Covert Value Judgments in Expert Testimony

Alex Worsnip

Forthcoming in Anthony Nadler, Molly O'Rourke-Friel & Doron Taussig (eds.),
Truth After Post-Truth: Finding A Way Forward, UMass Press
Please cite published version if possible

1. Prelude: experts and values in the COVID pandemic

Nothing has done more to put the role of experts in public life on the agenda than the COVID pandemic. On the one hand, a sizeable portion of the public, and some politicians—those who believed and pushed conspiracy theories that the COVID vaccine was killing people in droves, or even more wildly, that it was implanting microchips into its recipients—failed to adequately heed expert opinion. But on the other hand, a number of commentators and academics have recently suggested that a different contingent of the public, as well as many journalists and politicians, were *too* credulous of the experts on issues like masking, the origins of the virus, and—especially—school closures.[1]

Part of their point has to do with the ways in which experts are subject to the same pathologies that bedevil human psychology more generally: experts can be influenced by groupthink, tribalism, and political pressure; and are prone to overconfidence and intolerance of dissent. They argue that these tendencies led experts to rush too quickly to consensus about epidemiological matters surrounding COVID and policy responses to it, on the basis of insufficient evidence. However, another part of their point is about whether *even in principle*, we can or should expect scientific experts—even ones who handle the evidence in a maximally responsible way—to dictate matters of policy. Their point is that any question about what policy response to adopt unavoidably turns not just on scientific questions but also on questions of *value*. In the case of school closures for example, it depends on how we should strike the balance between preventing the spread of the virus, on the one hand, and number other costs such as the setbacks to learning that resulted from school closures, especially for those from disadvantaged backgrounds;[2] tolls on children's mental health from being kept at home; and more. These are difficult tradeoffs—and, crucially, as value questions, they are not themselves resolvable solely by science.[3]

Some of the public discourse around expert opinion in the pandemic belied this point. Most obviously, the slogan "follow the science"—and related proclamations like California governor Gavin

Work on this paper was supported in part by a subgrant from Arizona State University's *Humility in Inquiry* project, funded by the Templeton Foundation; I gratefully acknowledge the support. While this particular paper is solo-authored and the views expressed should not be attributed to them, I would like to thank my fellow team members on the grant project—Aditi Ahuja, Kurt Gray, Jeff Greene, Devin Lane, Giulia Napolitano, and Sam Pratt—for many excellent conversations about deference to experts, which have influenced my thinking enormously. For helpful feedback and questions about this paper (or related, now-abandoned predecessors to it), I would also like to thank audiences at King's College London, Tulane University, a workshop on "epistemic norms for the new public sphere" at the University of Warwick, and the Hettleman talks at UNC Chapel Hill, as well as Doug MacKay, Daniel Muñoz, the editors of this volume, and two anonymous reviewers. Special thanks to Jonathan Leader Maynard and Jessica Marsden for many stimulating conversations about this topic.

[1] See for example Macedo & Lee (2025), Russell & Patterson (2025), and Zweig (2025).

[2] See Mervosh, Miller & Paris (2024).

[3] Cf. Zweig (2025: 107; 285).

Newsom's claim that "science will determine when a school can be physically open"[4]—suggested that science alone could just directly instruct us on what policies to adopt, obscuring from view the ineliminable role that values play in any such policy decision, and the need for politics to serve as an arena in which to resolve such disputes.[5] Similarly, the suggestion that "we should leave it to health experts to tell us when the time is best to open up school buildings"[6] raises the question: why would health experts—who are focused almost singularly on the prevention of the spread of disease,[7] and have no special expertise at all on the educational or distributive justice effects of school closures—be in a position to decide how the tradeoff between the two should be struck? Indeed, why should *any* expert (or group of experts) be given the sole authority to decide questions that turn in part on value judgments?

No matter what politicians like Newsom may suggest, policy recommendations—that we *ought* to close schools, say—clearly involve value judgments. But scientists and other experts aren't always making policy recommendations. Sometimes they are—or seem to be—just telling us the state of the science. Here too, though, values can be built into the language experts use, just in a less overt way. Sticking with the pandemic example, consider expert pronouncements about what is "safe"—for example, about whether it is safe for schools to reopen, or indeed whether it is safe to go for a run with friends or have a drink with your neighbors on their back porch. Whether something is safe might seem to just be a scientific matter. But this appearance is misleading. Safety is scalar. Some behaviors are safer or riskier than others, but almost no behavior is absolutely guaranteed to lead to your immediate death, and similarly, almost no behavior is *completely* safe in the sense of carrying no risks. The same goes for policies. (Indeed, at *any* time, having schools open carries *some* risk of viral transmission—of flu, for example—a very small amount of which will lead to deaths.) In deeming something 'safe' in *absolute* terms, then, we seem to be saying something like this: the risks are small enough that the benefits outweigh them. But then pronouncements about what is safe *do* encode value judgments after all, about what kinds of benefits justify what degree of risk. Again, these are not purely scientific questions.[8]

Yet experts and others did make binary pronouncements about whether (for example) opening schools was "safe."[9] And these were sensitive not just to perceived risks but also to perceived benefits:

---

[4] https://x.com/gavinnewsom/status/1284213443457806337

[5] Cf. Mercuri (2020); Leonhardt (2022); Macedo & Lee (2025: 9-10); Zweig (2025: 117).

[6] Quoted from a joint statement from the American Academy of Pediatrics, American Federation of Teachers, National Education Association and School Superintendents Association, July 10, 2020.

[7] Interestingly, so much was recently admitted very candidly by Francis Collins, the Director of the NIH during the first two years of the pandemic. In reflecting on the public health establishment's response to COVID, he said:

> "if you're a public health person and you're trying to make a decision, you have this very narrow view of what the right decision is. And that is something that will save a life; it doesn't matter what else happens. So you attach infinite value to stopping the disease and saving a life. You attach zero value to whether this actually totally disrupts people's lives, ruins the economy, and has many kids kept out of school in a way that they never quite recover from. […] This is a public health mindset and I think a lot of us involved in trying to make those recommendations had that mindset and that was really unfortunate. It's another mistake we made." [quoted in Schorr 2023]

See also Carroll (2022); Zweig (2025: 255).

[8] Again, the point seems poorly recognized by politicians. Witness statements from former New York governor Andrew Cuomo (cited in Zweig 2025: 223-4): "Is it safe? You can determine that by science." "You reopen if it's safe. How do you know if it's safe? You look at the data." "It's purely on the numbers, it's on the science."

[9] See e.g. Graham (2020).

it was deemed "unsafe" for children to go to school, but not "unsafe" for (for example) "essential" workers—including food service workers and delivery workers—to go to work.

Let's call cases like these—where an expert says something that *appears* to be purely scientific, but tacitly encodes value assumptions—instances of *covert value judgments in expert testimony*.[10] In the remainder of this essay, I'm going to be exploring the complications that covert value judgments in expert testimony pose for the practice of *deference* to experts, which occurs when a layperson believes an expert's testimony on their say-so. Many philosophers (and others) have suggested that laypeople ought to defer to experts a lot more than they do. But as I've already been hinting, and will now argue in much more detail, this advice can become problematic when there are value judgments involved. My aim is not to argue that we should never defer to experts, but to raise to prominence a danger of deference (and a complication for its rationality) that I believe has received insufficient attention.[11] Awareness of this problem is important not just for laypeople considering whether to defer, but also for experts themselves. Accordingly, I will ultimately conclude with some reflections on how experts can be appropriately cognizant of the problem in issuing testimony and advice, and in so doing, make this testimony more worthy of deference.

2. The philosophers' advice: deference to experts

It's often said that distrust of experts is one of the deepest socio-political problems of our current age; perhaps, an aspect of a more general "epistemic crisis."[12] Empirical work does suggest fairly widespread (and increasing) distrust of establishment experts.[13] While this distrust is most pronounced among those on the right, there are some issues on which those on the left distrust the experts—notably, economics. For example, those on the left often reject the expert consensus[14] that developed countries dismantling protectionist measures is economically beneficial for both developed and developing countries.[15]

Many contemporary epistemologists (i.e., philosophers who specialize in the study of knowledge and rational belief) agree that this distrust of experts, and refusal to accept their claims, is not only socially harmful but also irrational. An increasing number of epistemologists (Huemer 2005; Zagzebski 2012; Ahlstrom-Vij 2015; Grundmann 2021; Matheson 2024) advocate a radical form of *deference* to experts, whereby on complex scientific and policy issues, one does not even try to evaluate the primary data and evidence for oneself, but rather quasi-automatically accepts the expert consensus. The key argument for this prescription starts with the claim that given the complex and technical

[10] Following philosophical tradition, I'm using the term 'testimony' here to include a wide variety of assertions that speakers make and invite their audiences to accept on that basis. I am not using 'expert testimony' in the narrow sense restricted specifically to courts of law.

[11] This is not to say that my argument is entirely without precedent. I will be building on insights of others, including Douglas 2008; Elliott & Resnik 2014; Bennett 2020: esp. 251-2; Harvard et al. 2021; Duijf 2021: esp. 9290; and Barnes forthcoming. But my precise framing of the problem as an issue for deference is, to my knowledge, largely original.

[12] E.g., Nichols (2017); Brooks (2020).

[13] See e.g. Kennedy, Tyson & Funk (2022); Kennedy & Tyson (2023).

[14] See Whaples (2006). Some recent work in economics (e.g. Autor et. al 2016) has put some pressure on the contention about the effect on developed countries, but the positive effect on developing countries is still very widely accepted.

[15] See e.g. Rankin (2001: 359); Guisinger (2017: 176-7). This consensus is also rejected by many on the contemporary populist right, and seems to be accepted most by centrists and moderates.

nature of the evidence pertaining to these issues, ordinary citizens typically lack the competence to understand this evidence and reach a well-founded evaluation of the conclusions it supports. For example, understanding climate data and drawing conclusions from it requires a level of technical proficiency that laypeople simply lack. When we lack this competence, these epistemologists claim, it is irresponsible to even try to gather and evaluate the evidence for ourselves—or, in a contemporary phrase, to "do our own research" (Levy 2022; Ballantyne et al. 2024). Instead, we should simply defer to those who do have the relevant competence. By doing so, we stand a better chance of arriving at a true belief.

This development in epistemology is a significant piece of intellectual history because it represents a turn away from a venerable tradition that stresses the primacy of intellectual autonomy—of interrogating, rather than uncritically accepting, the claims of those in authority.[16] This normative orientation continues to be tacitly assumed in the rhetoric and practice of much contemporary Western liberal education, which often presents itself as equipping students to "think for themselves," "challenge orthodoxy," or "reach their own conclusions."[17] And indeed, some philosophers continue to push back against the enjoinment to uncritically defer to experts, and to stress the value of at least a modest degree of epistemic self-reliance (Hazlett 2016; Lackey 2018, 2021). Others have identified ways in which deference to experts is not a panacea for our epistemic troubles—for example, by pointing out that figuring out *which* experts to defer to (or perhaps, who even counts as an expert in the first place) can sometimes be almost as hard as trying to reason for oneself about an issue (Millgram 2015: ch. 2, appendix A; Nguyen 2020; Ballantyne 2022).

As I've said, my own focus is going to be a reason for caution about deference to experts that is distinct from all those just mentioned: the problem of covert value judgments in expert testimony. To set the stage for this, let's examine a crucial distinction.

3. Descriptive vs. normative (value) judgments

Philosophers and social scientists frequently distinguish descriptive judgments and normative judgments. Very roughly, descriptive judgments are judgments about how the existing world *is*. Most descriptive judgments are empirical: they are to be settled through empirical observation and scientific methods. By contrast, normative judgments are judgments about how the world *ought* to be, what people ought to do (or believe, intend, hope for, etc.), or whether (aspects of) some existing or possible state of affairs is good or bad, valuable or disvaluable. Normative judgments include, though are not exhausted by, moral judgments. People's normative judgments are often also called their "values" or "value judgments."[18] I'll use the term 'value judgments' in what follows since it's more common both in the literature on expertise and in ordinary speech.

---

[16] A classic historical source is Kant (1784/1996). For an interesting discussion of some similar ideas in classical Chinese philosophy see Tiwald (2023).

[17] For example, here is the oft-quoted American Association of University Professors (AAUP) 1915 Declaration of Principles on Academic Freedom and Academic Tenure: "the university teacher […] should, above all, remember that his business is not to provide his students with ready-made conclusions, but to train them to think for themselves."

[18] Some philosophers use 'normative judgment' to refer to judgments about what people ought to do and 'value judgment' to refer to judgments about what is good and bad. Here I will use both terms (interchangeably) in a more capacious sense that includes both of these things.

To take an example that might arise in the context of expert testimony, the judgment that broccoli boosts immune health is a descriptive judgment, whereas the judgment that you *should* eat more broccoli is a value judgment. Of course, like many other value judgments, the value judgment that you should eat more broccoli might be partly *based* on the descriptive judgment that broccoli boosts immune health. But if it is, it is also—perhaps tacitly—based on further value judgment(s), which might include something like: "even if you dislike broccoli, boosting immune health is more important than avoiding the unpleasant taste." While this judgment may be very plausible, the crucial point for our purposes is that it must be added to the descriptive judgment that broccoli boosts immune health in order for the specific value judgment that you should eat more broccoli to logically follow. When someone makes a specific value judgment A on the basis of a descriptive judgment B, but some other value judgment C is required to derive the specific value judgment, we can say that they (again, perhaps tacitly) presuppose or assume C.

Philosophers sometimes call our most fundamental value judgments—ones that, plausibly, are not themselves based on *any* descriptive judgments—the "pure" value judgments, in contrast to more specific or "impure" value judgments (McGrath 2009). Pictorially, this model can be represented thus:

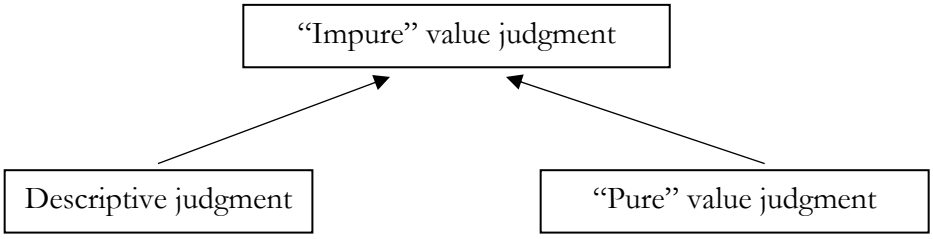

This model also accords with the familiar philosophical claim that one cannot derive an 'ought' (i.e., a normative or value judgment) solely from an 'is'. Value judgments are invariably based on further, more fundamental, value judgments.

4. Deference about value judgments?

Strikingly, while (as we've already seen) many philosophers think you're required to defer to experts about *descriptive* matters, many philosophers think that in sharp contrast, there's something deeply problematic or strange about deferring on value judgments, especially pure value judgments. The point is most often made with respect to moral judgment (Driver 2006; Hills 2009; McGrath 2009). Suppose, for example, that you are wondering whether meat-eating is morally permissible. And suppose you ask someone whether it is morally permissible and they say it isn't. It would be odd, these philosophers think, if you now just believed that meat-eating is not morally permissible purely on their say-so, without asking them for the reasons and trying to understand them. Note that this contrasts with many descriptive cases: it's *not* odd, the thought is, to simply accept the testimony of a climate scientist when they say that global temperatures are rising due to human activity without asking for or trying to comprehend the evidence on which they base this judgment.

A bit of qualification is required about the moral case. You might defer to someone on whether meat-eating is morally permissible without any oddness under certain conditions. Whether meat-eating is morally permissible is, arguably, an impure moral question: it depends on certain descriptive facts—

perhaps, on whether (certain) animals feel pain, or the conditions under which they are bred, or whether individual meat-eating causally contributes to the production of meat for consumption. Suppose that the reason why you are currently uncertain about whether meat-eating is morally permissible is *just* that you are uncertain about these descriptive facts. You might then defer to someone who is an expert about these descriptive matters about whether meat-eating is morally permissible, *if* you assume that this person shares your values, your pure moral judgments. Here, you are effectively deferring to them about whether the descriptive facts are such that (by your shared values) meat-eating is morally permissible. This is not odd. What is supposed to be odd, though, is deferring to pure moral judgments, or to impure moral judgments when the pure moral judgments they depend on are in question or not known to be shared.

Different philosophers have given different explanations of why this is odd, and in particular of why it is odder than deference about descriptive matters. An immediate candidate explanation might be that we can't gain moral knowledge by testimony. But participants in the literature almost universally reject this explanation. Instead, they appeal to other explanations: for example, that moral deference is problematic is because we can't gain moral *understanding* by testimony (Hills 2009), or because there is something distinctively *morally* problematic about deferring on moral matters (Hopkins 2007; Crisp 2014). Others think that the difference is simply due to the fact that while there are experts about (many) descriptive facts, there is no such thing as a moral expert: (pure) moral thinking is equally accessible to everyone, and requires no special expertise or access to any particular evidence (Williams 2005: 205; Davia & Palmira 2015). Still others think that it is because there are no objective moral facts to defer about; morality is relative to individual judgment or moral judgments are expressions of personal preferences (McGrath 2011, though she does not definitively endorse this explanation). And still others that think that in the final analysis deference about moral matters is not odd after all; or, at least, that even if it seems a little strange, there is nothing truly wrong or irrational about it (Sliwa 2012; Enoch 2014; Wiland 2017).

For our purposes, we do not need to settle this debate, for the crucial point is this: even if there are experts about morality (and, more broadly, matters of value, and even if it's rational to defer to *them*, there is little reason to think that those with expertise about descriptive (e.g., scientific) matters have special moral expertise about related (pure[19]) matters of value. Return to the broccoli example. A doctor or nutritional scientist might be an expert on whether broccoli boosts immune health. But do they really have any special expertise on the normative question of whether boosting your health is more important than avoiding foods you dislike? I think they do not: this is a value judgment that doctors have no special authority with respect to.

The broccoli example may seem silly or trivial, but we don't have to change it too much to get something that is more interesting. For example, a doctor might be an expert on the negative effects of consuming fatty foods or alcohol and how they can shorten your length of life. But there's actually,

[19] The 'pure' qualification is important. Experts will of course have relevant empirical expertise that informs their *impure* value judgments. Relatedly, they may have other forms of experience other than their technical training that is relevant. For example, a doctor may have interacted with many patients in the patient's situation, and may be able to make a better guess about what a patient might prefer later down the line than the patient herself can. (Thanks to the editors of this volume for this example.) This is all consistent with the doctor having no special expertise with respect to *pure* matters of value.

a genuine, non-obvious normative question about whether a long life is necessarily better than a shorter life where you get to eat and drink what you like. And doctors have no special expertise with respect to that.

Let's consider some higher-stakes examples. Economists may be experts with respect to the descriptive question of which policies will boost overall GDP, but they have no special expertise with respect to the normative question of how important boosting overall GDP is relative to ensuring equality (cf. Barnes forthcoming). To return to our opening case, epidemiologists might be experts about what measures most effectively combat the spread of infectious diseases, but they have no special expertise with respect to the normative question of the relative importance of avoiding the spread of these diseases as compared with other socially desirable outcomes such as presenting social isolation and mental health crises, the educational benefits of in-person instruction, and so on.[20] Climate scientists are experts with respect to whether and why climate change is happening, but they have no special expertise about our moral responsibilities to future generations and how to weigh these against short-term economic pains (cf. Broome 2012).

It's not, of course, that scientists and other experts should be forbidden from making judgments about these matters of value. Scientists are human beings too, and it is part of the life of a human being to make value judgments. What is problematic is their (perhaps tacitly) claiming expertise about such matters, or claiming that they should be *deferred* to about them. This is an instance of what Nathan Ballantyne (2019) calls "epistemic trespassing," where someone pronounces about a field outside their genuine field of expertise as if they were an expert. Moreover, even when scientific experts do not claim expertise about matters of value, ordinary people may treat them as having expertise about those matters, perhaps in part because they don't clearly distinguish between descriptive matters and related matters of value. This too is problematic, if scientific experts have no genuine expertise about these matters of value.

5. The problem of covert value judgments in expert testimony

So far, we might think insofar as there is a problem here, the solution is obvious. We should just defer to the technical or scientific experts about descriptive judgments (or, more precisely, the domain of descriptive judgments that they are experts about), but not about value judgments.

The problem, though—which our opening case-study of COVID already illustrated—is that almost by their very nature, scientific or policy *advice* and *recommendations* necessarily incorporate both descriptive and value judgments. For pieces of advice or recommendations are claims about what we *ought* to do—specifically, they tend to be what we earlier called "impure" value judgments, that themselves depend upon both descriptive and (perhaps tacitly) pure value judgments. Return to the school closure example. Suppose an epidemiologist says, in the middle of the COVID-19 pandemic, that we should keep schools closed rather than reopening them. This is an impure value judgment,

[20] They also plausibly lack special expertise about *descriptive* questions about what the effects of (say) school closures on outcomes such as educational attainment and mental health would be, since these are not matters of epidemiology. In pronouncing epidemiologists the primary or in some cases *only* experts about lockdown measures, then, there is something of implicitly privileging of one kind of relevant expertise over another. Thanks to an anonymous reviewer for urging this point.

partly based on descriptive judgments about the effects of reopening schools on the spread of COVID-19, but also partly based on further value judgments about the relative importance of preventing this spread as compared with the benefits of in-person instruction (on educational outcomes, on mental health, and so on). Thus, if we want experts to be in the business of giving *advice* at all—rather than just giving entirely neutral statements of the facts without making any consequent recommendations—it is hard to see how to defer to this advice without, in effect, deferring to them on both descriptive and value matters. We seem to face the choice of either deferring on both or deferring on neither—neither of which is ideal.

The problem is further intensified by the fact that often, the value judgments lurking in the background of experts' testimony may be unobvious, or covert. Let's explore a few ways in which this can be so, roughly from the most obvious to the most subtle.

First, and most simply, laypeople might simply not appreciate the point I just made, that there are always value judgments in the background of advice. To return to the mantra of "following the science," some people may be in the grips of the view that science *can* on its own—without any further value judgments—just dictate what ought to be done. In our earlier terms, they tacitly assume that an 'ought' *can* be derived from an 'is'. Those who don't see what's wrong with this picture may *unwittingly* defer to experts' value judgments, simply just not realizing that these judgments are there in the background of their advice.

Second, even when we do appreciate that there are value judgments in the background of expert testimony, it is sometimes impossible to know *which* value judgments are in the background. Return to the example of the epidemiologist who says we should keep schools closed. This advice seems to reflect some sort of judgment that the effects of re-opening schools on the spread of COVID-19 will be severe enough to outweigh the benefits of in-person instruction (both for students' educations and for their mental health). But the advice itself doesn't make clear what kind of value judgment about the relative priority of these goals is being made. It could be that the epidemiologist accepts that it can be legitimate to accept a moderate degree of spread of COVID-19 in order to secure the benefits of in-person instruction, but thinks that in these conditions the effect of re-opening on the spread of COVID-19 will be truly catastrophic. Or it could be that the epidemiologist thinks that prevention of the spread of COVID-19 should be given absolute (or almost absolute) priority over the benefits of in-person instruction, and so that even a modest increase in the spread of COVID-19 (as a result of re-opening) cannot be tolerated. We cannot tell, just from the advice, what value judgment is being made; and thus we cannot be sure whether we share the value judgment and hence can rationally defer to the overall advice.

Third, there are judgments involving concepts that look descriptive but are in fact (partly) value-laden, and thus that don't wear their value-ladenness "on their sleeves."[21] I've already suggested that 'safe' is like this. Just for a change, let me illustrate that via a slightly different context: the WHO's recent proclamation that "there is no safe level of alcohol consumption."[22] It might seem that judgments about what is safe, medically speaking, are straightforward descriptive judgments. But as

[21] In a previous paper (Worsnip 2017), I called these judgments—ones that are presented as descriptive but are in fact normative—"cryptonormative judgments."

[22] https://twitter.com/WHOWPRO/status/1655041779719364609

several commentators noted in the wake of the pronouncement, if 'safe' means 'completely risk-free', then it would also be true that there is "no safe level" of crossing the road or driving. Yet the WHO doesn't say that. This suggests that they are tacitly making a value judgment that those things are *worth* the risk whereas the enjoyment of alcoholic beverages (even in moderation) isn't. In that case, in their usage, 'safe' doesn't mean 'completely risk-free' but rather something more like 'low-risk enough that the benefits are worth it'. But then the concept of something's being 'safe' is actually covertly normative or value-laden, and their judgment that there is no safe level of alcohol consumption is a covert value judgment. Again, WHO officials have no special authority with respect to the value judgment that the risks of moderate alcohol consumption aren't worth the enjoyment of it, yet they presented themselves as experts on this matter giving official advice.

Fourth and finally, even when we confine ourselves to judgments that really *are* fully descriptive, many philosophers of science have suggested that the scientific *process* that produces these judgments is nevertheless unavoidably infused with value judgments (Longino 1990; Douglas 2000, 2008, 2009). Harvard et al. (2021) illustrate this point in detail with respect to a recent study of the risks and benefits of the AstraZeneca COVID-19 vaccine (which, readers may recall, was at least somewhat less effective than the competing Pfizer and Moderna vaccines [Cohen 2021] and has recently been withdrawn). As they point out, value judgments unavoidably entered into this study at several points. First, they entered at the point of deciding that the study was worth undertaking in the first place (as opposed to deciding whether to continue offering the AstraZeneca vaccine on the basis of the existing evidence and using scarce scientific resources to investigate other questions). Second, they entered at the point of deciding which variables and outcomes were worth including in the model: for example, in including age, sex and "frontline status" as input variables but not including race, income, occupation or household size. This assumes, roughly, that the effect of the former variables on vaccine efficacy is more important to know about than the effect of the latter (or, perhaps, that it was unlikely enough that the latter would make any difference, as to justify disregarding them). Similarly, some adverse side-effects that could potentially result from the vaccine were not included as outcomes in the model, which again reflects a value judgment that all-things-considered these outcomes were not crucial to investigate. Third, value judgments entered at the point of deciding what level of evidence or certainty was needed to draw a conclusion from the data. As many philosophers of science have emphasized (e.g. Douglas 2000, 2008), deciding how much evidence to require before drawing a conclusion requires considering the *risks* of drawing a false conclusion as well as those of failing to draw a true conclusion. Again, the evaluation of how significant these risks are and hence of how much evidence we *ought* to require before reaching a conclusion involves value judgment.

The point here is not that the value judgments these scientists made were mistaken—they may well have been well-justified and sensible. The point is rather that even to reach a descriptive judgment about the efficacy of the AstraZeneca vaccine, scientists had to make these value judgments. That complicates the picture for deference because when we defer to scientists even on these purely descriptive judgments, we defer to them on judgments that were themselves shaped by value judgments—which we may or may not share—and that might have been different had different value judgments been employed. It also casts doubt on any solution to the problem of covert value judgments that simply instructs scientists to stay out of making value judgments (and any kind of

advice or recommendation) and "stick to the facts." Given the unavoidability of value judgment in the scientific process, even that which is aimed at establishing purely descriptive facts, this may just not be possible.

Let's take a step back. In this section, I've surveyed four ways in which expert testimony can covertly encode value judgments. But why exactly is this a problem? The primary reason is that, as I've been continually stressing, scientists and other experts generally have no special expertise or authority with respect to (pure) value judgments, and so there is no special reason to *defer* to them about these judgments. Thus, when their judgments, advice and recommendations encode both descriptive judgments (that they do have expertise with respect to) and pure value judgments (that they do not), it is hard to reach an overall verdict about whether to defer to them. And, when the value judgments are covert, we may *think* we are deferring to them only about the descriptive issues when we are in fact deferring to them about the normative (value) ones as well.

This might not be such a big problem if we could generally expect scientists to make value judgments that the majority of laypeople do (or would) share. But there may be reasons to expect scientists' value judgments to diverge from those of many laypeople. A key point here is that scientists and other experts are not demographically representative of the general population. In particular, they tend to be socioeconomically privileged (Morgan et al. 2022). As such, they may place insufficient weight on risks that matter more to the socioeconomically underprivileged. If this is so, then it can be quite reasonable for many laypeople to be cautious about deferring to experts who do not share their values, when these values are encoded into these experts' testimony.

Indeed, we might worry that a society in which people deferred to experts on value judgments, or on judgments that tacitly encode value judgments, would fall short of important democratic ideals (Hazlett 2016; van Wietmarschen 2019). Even if we think that there is no democratic requirement for a society to be responsive to all of citizens' ill-informed or irrational descriptive judgments, it is much less controversial that a thriving democracy should reflect its citizens' *values*, and not just those of an expert elite.[23] (Indeed, perhaps in a truly ideal democracy, politics would serve as the sphere in which such values are contested and *shaped* through a public, deliberative-democratic process.[24]) If people were to uncritically defer to experts about matters that encode value judgments, and to shape their political behavior accordingly, they might unwittingly contribute to fostering a political system that fails to represent their own values. This provides further reason for caution about encouraging people to defer to experts uncritically, at least about matters that involve value judgment in whole or in part.

6. What to do

What can we do about the problem of covert value judgments in expert testimony? A preliminary and basic point is that our current public discourse about deference to experts is too coarse. Those who encourage us to place more trust in experts tend to, in effect if not design, encourage us to defer to them about not just descriptive but value judgments. This occurs whenever we are encouraged to uncritically defer to expert *advice*, which, as I've shown, by its nature presupposes value judgments. On the other hand, those who encourage us to be place *less* trust in experts tend to throw the baby out

[23] Cf. Christiano's (2008) case for "values-only voting."
[24] Thanks to an anonymous reviewer for urging this point.

with the bathwater, encouraging us to defer to them about neither value nor descriptive judgments. Public discourse around deference to experts needs to more clearly distinguish descriptive and value judgments—if not under that terminology, then under other terminology—which in turn requires us to educate people about the distinction between the two.[25] And it ought to recognize that scientific experts' expertise generally pertains to descriptive rather than value judgments, encouraging more deference about the former than the latter.

But this is not enough to address the problem, given the ways that expert testimony often—sometimes unavoidably—incorporates both descriptive and value judgments. In these cases, the problem can best be addressed by experts themselves, in the ways that they give testimony and advice. I've already indicated that I don't think the best solution is for scientists to *refrain* from making value judgments entirely. When scientific *process* incorporates value-judgments in the ways illustrated by the vaccine study, this may not even be possible; moreover, experts are human beings too, and like other human beings, it is natural for them to have (and share) opinions about what *ought* to be done. Moreover, in some cases, these opinions may be based on value judgments that are largely uncontroversial, or at least widely shared. Indeed, as we've seen, the value-laden nature of the scientific enterprise makes it effectively *impossible* for scientists to refrain from making value judgments that shape their ultimate judgments even about purely descriptive matters.

Instead, I join others (e.g. Douglas 2008; Elliott & Resnik 2014; Harvard et al. 2021) in calling on scientists and other experts to make the value judgments that their advice is based on more explicit—that is, *o*vert. Sometimes, this may require some reflection, unearthing value judgments that they do not realize that they are making or assuming. If experts do this, laypeople will be in a better position to determine whether they agree with the value judgments being made, and hence whether to defer to the experts' overall advice. And we will all be in a better position to highlight the ways in which experts' value judgments are contestable, and when appropriate, to democratically contest them.[26]

My argument for this conclusion has been based on the contention that if scientists make their value judgments explicit, people will be in a better position to make informed deference decisions. However, though this is speculative, it is possible that it will also have the benefit of enhancing public trust in experts about the (descriptive) issues that they really *are* experts about. When experts fold value judgments into their testimony, this may feed into the sense that they are lecturing people or bossing them around, which in turn fuels resentment and distrust of experts. This resentment and distrust is plausibly part of what causes people to refuse to defer to experts even about purely descriptive issues. Thus, if we are feeling optimistic, we might predict that if experts are more transparent about their values—highlighting what role they are playing in generating their advice, and crucially, illustrating how *different* values might lead to different conclusions—this might help to restore trust in experts

[25] K-12 education currently teaches a distinction between "fact" and "opinion," a distinction philosophers tend to dislike, both because it belies that one can have an opinion about a factual matter, and because it tacitly and contentiously assumes that there are no facts about value (especially moral) matters (McBrayer 2015). I suggest that the related, but different, distinction between descriptive and value judgments—without the assumption that all opinions are equally valid when it comes to value judgments—would be a better distinction to teach.

[26] Schroeder (2021) goes a step further and argues that, when their research requires them to make value judgments, scientists should actually replace their own value judgments with those of the public. I am more skeptical about this suggestion, but discussion of it will have to wait for another day.

somewhat. If we are more pessimistic, though, we might hypothesize that those who don't share experts' values may be even less likely to trust them when these values are explicitly disclosed, as compared with when those values are more covert.

Whether the optimistic or pessimistic view—or some more nuanced possibility—is the right one is worthy of empirical study.[27] But my point here has been that in cases where laypeople's value judgments really do differ crucially from those of experts, it's not clear that they *should* defer to those experts unquestioningly, at least not in all cases. And so our focus should not always be on increasing blanket trust in experts, but on equipping people to figure out whether (and when) experts are indeed worthy of their trust. Whether greater transparency would enhance public trust in experts or not, it would certainly provide laypeople with relevant and important information to inform their deference decisions.

References

Ahlstrom-Vij, Kristoffer (2015). The Social Virtue of Blind Deference. *Philosophy and Phenomenological Research* 91(3), 545-582.

Autor, David H.; Dorn, David; and Hanson, Gordon H. (2016). The China Shock: Learning from Labor Market Adjustment to Large Changes in Trade. *National Bureau of Economic Research Working Paper* 21906.

Ballantyne, Nathan (2019). Epistemic Trespassing. *Mind* 128(510), 367-395.

Ballantyne, Nathan (2022). Novices and Expert Disagreement. In Ballantyne & Dunning (eds.), *Reason, Bias, and Inquiry: The Crossroads of Epistemology and Psychology*. Oxford University Press, 227-253.

Ballantyne, Nathan; Celniker, Jared; and Dunning, David (2024). "Do Your Own Research." *Social Epistemology* 38(3), 302-317.

Barnes, Gordon (forthcoming). The Abuse of Expertise and the Problem with Public Economics. *Social Theory and Practice*.

Bennett, Matthew (2020). Should I Do As I'm Told? Trust, Experts, and COVID-19. *Kennedy Institute of Ethics Journal* 30(3), 243-263.

Brooks, David (2020). The Rotting of the Republican Mind. *The New York Times*, November 26, 2020.

Broome, John (2012). *Climate Matters: Ethics in a Warming World*. W.W. Norton & co.

Carroll, Aaron E. (2022). To Fight Covid, We Need to Think Less Like Doctors. *The New York Times*, January 14, 2022.

Christiano, Thomas (2008). *The Constitution of Equality: Democratic Authority and its Limits*. Oxford University Press.

Clark, Cory; Isch, Calvin; Everett, Jim; and Shariff, Azim (ms.) Politicization Undermines Trust in Institutions, Even Among the Ideologically Aligned Public. Preprint at https://doi.org/10.21203/rs.3.rs-3239561/v1.

Cohen, Jon (2021). AstraZeneca Lowers Efficacy Claim for COVID-19 Vaccine, A Bit, After Board's Rebuke. *Science*, March 25, 2021.

[27] Some empirical work has begun to examine this, including Elliott et. al (2017), Hicks & Lobato (2022), and Douglas, Turri & Buckwalter (ms.). Also relevant is Clark et al.'s (ms.) perhaps surprising finding that individuals trust institutions less when they perceive them as "politicized," *even when there is ideological alignment between the individual and the institution*.

Crisp, Roger (2014). Moral Testimony Pessimism: A Defence. *Proceedings of the Aristotelian Society Supplementary Volume* 88, 129-143.

Davia, Cory and Palmira, Michele (2015). Moral Deference and Deference to an Epistemic Peer. *The Philosophical Quarterly* 65(261), 605-625.

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

Douglas, Heather (2008). The Role of Values in Expert Reasoning. *Public Affairs Quarterly* 22(1), 1-18.

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

Douglas, Heather; Turri, John; and Buckwalter, Wesley (ms.). Values and the Authority of Science.

Driver, Julia (2006). Autonomy and the Asymmetry Problem for Moral Expertise. *Philosophical Studies* 128(3), 619-644.

Duijf, Hein (2021). Should One Trust Experts? *Synthese* 199(3-4), 9289-9312.

Elliott, Kevin C.; McCright; Aaron M., Allen, Summer; and Dietz, Thomas (2017). Values in Environmental Research: Citizens' Views of Scientists Who Acknowledge Values. PLoS ONE 12(10): e0186049.

Elliott, Kevin C. and Resnik, David B. (2014). Science, Policy, and the Transparency of Values. *Environmental Health Perspectives* 122(7), 647-650.

Enoch, David (2014). A Defense of Moral Deference. *Journal of Philosophy*, 111(5) 229-258.

Graham, David A. (2020). How Trump Closed Down the Schools. *The Atlantic*, July 17, 2020.

Guisinger, Alexandra (2017). *American Opinion on Trade: Preferences Without Politics*. Oxford University Press.

Grundmann, Thomas (2021). Facing Epistemic Authorities: Where Democratic Ideals and Critical Thinking Mislead Cognition. In Bernecker, Flowerree & Grundmann (eds.), *The Epistemology of Fake News*, Oxford University Press, 134-155.

Harvard, Stephanie; Winsberg, Eric; Symons, John; and Adibi, Amin (2021). Value Judgments in a COVID-19 Vaccination Model: A Case Study in the Need for Public Involvement in Health-Oriented Modelling. *Social Science & Medicine* 286, 114323.

Hazlett, Allan (2016). The Social Value of Non-Deferential Belief. *Australasian Journal of Philosophy* 94(1), 131-151.

Hicks, Daniel J. & Lobato, Emilio Jon Christopher (2022). Values Disclosures and Trust in Science: A Replication Study. *Frontiers in Communication* 7.

Hills, Allison (2009). Moral Testimony and Moral Epistemology. *Ethics* 120(1), 94-127.

Hopkins, Robert (2007). What is Wrong with Moral Testimony? *Philosophy and Phenomenological Research* 74(3), 611-634.

Huemer, Michael (2005). Is Critical Thinking Epistemically Responsible? *Metaphilosophy* 36(4), 522-531.

Kant, Immanuel (1784/1996). An Answer to the Question: What is Enlightenment? In *Practical Philosophy* (ed. & trans. Gregor), Cambridge University Press.

Kaya, Ahmet Ihsan; Liadza, Iana; Low, Hailey; Sanchez Juanino, Patricia; and Millard, Stephen (2023). Revisiting the Effect of Brexit. National Institute of Economic and Social Research.

Kennedy, Brian; Tyson, Alec; and Funk, Cary (2022). Americans' Trust in Scientists, Other Groups Declines. Pew Research Center.

Kennedy, Brian and Tyson, Alec (2023). Americans' Trust in Scientists, Positive Views of Science Continue to Decline. Pew Research Center.

Lackey, Jennifer (2018). Experts and Peer Disagreement. In Benton, Hawthorne & Rabinowitz (eds.), *Knowledge, Belief, and God: New Insights in Religious Epistemology*, Oxford University Press, 228-245.

Lackey, Jennifer (2021). Preemption and the Problem of the Predatory Expert. *Philosophical Topics* 49(2), 133-150.

Leonhardt, David (2022). Follow The Science? *The New York Times*, February 11, 2022.

Levy, Neil (2022). Do Your Own Research! *Synthese* 200(356).

Longino, Helen (1990). *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton University Press.

Macedo, Stephen & Lee, Frances (2025). *In COVID's Wake: How Our Politics Failed Us*. Princeton University Press.

Matheson, Jonathan (2024). *Why It's OK Not to Think for Yourself*. Routledge.

McBrayer, Justin P. (2015). Why Our Children Don't Think There are Moral Facts. *The New York Times*, March 2, 2015.

McGrath, Sarah (2009). The Puzzle of Pure Moral Deference. *Philosophical Perspectives* 23, 321-344.

McGrath, Sarah (2011). Skepticism about Moral Expertise as a Puzzle for Moral Realism. *Journal of Philosophy* 108(3), 111-137.

Mercuri, Mathew (2020). Just Follow the Science: A Government Response to a Pandemic. *Journal of Evaluation in Clinical Practice* 26(6), 1575-1578.

Mervosh, Sarah, Miller, Claire Cain & Paris, Francesca (2024). What the Data Says About Pandemic School Closures, Four Years Later. *The New York Times*, March 18, 2024.

Millgram, Elijah (2015). *The Great Endarkenment: Philosophy in an Age of Hyperspecialization*. Oxford University Press.

Morgan, Allison C.; LaBerge, Nicholas; Larremore, Daniel B.; Galesic, Mirta; Brand, Jeannie E.; and Clauset, Aaron (2022). Socioeconomic Roots of Academic Faculty. *Nature Human Behavior* 6, 1625-1633.

Nguyen, C. Thi (2020). Cognitive Islands and Runaway Echo Chambers: Problems for Epistemic Dependence on Experts. *Synthese* 197(7), 2803-2821.

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

Rankin, David M. (2001). Identities, Interests, and Imports. *Political Behavior* 23(4), 351-376.

Russell, Jacob Hale & Patterson, Dennis (2025). *The Weaponization of Expertise*. MIT Press.

Schorr, Isaac (2023). Former National Institutes of Health Director Admits to 'Narrow', 'Really Unfortunate' Pandemic Mindset: 'We Weren't Thinking' About Collateral Damage. *Mediaite*, December 28, 2023.

Schroeder, S. Andrew (2021). Democratic Values: A Better Foundation for Public Trust in Science. *British Journal for the Philosophy of Science*, 72(2), 545-562.

Sliwa, Paulina (2012). In Defense of Moral Testimony. *Philosophical Studies* 158(2), 175-195.

Tiwald, Justin (2023). "Getting It Oneself" (*Zide* 自得) as an Alternative to Testimonial Knowledge and Deference to Tradition. *Oxford Studies in Epistemology* 7, 306-335.

Whaples, Robert (2006). Do Economists Agree on Anything? Yes! *The Economists' Voice* 3(9), 1-6.
van Wietmarschen, Han (2019). Political Testimony. *Politics, Philosophy & Economics* 18(1), 2019, 23-45.
Wiland, Eric (2017). Moral Testimony: Going on the Offensive. *Oxford Studies in Metaethics* 12, 51-75.
Williams, Bernard (2005). Who Needs Ethical Knowledge? In *Making Sense of Humanity*. Cambridge University Press, 203-212.
Worsnip, Alex (2017). Cryptonormative Judgments. *European Journal of Philosophy* 25(1), 3-24.
Zagzebski, Linda (2012). *Epistemic Authority: A Theory of Trust, Authority, and Autonomy in Belief*. Oxford University Press.
Zweig, David (2025). *An Abundance of Caution: American Schools, the Virus, and a Story of Bad Decisions*. MIT Press.