infection and mortality rates. These decisions are sometimes said to be based on a trade-off between COVID-19 deaths now and deaths or decreased wellbeing in the future as a result of the economic impact of social distancing. However, the role of values is likely to be much more complex than this. Our values will affect, among other things, what we consider to be possible responses to a pandemic-generated recession, and political decisions will partially determine the outcomes of a recession, including who bears the brunt of its negative effects. Even setting the terms of the trade-off is influenced by our politics.

In response to such reasons to think science cannot be value free, some have called for greater transparency about the way in which values affect scientific results.[12] Greater transparency promises two benefits. First, transparency encourages novices to believe that scientists are honest about their procedures and findings. This belief in honesty is one of the crucial components in epistemic trust in science. Second, greater transparency could mitigate whatever tendency there might be among the public to think that the collapse of the value-free ideal of science warrants a blanket scepticism about all scientific claims. Values influence science in a variety of ways, some of which merit much less scepticism about the results of science than others. Heather Douglas (2008) has argued that there is a significant difference between direct and indirect influence of values on science. Values play an illegitimate direct role in science when they are used to warrant an empirical claim. Values play a legitimate indirect role in science when they determine what we accept as sufficient warrant for an empirical claim, with the warrant provided solely by epistemic evidence. Thus, for Douglas values should guide our judgements about acceptable levels of inductive risk, but they should not guide our judgement of the probability that an empirical claim is true. If greater transparency reveals that values play a legitimate role in science, then it can help to allay suspicion.

Transparency is by no means an unquestionably effective measure even for building epistemic trust (John 2018; Schroeder 2019). But aside from the problems that arise for using transparency to increase epistemic trust, there are three new problems that arise if we wish to use transparency to build recommendation trust. First, confidence in the sincerity of experts is not enough to meet conditions for rational trust in the

recommendations of experts. An individual expert or an expert community could disclose as much information as possible about their methods, procedures, and even value assumptions that lead to their results, without giving the public reason to accept their recommendations. In order to have recommendation trust in experts, we need to know not just that their testimony about their area of expertise is sincere, but also that their recommendations are in our interest. Our confidence in this is not secured by transparency alone.

Perhaps this requires simply that our attempts to increase public transparency of science focus specifically on the values, goals, and ends that inform expert recommendations, thus allowing the public to judge for themselves whether the values of the experts align with their interests. But this fails to address the other two problems with using transparency to build recommendation trust. The second problem for transparency is that the way that values inform scientific results is very often not as simple as we need it to be in order to have reliable ways of communicating this to non-experts. A basic principle of precautionary reasoning used in epidemiological modelling of infection rate may be simple enough to communicate effectively to non-epidemiologists, but transparency about the political assumptions involved in economic modelling of the effects of the pandemic could be much more difficult to achieve. Indeed, the role that values play in setting the parameters for an expert community may not even be transparent to the experts themselves; not all scientists are philosophers of science. Revealing the values that influence scientific recommendations and expert-led policy might sometimes be possible, but often such transparency is highly impracticable.

But even if we can achieve transparency, it is not clear that it will benefit recommendation trust in the same way it can benefit epistemic trust. This leads to the third problem with transparency. One of the purported benefits of transparency is to show non-experts that values are playing the role they should play in science. Provided that everything is running as it should—that the science in question is trustworthy—transparency can allay suspicion that values play an illegitimate direct role. But this relies on a distinction between direct and indirect use of value judgements that does not apply to expert recommendations. Values must play a direct role in a recommendation because values must feature in the reasoning that supports the recommendation. There is thus nothing to be gained simply by showing non-experts what role values play in an expert recommendation, because there is no relevant distinction between more or less legitimate roles for values to play in such recommendations.

In fact, showing the public that their interests are aligned with the values informing science is not just an unhelpful and often unworkable step, but it is an unnecessary step.[13] For my trust in an expert recommendation to be rational, I need to have well-placed confidence in the alignment of the ends of the recommendation and my own ends. But confidence in this can be secured through means other than first learning about the values that inform the expert recommendations and then comparing them with my own. Instead of allowing the public to examine recommendations through transparency of values, we could instead secure confidence in the alignment of values through building public values into the

procedures used to arrive at recommendations. If I have good reason to believe that my interests have played a significant role in the procedures used to arrive at expert-led policy, then I need not examine each element of the policy and each expert recommendation for alignment with my interests, because confidence in the procedure can provide confidence in the policy outcome.

## 5. DEMOCRATIZED SCIENCE AND ITS LIMITED VALUE FOR RECOMMENDATION TRUST

How might we build public interests into decision-making procedures in a way that supports recommendation trust in experts? Many have proposed a range of measures for democratizing science to build public trust in experts, but, again, not all of these proposals will help us deal with the extra demands when building trust in expert recommendations. One popular option is James Fishkin's deliberative polling (2009).[14] The deliberative polling process reaches policy verdicts through face-to-face deliberation between representative samples of the population, informed by relevant scientific information and bound by norms of appropriate discussion. The role of experts in the process is to explain the relevant scientific consensus, where there is one, and as much as possible the relevant concepts, methods, and standards of evidence used in reaching that consensus. Citizen participants are also informed of the most significant robust opposition to the consensus, with an opportunity to learn about the reasons that could be given for dissent, and to hear more from the experts about their defence against objections. Non-expert participants are thus given opportunities to resolve doubts as much as possible before coming to a collective conclusion about the correct policy decision, informed by the expert advice.

As Elizabeth Anderson has noted, one of the virtues of Fishkin's deliberative polling model is that there is evidence that such a procedure mitigates the effects of cultural cognition, a social psychological phenomenon in which individuals are more likely to accept information that is compatible with their existing values and to deny information that challenges those values (Anderson 2011, 158). Interaction with participants with other political views and a much more extensive interaction with experts than would usually be the case can help to increase the chances that a non-expert will respond to the data regardless of whether they affirm previously held convictions.

But the distinctive challenge of generating recommendation trust in experts is not to overcome cognitive bias. Even if the outgroup prejudice against experts is overcome through deliberative polling, a non-expert participant could still be warranted in refusing to accept what the experts tell them when the experts are issuing recommendations and not just giving factual testimony. Deliberative polling could successfully disabuse a citizen of epistemic prejudice against expert testimony without thereby giving them a reason to accept that the same epistemic-trustworthy expert is issuing recommendations that are in their interest. If this is the case, as I have argued above, the citizen will have good reason to epistemic-trust the expert without having good reason to recommendation-trust them.

Perhaps deliberative polling is used to tackle the problem not by showing citizens they have good reason to recommendation-trust, but instead by eliminating the need for recommendation trust in the first place. It could be that in deliberative polling experts no

longer play the role of issuing recommendations for policy, and instead they simply provide information for the non-expert deliberators to use in their collective policy verdicts. But though this might be valuable so far as it goes, it leaves us without a solution for the emergency situations in which expert-led policy, trusted by the public, is desirable and cannot be substituted with policy decisions made solely by novices. As I have argued in Section 2, this is precisely the kind of scenario we find ourselves in during the COVID-19 pandemic.

There is another way to secure confidence in the alignment of public interest and expert recommendation that relies on neither transparency nor novice participation in policy deliberation. We can look instead to the way in which science-led policy is communicated to the public. In addition to Fishkin's deliberative polling model, Anderson (2011) proposes two other means by which we can overcome cultural cognition and persuade the public to become better disposed to trusting experts where indeed they have good reason to do so.

The first, "expressive overdetermination," is borrowed from Braman, Kahan, and Grimmelmann (2005, 297): investing policy with multiple meanings such that it can be accepted from diverse political perspectives. Braman et al. cite French abortion law reform as an example of this. After decades of disagreement, France adopted law that made abortion permissible provided the individual has been granted an unreviewable certification of personal emergency. Evidence that such an approach would be effective pre-existed the new law, but such evidence had proved unconvincing to concerned parties. This new policy was sufficiently polyvalent to be acceptable to the most important parties to the debate; religious conservatives understood the certification to be protecting life, while pro-choice advocates saw the unreviewable nature of the certification as protection for the autonomy of women. With this framework in place, acceptable to all, opposing sides to the debate were able to converge on the details of implementation (ibid., 298).

Anderson's second alternative is to recruit spokespersons who are identifiable to diverse groups as similar to them in political outlook. The strategy is to overcome cultural cognition by convincing citizens that the relevant scientific consensus is likely not to be a threat to their values because that same consensus is accepted by those with similar values. Anderson cites Barack Obama establishing links with Evangelical Christians such as Rick Warren, one of the 86 evangelical leaders who had signed the Evangelical Climate Initiative two years before the beginning of Obama's presidency (Frazier-Crawford Boerl 2010, 153). She suggests that Obama's public association with Warren was an attempt to win over conservative Christians to his climate-change policy. Braman et al. call this "identity vouching" (2005, 297).

Can these strategies for communicating science to sceptical publics also work for building recommendation trust? I believe they can, with the right caveats. First, given that we are looking for measures to build well-placed recommendation trust, it must be the case that the evidence-based policy that is communicated through expressive overdetermination and identity vouching is in fact in the interests of the public whose trust we are soliciting. These strategies could also be used to mislead the public into accepting recommendations from

experts that are not in their interest, if, for instance, the spokespersons are dishonest about whether this is the case. But this would not be a strategy to cultivate well-placed recommendation trust, because the public would not have good reason to trust such recommendations.

Second, if we are to use these strategies specifically for recommendation trust, the goal of the strategies is different, and this will likely affect the details of how we implement them. Expressive overdetermination and identity vouching support epistemic trust by encouraging citizens to believe expert-testimony that would otherwise be resisted as a threat to those citizens' values. The same strategies support recommendation trust only if they encourage citizens to see that the relevant expert recommendations are in their interests. But it is not difficult to imagine how these strategies could do this. Braman et al. (2005, 297) suggest that French religious conservatives were more willing to accept the evidence in favour of abortion law reform once they saw that the law affirmed the sanctity of human life. However, it is also likely that this adjustment gave those same religious conservatives good reason to believe that this policy was in their interests ("If you want a policy that respects the sanctity of life, this is the policy for you."). Similarly, evidence-led policy could be presented to the public by spokespersons whose interests are already thought to align with the relevant members of the public. Thus, an American Evangelical Christian might think that if recommendations to cut carbon emission are good for Rick Warren, then they are good for them too.

With these caveats, using these communication strategies to cultivate recommendation trust is, I submit, preferable to other proposals. Expressive overdetermination and identity vouching avoid the problems of transparency measures: they demonstrate specifically that public interests align with expert recommendations, and they do so without relying on novices (or even experts themselves) understanding the complex ways in which expert values shape scientific recommendations. And these communication strategies also provide an approach that is preferable to deliberative polling in times of crisis. When an emergency situation precludes a lengthy process of expert-informed citizen deliberation, expert-led policymaking might be necessary. Where this is the case, trust must be won by demonstrating to the public that expert recommendations are in their interest.

## 6. CONCLUSION

The account I have given offers reasons for a cautious optimism about the possibility of building well-placed trust in science-led policy. I am afraid I will conclude on a more pessimistic note. I have argued for the view that there are additional challenges to rational trust in experts when those experts are leading public policy and not just providing policy-relevant information. The additional demands of this kind of trust boil down to three claims that I have argued for: recommendation trust is different from both epistemic and practical trust; the conditions for rational recommendation trust are more demanding than the conditions for rational epistemic trust; and many measures that have been proposed to cultivate trust in experts do not help to cultivate recommendation trust in science-led policy. The measures that can help involve communicating science-led policy to the public in

a way that allows the public to see, where it is indeed the case, that it is in their interest to follow the policy.

But the success of such a strategy is likely to be limited by factors that I have not had space to address. One such factor is the fragility of public confidence in the very idea of science-led policy. Where governments present their policy as technocratic, only for their claim to be “following the science” to be exposed as false, citizens will have good reason to be suspicious of future policy that is sold to them on a similar basis. Another significant factor is the broader climate of public trust, most notably the levels of public trust in politicians and government. Perhaps science-led policy is the product of experts and politicians working together. Perhaps science-led policy is the product of a full-blown technocracy, where experts become our political leaders. Either way, once science is corrupted by politics, its authority is vulnerable to public trust in the political system. Rebuilding this kind of trust is a much more difficult task.

## REFERENCES

- Anderson, Elizabeth. 2011. "Democracy, Public Policy, and Lay Assessments of Scientific Testimony." *Episteme* 8 (2): 144–64.
- Atchison, Christina, Leigh Bowman, Charlotte Vrinten, Rozlyn Redd, Philippa Pristera, Jeffrey W Eaton, and Helen Ward. 2020. "Perceptions and Behavioural Responses of the General Public During the COVID-19 Pandemic: A Cross-sectional Survey of UK Adults." medRxiv doi: 10.1101/2020.04.01.20050039
- Braman, Donald, Dan Kahan, and James Grimmelman. 2005. "Modelling Facts, Culture, and Cognition in the Gun Debate." *Social Justice Research* 18 (3): 283–304.
- Brennan, Johnny. 2020. "Can Novice Trust Themselves to Choose Trustworthy Experts? Reasons for (Reserved) Optimism." *Social Epistemology* 34 (3): 227–40.
- Burge, Tyler. 1995. "Content Preservation." *Philosophical Issues* 6: 271–300.
- Croce, Michel. 2019. "On What it Takes to be an Expert." *Philosophical Quarterly* 69 (274): 1–21.
- Douglas, Heather. 2000. "Inductive Risk and Values in Science." *Philosophy of Science* 67 (4): 559–79.
- . 2008. "The Role of Values in Expert Reasoning." *Public Affairs Quarterly* 22 (1): 1–18.
- Fishkin, James. 2009. *When the People Speak: Deliberative Democracy and Public Consultation*. Oxford: Oxford University Press.
- Frazier-Crawford Boerl, Christopher Wayne. 2010. "American Evangelicals and the Politics of Climate Change." *St Antony's International Review* 5 (2): 147–63.
- Fricker, Miranda. 1998. "Rational Authority and Social Power: Towards a Truly Social Epistemology." *Proceedings of the Aristotelian Society* 98 (2): 159–77.
- Hardwig, John. 1985. "Epistemic Dependence." *Journal of Philosophy* 82 (7): 335–49.
- . 1991. "The Role of Trust in Knowledge." *Journal of Philosophy* 88 (12): 693–708.
- Holton, Richard, and Jacopo Domenicucci. 2017. "Trust as a Two-place Relation." In *The Philosophy of Trust*, edited by Paul Faulkner and Thomas Simpson, 149–260. Oxford: Oxford University Press.
- Ipsos MORI. 2019. "Trust in Politicians Falls Sending Them Spiralling Back to the Bottom of the Ipsos MORI Veracity Index." <https://www.ipsos.com/ipsos-mori/en-uk/trust-politicians-falls-sending-them-spiralling-back-bottom-ipsos-mori-veracity-index> (Accessed 15 May 2020)
- John, Stephen. 2016. "From Social Values to P-Values: The Social Epistemology of the Intergovernmental Panel on Climate Change." *Journal of Applied Philosophy* 34 (2): 157–71.

Accepted for publication in Kennedy Institute of Ethics Journal as part of the special double issue on Ethics, Pandemics, and COVID-19.

——— 2018. “Epistemic Trust and the Ethics of Science Communication: Against Transparency, Openness, Sincerity and Honesty.” *Social Epistemology* 32 (2): 72–87.

Kappel, Klemens. 2014. “Believing on Trust.” *Synthese* 191 (9): 2009–28.

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

Kukla, Rebecca. 2007. “How do Patients Know?” *The Hastings Center Report* 37 (5): 27–35.

O’Neill, Onora. 2020. “Trust and Accountability in a Digital Age.” *Philosophy* 95 (1): 3–17.

Pew Research Center. 2019a. “Public Trust in Government: 1958–2019.” Accessed May 15, 2020. <https://www.people-press.org/2019/04/11/public-trust-in-government-1958-2019/>.

——— 2019b. “Trust and Distrust in America.” July. Accessed May 15, 2020. <https://www.people-press.org/2019/07/22/trust-and-distrust-in-america/>.

Rolin, Kristina. 2002. “Gender and Trust in Science.” *Hypatia: A Journal of Feminist Philosophy* 17 (4): 95–120.

Schroeder, S. Andrew. 2019. “Democratic Values: A Better Foundation for Public Trust in Science.” *The British Journal for the Philosophy of Science* axz023. <https://doi.org/10.1093/bjps/axz023>.

Watson, Jamie Carlin. 2020. “Hunting the Expert: The Precarious Epistemic Position of a Novice.” *Social Epistemology Review and Reply Collective* 9 (4): 51–58.

Wilsdon, James, and Rebecca Willis. 2004. *See-through Science: Why Public Engagement Needs to Move Upstream*. London: Demos.

Wylie, Alison. 2003. “Why Standpoint Matters.” In *Science and Other Cultures*, edited by Robert Figueroa and Sandra Harding, 26–48. New York: Routledge.



## NOTES

1. Ipsos MORI data on public trust in the UK in November 2019 showed that 14% of the public trusted politicians to tell the truth, matching previous recorded lows in 2016, 2011, and 2009 (Ipsos MORI 2019). Data from the Pew Research Center in April 2019 showed that 17% of respondents in the US trusted the government to do what is right at least “most of the time” (Pew Research Center 2019a), but this same figure has passed 50% only three times since 1972 (ibid.).
2. Ipsos MORI’s data show that 86% of respondents said they generally trusted professors to tell the truth, 84% scientists (Ipsos MORI 2019). Pew Research Center data from 2018 showed that 83% of respondents said they have at least a fair amount of confidence in scientists, with very little (1%) variation between age groups (Pew Research Center 2019b, 19).
3. What is an expert? The account of trust in experts I give here will remain neutral on debates about how best to define experts. However, I am inclined to what Michel Croce (2019) calls novice-oriented accounts because the experts I am concerned with are those that play the role of providing information and recommendations to non-experts.
4. Where decisions really are exclusively within the remit of politicians, science-led policy can only be such if politicians trust experts. The norms governing when it is reasonable for politicians to trust experts, and the measures that we can take to build this trust, might well differ from those relating to public trust in science.
5. It is not universally accepted that trust is tripartite. For more see Holton and Domenicucci (2017).
6. The domain restriction can also reflect a principle that underpins some standpoint epistemology: some facts about, for instance, social and political injustice are better understood by those who suffer those injustices (Wylie 2003). Thus marginalized groups might be more epistemic-trustworthy in relevant domains.
7. Hardwig is also not without his critics (see e.g., Kappel 2014; Rolin 2002). Nonetheless, the criticisms of Hardwig that I am aware of do not give us reason to adjust his principles in ways that would threaten the distinction in norms of rational trust that I am arguing for. Demonstrating this would, I suggest, take me beyond the scope of this paper.
8. Note my argument in this section is not an argument against Hardwig, because Hardwig does not claim his principle would apply to anything other than epistemic trust.

9. Does this mean that Hardwig's principle (T) is not true? (R) appears to be a more specific version of (T) insofar as it specifies the nature of the belief professed by B and accepted by A. But if (R) is entailed by (T), and (R) is false, then (T) must also be false. Given that the truth of (T) is not at stake in this paper, I will remain agnostic on this issue, though I provisionally suggest that (T) might be saved were we to specify that B's belief is about non-moral facts—this adapted (T) would not entail (R). I thank an anonymous reviewer for raising this.
10. Other concerns include the difficulty for novices of identifying experts. See e.g., Brennan 2020 and Watson 2020.
11. I am also setting aside matters of practical trust in experts, not to mention practical trust in government.
12. See e.g., Kitcher (2001, chapter 6) and Wilsdon and Willis (2004).
13. It could also be a dangerous step, as Schroeder (2019) has argued, insofar as it could encourage non-experts to be selective about the experts they prefer to follow, and as a result it might politicize science such that scientists and the public become divided along ideological lines.
14. This has been endorsed as a means to build trust by Anderson (2011) and Kitcher (2011, chapter 8).