Navigating the Novice-Expert Problem

How non-experts can reason about expert knowledge

Thesis Project Proposal

PhD in Philosophy

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

André Alessi

2024

Table of Contents

Working title for the thesis

Navigating the Novice-Expert Problem: How non-experts can reason about expert knowledge

A statement of the research topic

Some epistemologists have argued that "expert" is an epistemic status, and that non-experts (or novices) lack sufficient knowledge of the required field to identify an expert; novices are essentially reduced to guesswork or relying upon non-epistemic characteristics of experts, such as their institutional position, qualifications, behaviour, appearance, social or political authority, interpersonal trustworthiness, etc., to identify which experts are speaking from a position of epistemic authority. Call this the *novice-expert problem*. Taking an approach influenced by recent work in regulative epistemology, epistemic trust and epistemic trespassing, I argue that this is an overly strict analysis of the information that novices have available to them, and that it is generally possible for novices to identify which experts or expert claims they should trust in a way that is well-justified, even if these strategies aren't *completely* reliable. Rather than treating the incidental failures of this approach as evidence of its general unsuitability, I intend to show that these strategies can be relied on, but depend heavily on the epistemic environment in which they occur. I also show that this model predicts and explains some recent cases where novices have trusted the wrong people in place of genuine experts, and that it also suggests practical strategies for helping novices avoid trusting the claims of those that we shouldn't.

Significance of the proposed research topic

How non-experts (or "novices") reason about experts and expert claims has recently received a great deal of attention in social epistemology and philosophy of science, and in the wider public domain. When, how and why we should "trust the experts" or "trust science" when trying to inform our beliefs regarding topics such as the COVID-19 pandemic and its response, climate change, and even "meta" concerns like "fake news" and misinformation on social media, is a regular topic of discussion and debate. Which expert information a person relies on has real consequences for how people behave, and thus can lead to real world harm if the expert they trust is unworthy of that trust.

This discussion is sometimes framed as a debate about whether expertise is even necessary, whether our "trust in expertise" has declined globally, or whether ordinary people should "do their own research" on any given topic to reach informed positions. If we assume that people who believe true things will behave more appropriately (act in a more informed manner, avoid harm, etc) and if we assume that the beliefs of genuine experts are more likely to be true than those of novices relative to specific fields of expertise, then there are strong moral and practical incentives to ensure that novices are relying on the right experts when they make informed statements about their field of expertise, and to help the public identify and avoid relying on people who are not genuine experts.

Within epistemology[1], it has been noted that since "expert" is a status typically conferred by virtue of what someone knows, and since novices definitionally lack access to that knowledge themselves, it is not clear how a novice might identify whether a someone is genuinely making claims with sufficient expert knowledge to justify those claims. Instead, it is assumed that novices must weigh non-epistemic factors such as the purported expert's qualifications, status, appearance, behaviour, etc even though these factors may not reliably track someone's actual knowledge in relation to a given question. Whether or not we believe this is a serious epistemic problem depends largely on whether we think that these other factors might correlate with genuine expertise often enough to help novices evaluate experts in a robust way, or distinguish between genuine experts and epistemic trespassers when possible.

Starting from a definition of expertise that emphasizes its epistemic characteristics, I will argue that when novices try to identify the right expert for a given situation, they can use several different strategies that can, collectively, give them access to good information about which experts they should trust more often than not, including social judgements about experts that distribute epistemic responsibility across

[1] See e.g. (Goldman, 2001) for a summary of the problem, and (Levy, 2024) for an argument that the Novice-Expert Problem means that non-experts cannot form well-justified beliefs about expertise.

a community. That these strategies are not completely reliable is, I argue, not evidence that novices cannot make good judgements about the epistemic authority of particular experts, but rather that they represent a "good enough" approach that is as reliable as many other kinds of reasoning we ordinarily use. I stress that the reliability of these strategies is heavily dependent on the epistemic environment we happen to find ourselves in, especially the epistemic characteristics of the communities we belong to, and that this framework explains cases in the recent past where people who are not experts have nevertheless been taken to be acting as such.

Furthermore, this framework identifies the nature of this mistaken attribution of expertise on the part of novices. By showing that recent messaging within the practice of science communications that stresses trust in experts and in science as an all-or-nothing, unquestioning choice made by the general public, we have unwittingly been telling novices to pay attention to the wrong things in their evaluations of expert claims, leading to worse outcomes rather than better, more informed judgements about which experts to rely on.

Building off recent empirical research on how the public thinks about science[2], I intend to show that most people have complex, nuanced views about which expert claims to trust and when. I will also argue that acknowledging this nuance gives us a useful starting point for educating the public in how to better identify the right expert for the right moment without reducing their judgements down to a simple binary choice. In fact, I argue that this binary approach makes it easier for charlatans to mislead people as to their expertise, and needs to be replaced with a more practical description of the relationship between novices and experts (rather than just describing the characteristics of ideal experts.)

Along the way, I will be addressing two related topics in social epistemology: the nature of epistemic trust (or trust in epistemic authorities) and the phenomenon of epistemic trespassing, which involves a person (deliberately or through ignorance) presenting themselves as an expert when they are not.

I will argue that epistemic trust (better described in this context as *epistemic reliance*, in that it is an explicitly parasocial relationship) is only a useful frame for the kinds of judgements novices make when we acknowledge that "generic epistemic trust" is too broad a concept to reflect the kinds of judgements novices make about the trustworthiness of experts, and instead we need to consider a relatively fine-grained view of reliance in sources of expert claims. I intend to argue that an underexamined element of epistemic trust is that it is a relationship not merely to those with expertise, but also to other people

[2] E.g. (Dixson et al., 2022; Slater et al., 2021)

in our communities who we have reason to believe we can rely on to make judgements about experts for us that we could have made but didn't.

Additionally, I will develop a description of epistemic trespassing that directly follows from my description of what expertise is. My description of epistemic trespassing enables us to distinguish between the over-confident novice, the charlatan, and the misbehaving expert when it matters, and to treat them as a single category when it doesn't.

The conclusion of my argument—that given the right epistemic context and correct advice, people can make good, rational choices about which experts to trust and when—is intended to be a corrective to the view sometimes presented in social psychology research that treats public attitudes towards expert knowledge as largely irrational, driven by social identity or political association rather than reason. This latter view mistakes the necessary impact of social epistemic factors on individual reasoning for identity-driven bias. By showing that these factors are actually relevant and useful in our judgements about things that experts claim when we are within a healthy epistemic context, I argue that when a novice fails to identify the right expert for a given situation, it is typically because the socially-inherited epistemic norms they have applied in a given context are inappropriate for the information their community has available to it, and that interrogating these norms is part of the achievable work of being an epistemically responsible novice.

Significance of proposed research topic:

- Propose a solution for The Novice-Expert Problem with practical implications for discourse and science education
- Rethink how science communicators discuss expertise with the public
- Refine the concept of epistemic trust in relation to expertise to include horizontal social relations, and conditional reliance on experts
- Distinguish between the epistemic characteristics of different kinds of epistemic trespasser

Literature review

The current conversation in philosophy on the Novice-Expert Problem has many parents across a vast swathe of philosophy and sociology, and deciding which lineage to emphasise shapes the response in potentially quite radical ways, as each branching path comes with its own assumptions, history, framing and even terminology. In my case, I am interested in engaging with a particular thread of scholarship that arose in the mid-1980s in epistemology that starts from a recognition that individuals rely on others for the evidence for many of the beliefs that they hold, and questions that arise from this perspective: What is the nature of our reliance on the epistemic authority of experts?; How should we identify the correct experts for any given topic?; and, How should we engage with the epistemic authority of those experts in a responsible way?

John Hardwig's 1985 paper, "Epistemic Dependence", is where I have chosen to begin engaging with one specific strand of this discussion because it consists of a concise description of a way of conceiving the relationship between expert testimony and the evidence a novice has access to when justifying beliefs about matters where expert knowledge seems necessary.

In the paper, Hardwig discusses the seemingly obvious (but underexamined in traditional epistemology) claim that most of our evidence for the beliefs we hold depends on other people. If we take "evidence" for beliefs to consist of the kinds of things that Enlightenment philosophers imagined (first-hand or logically necessary propositions) then individual humans actually have very little evidence for most of our beliefs. Instead, we rely on the testimony of others to justify many of our beliefs, and (Hardwig argues) we should extend the scope of good reasons for belief to include these appeals to epistemic authority:

> In this paper, I want to consider the idea of intellectual authority, particularly that of experts. I want to explore the "logic" or epistemic structure of an appeal to intellectual authority and the way in which such an appeal constitutes justification for believing and knowing. I have divided the paper into three parts. In the first, I argue that one can have good reasons for believing a proposition if one has good reasons for believing that others have good reasons to believe it, and that, consequently, there is a kind of good reason for believing which does not constitute evidence for the truth of the proposition. In the second, I urge that because the layman is the epistemic inferior of the expert (in matters in which the expert is expert), rationality sometimes consists in refusing to think for oneself. In the third, I apply the results of these considerations to the concept of knowledge and argue that the expert-layman relationship is essential to the scientific and scholarly pursuit of knowledge. (Hardwig, 1985, pp. 335–336)

Hardwig's view is that, if we accept that individuals rely on others for much of the evidence for their beliefs, and if some of this evidence is inaccessible to non-expert individuals because it is highly technical in nature, then individuals must rely on these experts implicitly, without further qualification, to hold all kinds of beliefs that we typically hold. Other than simply identifying the correct expert, a novice cannot engage meaningfully with the evidence the expert has for the things they assert, so any beliefs that the novice has concerning these matters must depend on experts completely. Furthermore, he asserts that this kind of epistemic reliance or dependence can still be considered a good reason to believe something: while depending on others does not give us evidence that a belief is true, it can give us evidence that someone else has evidence that the belief is true, which is as good a reason as we can expect to have to believe many, many things.

Hardwig frames his argument as a response to an "extremely pervasive model of rationality" which says that "the very core of rationality consists in preserving and adhering to one's own independent judgment; for how can one be sure one is being informed, not misinformed, if one suspends judgment?" His view is "that this model provides us with a romantic ideal which is thoroughly unrealistic and which, in practice, results in less rational beliefs and judgements." (Hardwig, 1985, p. 340)

Hardwig briefly addresses the question of what happens if a novice doesn't know who the right experts are:

> ...[I]f I do not know and have no way of finding out who the experts are, I will have no way to appeal to the chain of authority. I will then not know who has good reasons to believe that *p*, to whom to defer, or whose opinion (if any) will give me good reasons for believing that *p*. This sometimes happens, and, when it does, rational deference becomes impossible. But generally I can find someone whose opinion is more informed than mine and who can refer me to someone who is knowledgeable about whether or not *p*. And even if a layman, because of his relative inability to discriminate among experts, ends up appealing to a lesser expert instead of a greater expert, the lesser expert's opinion will still be better than the layman's. (Hardwig, 1985, p. 341)

In a footnote to this paragraph, Hardwig nods towards "...the often-excruciating practical problem of identifying who the real or best experts are--e.g., what is the patient faced with conflicting medical opinions to do? But these are logically posterior issues and problems; the argument of this section of the paper is that in any case he should not make his own diagnosis, nor even read up some about his problem and then make his own diagnosis." (Hardwig, 1985, pp. 341, footnote 6) That is, Hardwig's position is that the epistemically responsible approach to any question that requires expertise to answer

is to remain undecided unless an appropriate expert is identified, at which point a novice should accept the expert's view as correct.

Hardwig's argument in "Epistemic Dependence" also involves a reflection on the inter-disciplinary nature of science. If we must, individually, rely on the knowledge of experts on matters where we are novices, then experts themselves must also rely on the expertise of other experts because expertise is often narrowly defined, so even scientists working on relatively narrow questions must rely on others with slight different domains of expertise to effectively advance inquiry, both pragmatically (an experiment might need both particle physicists and computer technicians, for example) and more generally in terms of how experts literally reference other experts in their work. Individually, even experts might only be said to know a few things, but collectively the scientific community can be said to know a great deal, because everyone, expert or novice, depends on others for much of their evidence, and the rest of us who depend on the scientific community can rely on them for evidence that our beliefs are well-justified.

An odd consequence of this approach, Hardwig notes, is that in asserting that much of our evidence for beliefs involves nothing but dependence on others, and if we think that this kind of dependence is so common as to be inescapable, then if we want say that people know things based on this dependence, rather than that they merely have good reason to believe them, it follows that it is conceivable to say that a person can know something without having evidence for the truth of it, or even without understanding it:

> Suppose someone tells me something that is true without giving me evidence for its truth. Perhaps *A* tells me that laetrile does not cure cancer without giving me the studies that prove this, much less the concrete data on which those studies were based. But suppose I have good reasons to believe that *A* is an authority in the field of cancer research and so I believe what he tells me. Do I then *know* that laetrile does not cure cancer, or have I achieved something much less than knowledge (perhaps only right opinion or rational belief)? If I then know, is it possible for one to know that *p* without possessing evidence for the truth of *p*. But that seems paradoxical or counterintuitive; for, in the cases we are considering, evidence is relevant to establishing knowledge, but we are asking whether it is possible to have this knowledge without the relevant evidence.
>
> Even more paradoxical is the idea that *B* can know that *p* even though he doesn't understand that *p*. Suppose an eminent authority in particle physics tells me that a quark is a fundamental particle, and suppose this is true. But I don't even understand what that means, because I have

> no notion of what a quark is or what counts as a fundamental particle. However, I check up on the physicist, and, as a result, I *know* that a quark is a fundamental particle, thought I don't even understand what I know? (Hardwig, 1985, p. 345)

Hardwig offers two potential ways of interpreting what is happening here: First, it might be that we can be said to know things vicariously through chains of dependence: If I know that you know that *p*, and if you are indeed an expert who knows that *p*, then I can be said to know that *p*, too. Alternatively, it might be that, past some point in the steps along this chain, it cannot be said that I know that *p*, but rather that the community I belongs to knows that *p*. Each approach has profound implications for the study of epistemology, which has traditionally been framed as analysis of the knowledge of individual knowers, rather than analysis of how much of our knowledge depends on the people around us.

In a second paper, "The Role of Trust in Knowledge" (Hardwig, 1991) Hardwig explores the nature of this dependence in more detail. Adding to his previous argument, he notes that the nature of our epistemic dependence requires not merely deference but trust, and all the complicated dynamics (involving the novice, the expert and their community) that come along with it. Novices are responsible for correctly identifying the appropriate epistemic authorities to place their trust in; experts are responsible for ensuring that they are, in fact, worthy of that trust; and institutions are responsible for building mechanisms of accountability that ensure that the experts they credential act appropriately. The underlying epistemic relationship remains unchanged in that our reliance on others for evidence we do not ourselves possess is still one of deference to epistemic authority, but considering trust adds an ethical layer to knowledge:

> Clearly, the implications of the role of trust in knowledge will reach beyond epistemology and the philosophy of science into ethics and social philosophy. I close with just one example. The prevailing tenor of twentieth-century Anglo-American philosophy has been that epistemology is more basic than ethics. On this view, ethics must meet epistemological standards on pains of bankruptcy. And the prevailing suspicion in our culture—a suspicion nurtured by philosophy—is that ethics cannot pass the epistemological test, and that there is thus no ethical knowledge. Science, in contrast, is commonly believed to be too "hard" and "objective" to require anything as mushy and subjective as ethics.
>
> Bust scientific realism—indeed any theory that grants objectivity to scientific judgements—turns out to be incoherent when combined with subjectivism and skepticism in ethics. It remains true, of course, that ethical claims must meet epistemological standards. But if much of our knowledge rests on trust in the moral character of testifiers, then knowledge depends on

> morality and epistemology also requires ethics. In order to qualify as knowledge (or even as rational belief), many epistemic claims must meet ethical standards. If they cannot pass the ethical muster, they fail epistemologically. (Hardwig, 1991, p. 708)

Although Steve Fuller is now more well-known for his arguments in favour of including Intelligent Design in educational programs that teach evolution [3] (Fuller, 2008), in an earlier work titled "Social Epistemology" (Fuller, 2002 2nd ed.) Fuller responds to Hardwig's view that knowledge of the sort that relies on expertise cannot be adequately independently evaluated by a novice. Fuller, in the concluding chapter titled "Social Epistemology and the Problem of Authoritarianism", describes Hardwig's argument in "Epistemic Dependence" as:

> ...the *Authoritarian Theory of Knowledg*e. In a nutshell, it says that the rationality of thinking for oneself diminishes as society's knowledge gathering activities expand to the point of requiring a division of cognitive labor into autonomous expertises. (Fuller, 2002, p. 278)

Fuller goes on to argue that ATK arises due to a conflation of different aspects of the problem: by confusing what we think an expert must be with the evidence we have for how novices actually rely on experts in the real world, proponents mistakenly assume that categorical deference to epistemic authority is the best choice for a novice to make in every situation, rather than an ideal that relies on a lot of other things being actually (not merely possibly) true.

Although I will not focus on the rest of Fuller's arguments in my project, I find his distinction between analytic, empirical and normative claims about expertise to be a useful framework, and one I will come back to repeatedly.

In a relatively technical paper called "Epistemic Paternalism: Communication Control in Law and Society", Alvin Goldman (Goldman, 1991) discusses variations of a principle in epistemology called

---

[3] The reader might at this point wonder why I would include a reference to a person who is apparently no longer well-regarded in the field. Although I do not intend to spend much time relying on Fuller's work directly, an element of the argument I intend to make is that our evaluations of the epistemic authority of experts is typically not the all-or-nothing defence Hardwig and others describe, but rather a more fine-grained judgement about specific claims in relation to the apparent expertise of the individual in question. Fuller's discussion of categorical deference is included here because it is a reasonable response to concerns about the trustworthiness of individual experts, and the trustworthiness of the scientific community as a whole, even if the solutions he proposes are likewise a little too broad to be practically useful. It is also important, I think, to address the actual ideas of people who might be "outside the tent" if my concern is indeed in giving practical advice to people who are already deeply suspicious of credentialed institutional expertise, a community which I think Fuller can be said to at least be sympathetic to. For specific critiques of Fuller's work on ID, I recommend Kevin Lambert's "Fuller's Folly, Kuhnian Paradigms, and Intelligent Design" (Lambert, 2006)

"the requirement of total evidence", or RTE. This principle can be formulated in different ways, but roughly works out to the assertion that the ideal environment for an individual to correctly formulate beliefs or subjective probabilities is one in which the individual has access to all relevant information about the matters in question. A particular version of this principle, which Goldman calls C-RTE (for "control") says that if a given expert or other authority has control over information relevant to a particular matter, they should be expected to provide an individual with all that relevant information in the name of enabling the individual to reasonably formulate beliefs that are most likely to be correct. Goldman objects to C-RTE as both impractical and epistemically dubious; individuals should not expect to have access to *all* the evidence they need to impartially make up their minds on a specific matter because their evaluation of what information is relevant or not is not may in fact undermine epistemically valuable outcomes (i.e. a novice might misunderstand the meaning of some evidence and come to hold beliefs that are false.) He illustrates this with examples from law and education where withholding information might be legally or ethically important, or where it might actually be practical to prevent individuals from being misled (such as not inviting a drug dealer to a classroom lesson about drug safety to give "the other side of the story.")

Goldman refers to this control of information as "epistemic paternalism":

> I shall think of communication controllers as exercising epistemic paternalism whenever they interpose their own judgment rather than allow the audience to exercise theirs (all with an eye to the audience's epistemic prospects). Thus, the exclusion of evidence of doubtful veracity (e.g., hearsay evidence) also qualifies as epistemic paternalism. The courts apparently feel that jurors cannot be counted on to discount hearsay testimony adequately. So they substitute their own wisdom for that of the jurors. Similarly, when judges exclude evidence as irrelevant, they must use their assessment of whether the evidence affects the probabilities of the propositions in question (the jurors' assessments might have been different). When these categories of exclusion are added, it is clear that courts engage in a substantial amount of epistemic paternalism.
>
> Is such paternalism really warranted? Are these rules good rules from an epistemic point of view? This is open to dispute…
>
> … I am inclined to think that some paternalism is appropriate in this arena, although I shall not take a firm stand on specific policies. What I want to do is identify the questions clearly, and put them in the framework of a wider set of questions. What we have here is a set of rules or practices whose adoption has an impact on the truth values of the doxastic decisions that

> cognizers make. One question is: For each such rule, how good is its impact from a veritistic point of view, that is, in terms of the likelihood of getting truth and avoiding error? Would alternative rules or practices have better veritistic properties? (Goldman, 1991, pp. 119–120)

Goldman's 2001 paper, "Experts: Which Ones Should You Trust?" (Goldman, 2001) is often cited as the paper that laid out the Novice/Expert Problem as an ongoing research project:

> ...some issues in epistemology are both theoretically interesting and practically quite pressing. That holds of the problem to be discussed here: how laypersons should evaluate the testimony of experts and decide which of two or more rival experts is most credible. It is of practical importance because in a complex, highly specialized world people are constantly confronted with situations in which, as comparative novices (or even ignoramuses[4]), they must turn to putative experts for intellectual guidance or assistance. It is of theoretical interest because the appropriate epistemic considerations are far from transparent; and it is not clear how far the problems lead to insurmountable skeptical quandaries. This paper does not argue for flat-out skepticism in this domain; nor, on the other hand, does it purport to resolve all pressures in the direction of skepticism. It is an exploratory paper, which tries to identify problems and examine some possible solutions, not to establish those solutions definitively.
>
> The present topic departs from traditional epistemology and philosophy of science in another respect as well. These fields typically consider the prospects for knowledge acquisition in "ideal" situations. For example, epistemic agents are often examined who have unlimited logical competence and no significant limits on their investigational resources. In the present problem, by contrast, we focus on agents with stipulated epistemic constraints and ask what they might attain while subject to those constraints. (Goldman, 2001, p. 85)

In his discussion of Hardwig's portrayal of the problem in "Epistemic Dependence", Goldman describes Hardwig's view of the relationship between reliance on experts and knowledge as a kind of "blind" trust where the only information a novice might have about a given expert's authority to make the claim is that they happened to say it while being experts:

> Hardwig is intent on denying full-fledged skepticism; he holds that the receiver of testimony can acquire "knowledge" from a source. But by characterizing the receiver's knowledge as "blind",

[4] Goldman's use of "ignoramuses" is, I would suggest, a rare case where the term is used descriptively rather than pejoratively, at least in modern writing.

> Hardwig seems to give us a skepticism of sorts. The term "blind" seems to imply that a layperson (or a scientist in a different field) cannot be rationally justified in trusting an expert. So his approach would leave us with testimonial skepticism concerning rational justification, if not knowledge. (Goldman, 2001, p. 86)

Goldman discusses other, slightly less strict, interpretations of Hardwig's position which assert that knowledge can be justified by the assertions of others whether or not we are in a position to evaluate the epistemic characteristics or reliability of the speaker. Goldman frames these positions as being opposed to the *reductionist* or *inductionist* view of testimony: this is the view that the reliability (or otherwise) of a given testimonial claim is much like any other kind of inference we make, where "Bob says X" is reliable justification to believe X iff Bob is usually a reliable source in relation to X.[5] Goldman's own position is that anti-reductionist views of this sort are apparently inadequate, leaving some kind of reductionist theory as the most likely alternative:

> I shall not myself be offering a full-scale theory about the justification of testimonial belief. In particular, I do *not* mean to be advancing a sustained defense of the reductionist or inductivist position. Of greater concern to me is the recognition that a hearer's evidence about a source's reliability or unreliability can often bolster or defeat the hearer's justifiedness in accepting testimony from that source. (Goldman, 2001, p. 88)

Goldman then describes what he believes the core of the Novice-Expert Problem to be:

> There are, of course, degrees of both expertise and novicehood. Some novices might not be so much less knowledgeable than some experts. Moreover, a novice might in principle be able to turn himself into an expert, by improving his epistemic position vis-a-vis the target subject-matter, e.g., by acquiring more formal training in the field. This is not a scenario to be considered in this paper, however. I assume that some sorts of limiting factors-whether they be time, cost, ability, or what have you-will keep our novices from becoming experts, at least prior to the time by which they need to make their judgment. So the question is: Can novices, while remaining

[5] This description of reductivism in relation to testimony is best outlined by C. A. J. Coady's "Testimony and Observation" where Coady describes David Hume's (Hume, 1748[2012]) position on the matter as reductionist, before going on to criticise it on that basis, leading to him assert: "Hence, I suspect that the problem of justifying testimony is a pseudo-problem and that the evidence of testimony constitutes a fundamental category of evidence which is not reducible to, or justifiable in terms of, such other basic categories as observation or deductive inference." (Coady, 1973, p. 154)

> novices, make justified judgments about the relative credibility of rival experts? When and how is this possible? (Goldman, 2001, p. 89)

Goldman, unlike Hardwig, sees the problem as being one of a novice attempting to rationally choose between *two* seemingly equally reliable experts, rather than merely considering whether appeals to the epistemic authority of any expert at all count as evidence:

> ...in what I am calling the novice/expert problem (more specifically, the novice/2-expert problem), the novice is not in a position to evaluate the target experts by using his own opinion; at least he does not think he is in such a position. The novice either has no opinions in the target domain, or does not have enough confidence in his opinions in this domain to use them in adjudicating or evaluating the disagreement between the rival experts. He thinks of the domain as properly requiring a certain expertise, and he does not view himself as possessing this expertise. Thus, he cannot use opinions of his own in the domain of expertise-call it the *E-domain*-to choose between conflicting experts' judgments or reports. (Goldman, 2001, p. 90)

In this context, the "experts" in question are cognitive experts, in that they are distinguished by what they *know* and what they can do with that knowledge, rather than the broader sense of "expertise" which would include people with a high level of practical skill. Cognitive experts have access to a large number of true facts about their given domain of expertise, along with the ability to make good inferences from those facts to answer other questions that might arise in relation to their domain.[6] It is this focus on knowledge that makes the Novice-Expert Problem a specifically epistemological problem, rather than merely a practical problem that can be evaluated by comparing the material consequences of the choices the novice makes. Writes Goldman:

> The novice/2-experts problem is whether a layperson can *justifiably* choose one putative expert as more credible or trustworthy than the other with respect to the question at hand, and what might be the epistemic basis for such a choice? (Goldman, 2001, p. 92)

The remainder of the article involves Goldman evaluating five possible sources of evidence a novice might have for evaluating their choice between two experts:

---

[6] "...an expert (in the strong sense) in domain D is someone who possesses an extensive fund of knowledge (true belief) and a set of skills or methods for apt and successful deployment of this knowledge to new questions in the domain. Anyone purporting to be a (cognitive) expert in a given domain will claim to have such a fund and set of methods, and will claim to have true answers to the question(s) under dispute because he has applied his fund and his methods to the question(s)." (Goldman, 2001, p. 92)

> (A) Arguments presented by the contending experts to support their own views and critique their rivals' views.
>
> (B) Agreement from additional putative experts on one side or other of the subject in question.
>
> (C) Appraisals by "meta-experts" of the experts' expertise (including appraisals reflected in formal credentials earned by the experts).
>
> (D) Evidence of the experts' interests and biases vis-a-vis the question at issue.
>
> (E) Evidence of the experts' past "track-records'' (Goldman, 2001, p. 93)

Goldman's evaluation of each of these sources of potential evidence is that it might sometimes be possible for novices to reason well about experts, but it is often very, very difficult, mostly for reasons beyond the novice's control. He concludes:

> My story's ending is decidedly mixed, a cause for neither elation nor gloom. Skeptical clouds loom over many a novice's epistemic horizons when confronted with rival experts bearing competing messages. There are a few silver linings, however. Establishing experts' track-records is not beyond the pale of possibility, or even feasibility. This in turn can bolster the credibility of a wider class of experts, thereby laying the foundation for a legitimate use of numbers when trying to choose between experts. There is no denying, however, that the epistemic situations facing novices are often daunting. There are interesting theoretical questions in the analysis of such situations, and they pose interesting practical challenges for "applied" social epistemology. What kinds of education, for example, could substantially improve the ability of novices to appraise expertise, and what kinds of communicational intermediaries might help make the novice-expert relationship more one of justified credence than blind trust. (Goldman, 2001, p. 109)

Goldman's formulation of the Novice-Expert Problem portrays the issues involved as relatively narrowly constructed: curious, well-meaning novices trying to identify the expert that is most likely to be correct from a limited number of experts who are all acting reasonably responsibly. In subsequent years, and particularly following the 2016 US Presidential Election, Brexit and the COVID-19 Pandemic, emphasis has shifted to cases where novices might not be motivated *merely* by a desire to know the truth (De Cruz, 2020; D. M. Kahan, 2015; D. M. Kahan et al., 2011; Levy, 2019), where experts might not be true authorities in the topic in question either due to the narrowness of their expertise or their own overconfidence (Ballantyne, 2015, 2018, 2019; Ballantyne et al., 2024; Millgram, 2015), or where the

social epistemic context influences the kinds of information novices have available to them to correctly reason about who the right epistemic authorities must be (Dentith, 2018; Levy, 2007, 2021, 2024; C. Thi Nguyen, 2020, 2021; Slater et al., 2019, 2022)

Neil Levy has written extensively on the topic of why people might come to trust the wrong epistemic authorities, or to distrust the right ones. In "Due deference to denialism: explaining ordinary people's rejection of established scientific findings" (Levy, 2019), Levy argues that, rather than being the result of a deficit in information or rationality, people who come to believe in things like the idea that anthropogenic climate change is a hoax are for the most part merely "epistemically unlucky": the cues that we all use to identify epistemic authorities that are trustworthy, competent and benevolent are highly dependent on the collective epistemic work done by the communities we belong to, and as a result can be interfered with by political polarisation in relation to those specific groups we align with. That is, if we belong to a group (such as political conservatives) our expectations of what a reliable expert looks like are distorted by our "knowledge" that our views on certain topics can only be justified by evidence from people who agree with our group, and experts that disagree with the views of our group are treated as untrustworthy. Levy argues that this process occurs prior to any individual reasoning we might do, so it cuts us off from genuinely good evidence that would allow us to identify the right experts:

> I have suggested that we are disposed to be epistemic individualists not despite but because we are reliant on collective cognition: epistemic individualism is an adaptation for group deliberation. Under a variety of circumstances, epistemic individualism must be appropriately constrained for it to be adaptive for group deliberation: our disposition to defer to those with the appropriate markers of competence and benevolence allows epistemic individualism to be epistemically virtuous.
>
> When cues to reliability dissociate from genuine expertise, epistemic individualism may not be appropriately constrained, despite our continuing to defer. When societies become polarized and matters of fact deeply politicized, our dispositions to defer may produce systematic divergences on matters of fact. This is not an individual failing. Not only are the conservatives who reject climate science neither less rational nor always less well-informed than the liberals who accept it, the mechanisms that dispose them to defer to benevolent and competent individuals are working as designed. It is further facts, about earlier links in the chain of deference, that entail that their beliefs are unwarranted, not facts about how they deploy their cognitive resources. (Levy, 2019, p. 324)

Helen De Cruz's "Believing to belong: Addressing the novice-expert problem in polarized scientific communication" (De Cruz, 2020) attempts to include the non-epistemic aspects of the Novice-Expert Problem in her analysis (echoing Hardwig's exhortation that ethical concerns are inextricable from epistemic ones):

> The novice-expert problem is usually framed in purely epistemic terms, but, as I will argue, accepting testimony poses non-epistemic demands as well. These include social and moral considerations. As a result of such non-epistemic factors, people will sometimes accept discredited scientific views if doing so helps them to coordinate better with individuals who hold similar beliefs. I will argue that to adequately explain non-specialists' deference to fringe scientific theories such as climate change denial or creationism we need to understand both their epistemic and non-epistemic motivations. (De Cruz, 2020, p. 2)

De Cruz responds to Levy's (Levy, 2019) discussion of denialism (discussed above) above by pointing out that it isn't simply the *epistemic* features of the communities we belong to that shape our perceptions of who should count as a relevant epistemic authority, but also non-epistemic factors such as perceived group expectations, meaning that an explanation that aligns these factors with Levy's narrower view of epistemic heuristics concerning expertise provides a better model for understanding recent experimental results that study partisanship and polarisation in scientific debates:

> I have examined why people defer to scientific fringe theories. A popular explanation of this phenomenon (termed denialism) is that people find themselves at the end of testimonial chains that trace to merchants of doubt and pseudo-experts. Levy has recently argued that these deferential chains can be explained as a result of epistemic bad luck, as people use sensible heuristics (benevolence and competence) to assess scientific testimony which go awry in the case of denialism. As I have argued, this fails to account for social considerations in the evaluation of testimony. An alternative explanation says that people are driven by cultural values and perceive scientific information through this cultural lens. I have argued that this fails to account for increasing partisanship in the reception of scientific information, especially in the US, and that it fails to explain a lingering concern for epistemic considerations, recently emphasized by failures to replicate the backfire effect. (De Cruz, 2020, p. 14)

De Cruz's diagnosis is even more bleak than Levy's: barring large scale interventions in social groups, traditional and social media, it is unlikely that individuals have much hope of identifying the right experts at all, as individual reasoning in this context is ineffective.

Both Levy and De Cruz reference another similar approach popularised by Dan Kahan (D. Kahan & Braman, 2005; D. M. Kahan, 2015; D. M. Kahan et al., 2010, 2011) that Kahan describes as "cultural cognition":

> The cultural cognition thesis holds that individuals form risk perceptions that reflect their commitments to contested views of the good society. (D. M. Kahan et al., 2010, p. 501)

That is, cultural cognition operates as a kind of group evaluation of value and risk that biases individuals away from the idealised, individualised reason assumed by traditional epistemology:

> Essentially, cultural commitments are prior to factual beliefs on highly charged political issues. Culture is prior to facts, moreover, not just in the evaluative sense that citizens might care more about how gun control, the death penalty, environmental regulation and the like cohere with their cultural values than they care about the consequences of those policies. Rather, culture is prior to facts in the cognitive sense that what citizens believe about the empirical consequences of those policies derives from their cultural worldviews. Based on a variety of overlapping psychological mechanisms, individuals accept or reject empirical claims about the consequences of controversial polices based on their vision of a good society. (D. Kahan & Braman, 2005, p. 150)

Given that the issues we are now concerned with are highly "politicised" rather than Goldman's portrayal of merely technically challenging scientific knowledge, Kahan claims that cultural cognition operates as an unavoidable distortion of individual judgements about epistemic authority, preventing good reasoning when it matters.

Another issue that has been identified in the time since Goldman's paper is the question of what happens when an expert steps outside their domain of expertise? In "Epistemic Trespassing", Nathan Ballantyne (Ballantyne, 2018) describes the phenomenon:

> Epistemic trespassers are thinkers who have competence or expertise to make good judgments in one field, but move to another field where they lack competence—and pass judgment nevertheless. We should doubt that trespassers are reliable judges in fields where they are outsiders. (Ballantyne, 2018, p. 367)

Unlike concerns about the possibly poor or unreliable reasoning done by novices who are trying to identify the right expert, concerns about experts-who-aren't have a bearing on whether novices can even spot when experts are making this sort of mistake. Epistemic trespassing tends to occur in highly visible

spaces, via media appearances or self-organised public talks outside the restrictive confines of institutions, meaning it isn't easy for a novice to realise that such trespassing is occurring, since the non-epistemic conduct of the purported expert seems to closely match what a novice might think an expert looks like:

> Experts on a public stage are cast in the role of the 'public intellectual' or 'celebrity academic'. They may find trespassing all but impossible to resist. Microphones are switched on, TV cameras zoom in, and 'sound bites' come forth, coaxed out of the commentators by journalists. So what do you have to say about philosophy, Neil deGrasse Tyson? And what about arguments for the existence of God, Professor Dawkins? (Ballantyne, 2018, p. 369)

Ballantyne notes that trespassing occurs fairly frequently, and is complicated by the fact that some questions are "hybridized": they require evidence that can only be obtained by consulting experts across domains:

> We need to see how these connections cause experts unwittingly to overstep their limits. This easily happens when experts investigate what I will call hybridized questions—ones addressed and answered by combining evidence and techniques from two or more fields.
>
> Fields are fixed by a set of questions and expertise is fixed by bodies of evidence and skills needed to answer a field's questions. But note that sometimes fields overlap or converge and come to share a question. This may happen in one of three ways: (a) the evidence required to answer a question reliably or responsibly comes from two or more fields; (b) the skills required to evaluate the evidence well come from two or more fields; or (c) both the relevant evidence and the relevant skills required to answer a question reliably or responsibly come from two or more fields. In such situations, the experts in one field will not all satisfy the same evidence or skill conditions as the experts in an- other field. Since the experts are in different fields, the evidence and skills that constitute their expertise differ. A hybridized question, then, is one that experts in distinct fields could try to answer using their own resources. Take the question, 'What caused the Cretaceous-Paleogene extinction event?' It is addressed by experts in palaeontology, geology, climatology, and oceanography, among other fields. The question is thoroughly hybridized and answering it calls for a host of evidence and skills. A question can be hybridized even if it is experts from one field only who address the question. But once investigators recognize that a question is hybridized, they should think that answering it reliably and responsibly calls for cross-field resources.
>
> These days, hybridized questions abound. (Ballantyne, 2018, p. 372)

While this does count as a form of trespassing, it is also an inevitable consequence of the interdisciplinary nature of science, often due to what Elijah Millgram calls "Hyperspecialization". (Millgram, 2015) An expert might not initially realise that they are trespassing, but once they do, they should either take steps to defer to a more appropriate expert, to expand their own expertise through research and study, or to interrogate whether the expertise that they already do have is meaningfully transferrable to the domain they are trespassing in. In this paper, Ballantyne only considers experts who believe they are acting responsibly when trespassing, rather than thornier questions involving non-experts who incorrectly think they are (or don't need) experts[7], which is the matter he discusses in "Do Your Own Research" (DYOR), written with Jared B. Celniker and David Dunning. (Ballantyne et al., 2024):

> What's notable is that people know there is a crisis about expertise—its defenders and opponents alike. But they differ sharply about what the nature of the crisis is and how it should be addressed, either by individuals or society as a whole. This essay is an effort to make sense of one prominent idea that plays a fascinating role in these debates and controversies.
>
> That idea is DYOR. (If the reader has not heard of it, we recommend: Do Your Own Research.) The slogan is flexible and versatile, used frequently on social media platforms, in messages about topics from medical science to financial investing to conspiracy theories. Interestingly, some proponents of the open science movement seem to presuppose that laypeople are proficient in doing their own research—that is why, according to them, articles and data on primary and unfiltered scientific research should be available online to anyone. (Ballantyne et al., 2024, p. 302)

Ballantyne and his co-authors consider why people might prefer DYOR to relying on expert authority (it is often tied up with both political identity, with subjective impressions of competence, and with idealistic views of what good critical thinking involves) and whether DYOR is ever the epistemically right approach (it might be but often isn't, and requires a kind of epistemic humility that may not come naturally to those attracted to the approach.) DYOR is a relatively new area of enquiry for social epistemology, and that part of our response to it should be to attempt to give advice to those who DYOR so they make better choices about the sources they trust:

---

[7] As far as I have been able to determine, there is little recent work in social epistemology that addresses the problem of *deliberately* malicious experts. Jennifer Lackey's recent paper "Preemption and the Problem of the Predatory Expert" (Lackey, 2021) discusses genuine experts who use their legitimate status for immoral personal benefit, which is one aspect of this, while Christoph Jäger's "False Authorities" (Jäger, 2024) considers non-experts who present themselves as experts, either unintentionally or intentionally. This will be a topic that I will return to in my own work.

> According to us, better research requires competence—skills and scaffolding—as well as a healthy dose of humble inquiry. To improve the efforts of DYOR enthusiasts, competence and humility need to be more accessible and commonplace. (Ballantyne et al., 2024, p. 314)

Neil Levy has recently written a paper[8] that discusses the Novice-Expert Problem directly, titled "No Trespassing! Abandoning the Novice/Expert Problem." (Levy, 2024) As the title suggests, Levy extends his argument from (Levy, 2019) and asserts that the Novice-Expert Problem, as (Goldman, 2001) is not solvable because novices can never be said to have the kinds of evidence required to identify the correct experts in an epistemically responsible way:

> Novices must trust experts if they are to come to hold justified beliefs about a range of very significant issues. They must trust experts because they can't judge these questions for themselves: they lack the evidence and the capacity required to assess the issues directly. Goldman, and those he inspired, aim to develop criteria novices can use to assess expert reliability that do an endrun around our incapacity to assess the issues. They haven't succeeded at this aim: when novices find themselves uncertain who to trust, these criteria won't come to their aid. To apply each criterion, we need deep knowledge of the science and its structure – knowledge that is available only to experts. (Levy, 2024)

Levy's solution is that instead novices should put their trust in the institutions that credential experts, and that it is this trust in institutions, not reasoning about evidence from the perspective of a novice, that is needed to do as well as we are able to identify the right experts:

> 'Trust well' may not be very useful advice, but it has the virtue of being the correct advice. Just as scientists can conduct research and extend human knowledge only by trusting well, so we laypeople must trust well. If we don't already trust well, the criteria won't help us trust better. (Levy, 2024)

In practice, this approach works out to something like Hardwig's view that the beliefs we hold on the basis of expert testimony are held blindly, while ethical commitments between ourselves and the institutions in question have to be seen as integral to the process of forming these beliefs.

---

[8] Levy has also recently published a book on a range of related topics, *Bad Beliefs*. (Levy, 2021) Levy's arguments there on the Novice-Expert Problem are very similar to those in "No Trespassing!" (Levy, 2024) so I won't describe them here, except to note that Levy's wider theory of how heavily people rely on the nature of the epistemic environment they find themselves in has implications for much of the broader project my argument occurs within, and criticisms of that view will be necessary for my argument to be successful.

Responding to an earlier paper of Levy's which makes a similar point about the importance of trust in institutions[9], M Dentith has argued in "Expertise and conspiracy theories" that in at least some cases, such as in relation to conspiracy theories[10], it is not possible to unproblematically appeal to experts, as the conspiracy theories themselves include reasons to think that *any* relevant epistemic authority is suspect. Assuming that the kind of conspiracy theories we are concerned about can be contrasted with some official version of events:

> If there is a case for preferring official theories over conspiracy theories, it will be grounded in the understanding that only some official theories are epistemically superior; the officialness of a theory does not necessarily tell us anything about its epistemic merits. Officialness in this case only tells us that the theory has been endorsed by some influential institution. Given institutions are many and varied, *some* endorsements will be epistemic whilst others will be merely political or pragmatic...
>
> ... The telling point is that being an authority—a member of an influential institution— does not necessarily make one an expert; a theory can be labelled as 'official' just because it has been endorsed by some influential authority. This does not tell us about the epistemic nature of the endorsement, which might be political or pragmatic. As such, we should not confuse a theory's endorsement with it having any special epistemic character. An endorsement tells us little about whether the evidence supports the theory. It simply tells us that someone or some influential institution has lent support to it. (Dentith, 2018, pp. 197–198)

Dentith argues that, at least in the case of conspiracy theories, it isn't clear who the relevant epistemic authorities might be. Typically, those in good standing within the kinds of institutions we normally trust for expert advice might not be familiar with any given conspiracy theory because their own domain of expertise doesn't easily map on to the claims made by the conspiracy theory (although they might be able to offer expert opinion on some of the claims involved.) Other experts, who might portray themselves as experts regarding conspiracy theories, would seem to require a degree of knowledge about new conspiracy theories that is grounded in knowledge of established conspiracy theories, but it isn't necessarily the case that familiarity with one conspiracy theory involves transferrable expertise in

---

[9] "Radically Socialised Knowledge and Conspiracy Theories", (Levy, 2007)

[10] It seems fair to include conspiracy theories in this discussion at this point because many of the concerns we have about this terrain in recent years (namely which sources of information to trust) overlap considerably with concerns about conspiracy theories. The assumption much of the literature makes (often explicitly made in Ballantyne and Levy's works) is that novices are confronted with a choice between the "official position" on many matters and explanations or evidence that amount to (or rely on) conspiracy theories, such as beliefs about vaccination safety or gender affirming care.

other conspiracy theories. Furthermore, if we are to start from a position that assumes a given conspiracy theory is an alternative we should take seriously[11] until we have examined the evidence for and against it, then we need to also take into account the possibility that either any experts we might consult could be untrustworthy (which, as Hardwig and Levy point out, seems like a problem for any advice to novices who might wish to identify the right expert *to* trust.)

Another potential wrinkle in the debate is raised by the work of Matthew Slater, Joanna Huxster, Julia Bresticker, Emily Scholfield and others (Huxster et al., 2018; Slater et al., 2019, 2022; Slater & Scholfield, 2022) who have done a significant amount of empirical research[12] that examines the question of what we mean when we say that a person should "trust science" (as a variant on "trust the experts"):

> What precisely would it mean to ask whether some group— voting-age Americans without scientific training, say— "trusts science"? No one thing, presumably. Indeed, not many of the things that "trusting science" could mean are very plausible. The term might even seem at first glance to be a category mistake. Science is some- thing that people do; what would it mean to extend epistemic trust to an activity? (Slater & Scholfield, 2022, p. 1045)

Trust is often conceptualised as a reciprocal relationship; that is, a relationship between two individuals that involves some degree of reciprocal good will and responsibility.[13] This is contrasted with reliance, which involves one person observing another person's (or object's) habits whether or not that second individual is aware that the first is doing so. At least in this sense, many of the relationships we have to experts that have been described above, especially by Hardwig and Levy, don't quite seem to be ones of trust, since we aren't trusting an individual expert in most cases, but rather an institution or "the scientific consensus", and these experts and institutions are almost certainly not aware of us as

---

[11] Dentith has discussed a certain kind of conspiracy theory that should *not* be taken seriously due to its specific similarities with other unsupported conspiracy theories in "Suspicious Conspiracy Theories" (Dentith, 2022). The kind of conspiracy theory being gestured at here is not that kind, in that it is an (at least superficially) reasonable alternative.

[12] Slater, et al.'s research is focused almost exclusively on attitudes about trust in science in the USA, so here I rely mostly on their conceptual discussions rather than their empirical findings. For empirical research on similar matters here in Aotearoa New Zealand, I will be relying on Dixson, et al's recent paper "Trust in science and scientists: Effects of social attitudes and motivations on views regarding climate change, vaccines and gene drive technology" (Dixson et al., 2022) and related works.

[13] This is a view espoused in Annette Baier's classic paper, "Trust and Antitrust" (Baier, 1986) and has largely been how trust and reliance are distinguished. Some recent work, e.g. (C. T. Nguyen, 2022) has challenged this distinction, but for the purposes of this discussion I will defer to Slater's approach, which distinguishes trust and reliance more or less as Baier does. Onora O'Neill's "A Question of Trust" (O'Neill, 2002) is another work that addresses this topic.

individuals. If instead we frame this relationship as a kind of reliance, then we seem to be more accurately describing the one-way nature of this epistemic dependence in practice, although we give up any assumption of explicit benevolence towards novices on the part of certified experts, which might lead us to wonder whether assuming benevolence is reasonable. This may be something that could be resolved by asserting that experts do have some kind of explicit responsibility to the public in general, but that then leaves us to ask whether this is actually the case, and who gets to be included in "the public".

Another issue Slater, et al. address is the question of what it might mean to trust in "science" or "the scientific consensus." Hardwig, Goldman, Levy and others have assumed that any disputes or disagreements between experts are, for the most part, outside the scope of a novice's ability to adjudicate (although Goldman is more of an optimist here than the others.) But when we consider what the scientific consensus is, it becomes less clear how a novice might correctly identify it in the first place. Rather than looking at trust in/reliance on individual scientists, which is how Goldman frames the problem, and which opens up issues with the trustworthiness of an individual scientist considered independently of the institutions that might credential them:

> Recognizing that the aims and values of a given scientist may not cohere with one's own — and that, being people, scientists are as apt as anyone to dissemble or mislead (given the right incentives and character flaws) — might lead one to withhold their epistemic trust.
>
> Such complications at the individual level suggest an alternative locus for the prima facie trustworthiness of science: the scientific community (as a somehow united whole) — or, to construe things more narrowly: scientific consensus (concerning a particular issue). It is at this community level that particular scientific claims are vetted via peer-review and less formal post-peer-review practices. It is at this level that replications are attempted, disputes are prosecuted, papers are cited (positively and critically), results used as a platform for further work, and so on. When "the knowledge machine"...of the scientific enterprise is firing on all cylinders, it is arguably reasonable to identify a kind of social objectivity attached to results on which there is robust scientific consensus... (Slater et al., 2022)

That is, if our trust is in the institutions of science rather than in those of individual scientists, we seem to be able to take advantage of the kinds of intra-community processes within the scientific community that would lead to more robust and reliable knowledge than would be produced by any one individual expert. But this leap requires a certain level of scientific literacy on the part of the novice, a familiarity that we cannot assume that all novices possess (and empirical research suggests that they often don't.)

Slater, et al.'s suggestion is that we should be encouraging this trust in scientific communities rather than individual experts, but that this requires a concomitant commitment to educating the public about how scientific communities work, and what their findings typically look like in that context. Suggesting that scientific institutions adopt a consensus-messaging strategy (CMS):

> ...if a CMS is sophisticated (abjuring a facile identification of scientific consensus and mere agreement), then the results of our study lead us to doubt that it will be effective; if the CMS, on the other hand, takes the simple approach and equates consensus and agreement, then it will either be difficult to employ in a broad range of cases or will transgress the sincerity norm in science communication (for communicators who accept our earlier points, anyway). The conclusion of the previous section was that even if this norm admits of exceptions in certain cases, we need to be cautious about potential downstream consequences for public trust and contributing to a more challenging communication environment overall. (Slater et al., 2022)

Any approach taken to try to better communicate the actual nature of scientific consensus to the general public needs to be carefully considered as there are possible misunderstandings that would lead to less, not more, trust in scientific expertise by giving the public either not enough information, or too much potentially misleading information.

Finally, I want to draw attention to a discussion that was not originally part of the literature on the Novice-Expert Problem as envisaged by Hardwig and Goldman, but was gestured at by Fuller and later discussed (at least in passing) by Levy, De Cruz, Ballantyne and Dentith: the nature of the *difference* in attitudes to the trustworthiness of scientific expertise, or in the standards used to choose specific experts over others within different communities. A person in a Western democracy who is nominally liberal, feminist, educated, middle class, etc probably won't identify the *same* experts as a Trump-voting conservative antivaxxer (for example.) Yet much of the discussion above has assumed that there is one right way to identify experts, and that either everyone should be making the same choices or that the different choices being made are the result of extra-rational factors like political identity/cultural cognition (Kahan, De Cruz, etc) or because the epistemic resources one has access to in virtue of belonging to these communities means that the choices regarding epistemic authority are beyond the awareness of the individual (Levy.) C. Thi Nguyen has offered a slightly different explanation for these differences in the paper "Echo Chambers and Epistemic Bubbles" (C. Thi Nguyen, 2020). In distinguishing between the different *epistemic characteristics* of communities[14], rather than merely the

[14] Levy's *Bad Beliefs* is, purportedly about this same problem, in that he wants to address people who hold said "bad beliefs", which are categorically different from merely "untrue" beliefs and include things like antivax and

content of their beliefs or identity, Nguyen suggests that these differences can shape our assessments of epistemic authority in ways that may or may not be inescapable depending on the *specific* characteristics of the group we are in. Most of us belong to epistemic bubbles, where the information we have available to us is (typically inadvertently) limited or distorted, but where we still have some degree of access to contrary information that might change our views from experts who might not share our beliefs or identity. But some of us instead find ourselves in echo chambers (e.g. a particularly radical political group or a cult), which are typically deliberately constructed to exclude contrary information pre-emptively, meaning that only those experts who are already approved by the community will be taken seriously, reducing or eliminating the possibility of changing our mind independently:

> Compare this process of credence manipulation to the process of omission found in epistemic bubbles. In one standard scenario, I add others as trusted members of my epistemic network based on agreement. I am then less likely to encounter an outside voice — but when I do actually have such an encounter with an outsider, I have no background reason to dismiss them. Bubbles restrict access to outsiders, but don't necessarily change their credibility. Echo chambers, on the other hand, work by offering a pre-emptive discredit towards any outside sources. (C. Thi Nguyen, 2020, p. 151)

This gives a potential model for discussing the differences between communities in terms of their epistemic characteristics, not the mere fact that they happen to believe different things:

> It is important to note that the epistemic mechanisms by which echo chambers work, though problematic, are not *sui generis*. They are perversions of natural, useful, and necessary attitudes of individual and institutional trust. The problem isn't that we trust and distrust groups and institutions. In fact, we must do so. (C. Thi Nguyen, 2020, p. 154)

Nguyen's conclusion is that people within echo chambers have been manipulated into a self-reinforcing, almost inescapable relationship of trust and distrust with epistemic authorities that most of us do not experience. Addressing problematic relationships with epistemic authority rests on our success in rebuilding relationships of trust counter to the strong forces at work within a given community.

It is in this context that I will be developing my own argument: by pulling in strands that focus on trust, reputation and social identity as they relate to the decisions made by novices; by considering how

---

anti-climate change beliefs. But much of the book focuses on a general theory of how epistemic environments shape the beliefs of their members, with noticeably little discussion specifically on the nature of the epistemic differences between groups.

different epistemic communities approach questions of expertise, epistemic authority and truth in different ways that lead to some of the specific, historically contingent phenomena we see in the world today (extremist conspiracy theory groups, antivax communities, etc.) I will also touch on questions of testimony and what it means to have evidence for a belief about a technical topic, and whether we can be said to know something without understanding it.[15]

[15] On testimony, I will be referring to the works of Jennifer Lackey, especially her paper "Testimonial Knowledge and Transmission" (Lackey, 1999) and her book *Learning from Words: Testimony as a source of Knowledge*. (Lackey, 2010) For questions of whether we can know something that we don't believe, I will be looking at the works of Dan Sperber, particularly *Rethinking Symbolism* (Sperber, 1975) and Francois Recanati's response, "Can we believe what we do not understand?" (Recanati, 1997)

Statement of research questions or hypotheses

My central research question is: can novices make epistemically good judgements about experts? My thesis is that they can, given some seemingly uncontroversial assumptions about the nature of expertise and the information available to novices.

Primary research questions:

- What are the epistemic characteristics of an expert or a non-expert?
- What information is typically available to a non-expert to make judgements about whether an expert should be trusted?
- Is the information available to non-experts typically sufficient to make good judgements about when to trust an expert?
- What effect does a non-expert's epistemic context have on their judgements about experts and expert claims?

Auxiliary research questions:

- What is the nature of an epistemic community?
- What are the epistemic characteristics of an epistemic trespasser?
- From the perspective of a non-expert, what information is available that distinguishes between a genuine expert and a trespasser?
- What advice can we provide to non-experts to help them correctly identify experts?

Methodology and forms of analysis employed

My research will involve an extensive reading of the literature and careful analysis of the various positions and arguments that have been put forward therein. I will offer a critique of these positions and develop a position of my own. In doing so, I will apply the approach of regulative epistemology, according to which the study of knowledge can provide practical guidance for inquiry. Regulative epistemology will significantly inform my descriptions of the epistemic nature of the Novice-Expert Problem and its implications for helping the general public make better choices about expertise.

Since my research is not currently expected to involve empirical research, I will not require ethical approval. However, if matters change and I do require empirical research, I will submit the appropriate application to the Division's Human Research Ethics Committee.

Analytical outline for the proposed thesis

Chapter 1 – Defining the Novice-Expert Problem

My main aim in this chapter will be to outline the problem I intend to tackle: that "expert" is an epistemic status, and that a "novice" is (apparently) a person who does not have access to the epistemic facts that determine this status, so must rely on non-epistemic evidence when identifying an expert. The implication is that novices are not able to form well-justified views of who might be an expert, as they must rely on social signalling or other potentially unreliable heuristics, and that this results in people trusting the wrong experts, which has the potential to lead to genuine harm. I describe some approaches to what counts as a well-justified belief about expertise.

Chapter 2 – Experts and Epistemic Communities

One of the key things we rely on when making judgements about expertise is the functioning of the relevant epistemic communities we belong to. I outline a theory of what constitutes an epistemic community that identifies the specific features that relate to questions of expertise, and how these distributed judgements interact with the beliefs of the individuals within the group. I discuss whether this interaction should be treated merely as a kind of testimony and argue that it is better treated as a distinct phenomenon. I note that some philosophers have appealed to the nature of these epistemic communities as a potential solution to the Novice-Expert Problem. I consider those solutions to be only partially correct because they do not distinguish between good and bad epistemic communities, and do not untangle the complex relationship between the epistemic responsibility of a community and that of its members.

Chapter 3 – Experts and Expert Claims

I discuss a key distinction when considering how novices think about experts: the nature of judgements about experts (or groups of experts) by novices, and how novices engage with individual claims that rely on expert knowledge. These two scenarios involve different calculations on the part of the novice, although they do interact with each other in that individual instances of expert claims are what communities use to identify reliable experts, and judgements about a specific claim are more likely to be correct when the novice has access to information about the ongoing reliability of the expert. I argue that judgements about general expert status are more stable judgements that rely on distributed epistemic responsibility across a community and over a period of time. For this reason, they depend more on the quality of the community's epistemic norms. By contrast, judgements about specific claims may occur when the novice has little to no good information available to them to inform their

assessment, and thus require more careful reasoning from the individual, even if the source is otherwise considered trustworthy by the community.

Chapter 4 – Domains of Expertise and Epistemic Trespassing

Experts are experts in relation to a specific domain, which is typically defined by the experts themselves, and the question of whether a given claim is being made by a relevant expert is tangled up with the question of whether a novice is able to identify someone who is trespassing beyond their domain. One possible corollary of the Novice-Expert Problem is that novices cannot spot epistemic trespassers, or charlatans, or well-meaning experts who misjudge the extent of their own expertise. I discuss what information might be available to a novice to avoid this issue and argue that the only real way out of this problem is for individuals to interrogate the epistemic norms of the communities they belong to, since there may or may not be enough direct evidence available to an individual to spot a bad actor. I also argue that these epistemic norms are vital to the health of the beliefs of the individuals within communities, and that bad community epistemic norms, rather than mere false beliefs, are the most salient factor in whether people might end up relying on the wrong experts (or on no experts at all.)

Chapter 5 – What Good Reasoning About Expertise Looks Like

I discuss how epistemic responsibility falls upon both epistemic communities and the individual members of those communities. In making judgements about experts and expert claims, novices heavily rely on others to lighten the cognitive load of policing the reliability of expert sources of information. However, individuals still have a duty to understand how we come to identify experts collectively, and to consider whether the norms our communities use are always fit for purpose, because a novice who relies on a community of people with poor norms about expertise will end up at best unable to trust the people they should be able to rely on, or at worst mislead by others to mistrust good experts and trust bad ones.

Timetable for the project

Timeline / Task	Year 1				Year 2				Year 3				Year 4			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
Complete draft of thesis proposal	✓	✓														
Confirmation meeting			✓													
First draft: Defining the Novice-Expert Problem(s)			✓													
Second draft: Defining the Novice-Expert Problem(s)			✓													
Final draft: Defining the Novice-Expert Problem(s)				✓												
First draft of Experts and Epistemic Communities				✓												
Second draft of Experts and Epistemic Communities					✓											
Final draft of Experts and Epistemic Communities					✓											
First draft of Experts and Expert Claims						✓										
Second draft of Experts and Expert Claims						✓										
Final Draft of Experts and Expert Claims							✓									
First draft of Domains of Expertise and Epistemic Trespassing							✓									
Second Draft of Domains of Expertise and Epistemic Trespassing								✓								
Final Draft of Domains of Expertise and Epistemic Trespassing								✓								
First draft of What Good Reasoning About Expertise Looks Like									✓							

Second Draft of What Good Reasoning About Expertise Looks Like									✓							
Final Draft of What Good Reasoning About Expertise Looks Like										✓						
First draft of Conclusion										✓						
Second draft of Conclusion											✓					
Final draft of Conclusion											✓					
First draft of Introduction												✓				
Second draft of Introduction												✓				
Final draft of Introduction													✓			
Final review and comments													✓			
Final revisions														✓		
Confirmed final version														✓		

References

Baier, A. (1986). Trust and Antitrust. *Ethics*, *96*(2), 231–260.

Ballantyne, N. (2015). THE SIGNIFICANCE OF UNPOSSESSED EVIDENCE. *The Philosophical Quarterly*, *65*(260), 315–335.

Ballantyne, N. (2018). Epistemic Trespassing. *Mind*, *128*(510), 367–395.

Ballantyne, N. (2019). *Knowing Our Limits*. New York, NY, USA: Oxford University Press.

Ballantyne, N., Celniker, J. B., & Dunning, D. (2024). "Do your own research." *Social Epistemology*, *38*(3), 302–317.

Coady, C. A. J. (1973). Testimony and Observation. *American Philosophical Quarterly*, *10*(2), 149–155.

De Cruz, H. (2020). Believing to Belong: Addressing the Novice-Expert Problem in Polarized Scientific Communication. *Social Epistemology*, *34*(5), 440–452.

Dentith, M. R. X. (2018). Expertise and Conspiracy Theories. *Social Epistemology*, *32*(3), 196–208.

Dentith, M. R. X. (2022). Suspicious conspiracy theories. *Synthese*, *200*(3), 243.

Dixson, H. G. W., Komugabe-Dixson, A. F., Medvecky, F., Balanovic, J., Thygesen, H., & MacDonald, E. A. (2022). Trust in science and scientists: Effects of social attitudes and motivations on views regarding climate change, vaccines and gene drive technology. *Journal of Trust Research*, *12*(2), 179–203.

Fuller, S. (2002). *Social Epistemology*. Indiana University Press.

Fuller, S. (2008). *Dissent Over Descent: Intelligent Design's Challenge to Darwinism*. Icon.

Goldman, A. I. (1991). Epistemic Paternalism: Communication Control in Law and Society. *The Journal of Philosophy*, *88*(3), 113–131.

Goldman, A. I. (2001). Experts: Which Ones Should You Trust? *Philosophy and Phenomenological Research*, *63*(1), 85–110.

Hardwig, J. (1985). Epistemic Dependence. *The Journal of Philosophy*, *82*(7), 335–349.

Hardwig, J. (1991). The Role of Trust in Knowledge. *The Journal of Philosophy*, *88*(12), 693–708.

Hume, D. (1748[2012]). *An enquiry concerning human understanding* (S. Buckle, Ed.). Cambridge University Press. https://doi.org/10.1017/cbo9780511808432

Huxster, J. K., Slater, M. H., Leddington, J., LoPiccolo, V., Bergman, J., Jones, M., McGlynn, C., Diaz, N., Aspinall, N., Bresticker, J., & Hopkins, M. (2018). Understanding "understanding" in Public Understanding of Science. *Public Understanding of Science* , *27*(7), 756–771.

Jäger, C. (2024). False Authorities. *Acta Analytica*. https://doi.org/10.1007/s12136-024-00594-3

Kahan, D., & Braman, D. (2005). Cultural Cognition and Public Policy. *Yale Law & Policy Review*, *24*, 5.

Kahan, D. M. (2015). Climate-science communication and the *measurement problem*. *Political Psychology*, *36*(S1), 1–43.

Kahan, D. M., Braman, D., Cohen, G. L., Gastil, J., & Slovic, P. (2010). Who fears the HPV vaccine, who doesn't, and why? an experimental study of the mechanisms of cultural cognition. *Law and Human Behavior*, *34*(6), 501–516.

Kahan, D. M., Jenkins-Smith, H., & Braman, D. (2011). Cultural cognition of scientific consensus. *Journal of Risk Research*, *14*(2), 147–174.

Lackey, J. (1999). Testimonial Knowledge and Transmission. *The Philosophical Quarterly*, *49*(197), 471–490.

Lackey, J. (2010). *Learning from Words: Testimony as a Source of Knowledge*. OUP Oxford.

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

Lambert, K. (2006). Fuller's Folly, Kuhnian Paradigms, and Intelligent Design. *Social Studies of Science*, *36*(6), 835–842.

Levy, N. (2007). Radically Socialized Knowledge and Conspiracy Theories. *Episteme*, *4*(2), 181–192.

Levy, N. (2019). Due deference to denialism: explaining ordinary people's rejection of established scientific findings. *Synthese*, *196*(1), 313–327.

Levy, N. (2021). *Bad Beliefs: Why They Happen to Good People*. Oxford University Press.

Levy, N. (2024). No Trespassing! Abandoning the Novice/Expert Problem. *Erkenntnis. An International Journal of Analytic Philosophy*. https://doi.org/10.1007/s10670-024-00794-8

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

Nguyen, C. T. (2022). Trust as an unquestioning attitude. *Oxford Studies in Epistemology*. https://philpapers.org/rec/NGUTAA

Nguyen, C. Thi. (2020). ECHO CHAMBERS AND EPISTEMIC BUBBLES. *Episteme*, *17*(2), 141–161.

Nguyen, C. Thi. (2021). Was it Polarization or Propaganda? *J. Philos. Res.*, *46*, 173–191.

O'Neill, O. (2002). *A question of trust*. Cambridge University Press.

Recanati, F. (1997). Can we believe what we do not understand? *Mind & Language*, *12*(1), 84–100.

Slater, M. H., Huxster, J. K., & Bresticker, J. E. (2019). Understanding and Trusting Science. *Journal for General Philosophy of Science. Zeitschrift Fur Allgemeine Wissenschaftstheorie*, *50*(2), 247–261.

Slater, M. H., Huxster, J. K., & Scholfield, E. R. (2022). Public Conceptions of Scientific Consensus. *Erkenntnis. An International Journal of Analytic Philosophy*, 1–22.

Slater, M. H., & Scholfield, E. R. (2022). Trust of Science as a Public Collective Good. *Philosophy of Science*, *89*(5), 1044–1053.

Slater, M. H., Scholfield, E. R., & Moore, J. C. (2021). Reporting on Science as an Ongoing Process (or Not). *Frontiers in Communication*, *5*. https://doi.org/10.3389/fcomm.2020.535474

Sperber, D. (1975). *Rethinking Symbolism*. CUP Archive.