



ARTICLE

The Rationality of COVID-19 Vaccine Hesitancy

Joshua Kelsall 

Politics and International Studies, University of Warwick, Coventry CV4 7AL, UK
Email: Josh.Kelsall@warwick.ac.uk

(Received 15 November 2022; revised 26 June 2023; accepted 30 August 2023)

Abstract

Some vaccine-hesitant people lack epistemic trust in the COVID-19 vaccine recommendation that because vaccines have been shown to be medically safe and effective, one ought to get vaccinated. Citing what I call *exception information*, they claim that whatever the general safety and efficacy of vaccines, the vaccines may not be safe and effective for them. Examples include parents citing information about their children's health, pregnant women's concerns about the potential adverse effects of treatment on pregnant women, young people citing their relative invulnerability to extreme COVID-19 symptoms, or members of vulnerable racial groups citing epistemic injustice, such as a lack of representation in COVID-19 vaccine trials. This paper examines the extent to which a lack of epistemic trust in vaccine recommendations, based on such exemption information, is rational.

Keywords: Rationality; epistemic trust; COVID-19; vaccine hesitancy; science communication; communication ethics

1. Introduction

A COVID-19 vaccine that has passed clinical trials and been approved by accredited institutions like the FDA should motivate vaccine uptake. *Ceteris paribus*, if the relevant experts have shown that vaccines are generally safe and effective, then one *ought* to believe that vaccines will be safe and effective for oneself. Indeed, institutions such as the FDA, NHS, WHO, and the CDC make such recommendations (CDC 2022a; FDA 2022; NHS 2022; WHO 2022). Yet some vaccine-hesitant people reject the jump from the general to the individual case, citing what I call *exception information*; information that they believe gives them a reason to believe that a vaccine is unsafe or ineffective *for them*. Examples include parents concerned about their children's health (Vidgen *et al.* 2022); pregnant women's concerns about the adverse effects of treatment for their unborn babies (Firouzbakht *et al.* 2022); young people citing their relative invulnerability to extreme COVID-19 symptoms (Jones *et al.* 2021); or members of vulnerable racial groups citing epistemic injustice, such as a lack of representation in clinical trials (Treweek *et al.* 2020).

© The Author(s), 2023. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

I address whether it is rational to lack epistemic trust in the recommendation that *because* a COVID-19 vaccine is medically safe and effective, one *ought* to get vaccinated. The answer turns on three questions.

- (1) What does it mean to lack epistemic trust?
- (2) What does it mean for epistemic trust to be rational?
- (3) To what extent can laypersons evaluate the epistemic trustworthiness of experts?

To answer (1), I use Pettit (1997) and Lenard's (2008) distinction between distrust and mistrust. Roughly, distrust occurs when one *judges* that a trustee is untrustworthy; mistrust occurs when one is *uncertain* whether a trustee is trustworthy.

To answer (2), I argue that rationality is a normative concept that we use to evaluate the connection between the ways that laypersons form epistemic trust, mistrust, and distrust, and the probability of satisfying their practical and epistemic goals. The relevant practical goal is protecting one's health, either from a dangerous virus, or a dangerous or ineffective vaccine; the relevant epistemic goal is having true beliefs about the safety and efficacy of COVID-19 vaccines.

I consider two answers to (3). First, the second-order approach. According to this, laypersons can only indirectly evaluate expert testimony for features of epistemic trustworthiness (Anderson 2011; Goldman 2001). Second, the first-order approach. According to this, laypersons can evaluate the epistemic trustworthiness of experts (to some extent) through direct testimonial engagement (Lane 2014).

I argue that both the first and second-order approaches have shortcomings in the case of COVID-19, due to the epistemic limitations of laypeople. Thus, rationality puts strong constraints on the extent that which laypersons can rationally evaluate scientists' vaccine recommendations. Nevertheless, I show that in most cases where laypersons have exceptional information, they can rationally mistrust, but not distrust experts. Rationality then requires laypersons to engage with mainstream medical professionals to settle this mistrust.

2. Epistemic trust, distrust, and mistrust

Hardwig (1991) observes that our capacity to know all truths directly is limited, especially the complicated truths in the domain of science. Much of our knowledge of science is predicated on epistemic trust in the testimony of experts. If our trust in experts is to be rational, then we must have reason to believe that our experts are honest, and competent within their domains of expertise (Anderson 2011; Goldman 2001; Hardwig 1991: 700; Keren 2007). When we trust, we are disposed to take trusted experts' testimony to carry more weight than our own, and revise our beliefs accordingly (Kelsall 2021: 294); we also aim to trust those who are trustworthy.

Distrust is the contrary of trust (Hawley 2014). To epistemically distrust an expert is to judge that the expert is either dishonest, incompetent, or wrong. We should distinguish an expert's being wrong from an expert's incompetence, since one can be wrong without necessarily being incompetent.¹ If we are rational, then we are disposed to take distrusted experts' testimony to carry at least no more weight than our own. We aim to distrust those that are untrustworthy.

¹Suppose a mathematician is competent but makes a typo in his calculation, which undermines his proof. One might argue that this is incompetence, but unless one wants to be a perfectionist about competence, this is too strong.

Mistrust is a cautious attitude towards others in which one is uncertain about whether a potential trustee is trustworthy (Lenard 2008: 313; Pettit 1997: 236). It is what Friedman would call an interrogative attitude (2019). Interrogative attitudes have a question as their content, and aim towards resolving this question, thus becoming settled (Friedman 2019: 299). In mistrust, the trustworthiness of the trustee is in question, and the mistruster works towards accurately realising the end goal of trust or distrust. Epistemic mistrust occurs when one is unsure about whether an expert is either honest, competent, or wrong.

Laypersons may trust, distrust, or mistrust two pieces of scientific testimony regarding vaccines. The first is descriptive: it is the claim that some vaccine is generally medically safe and effective. The second is normative: it is the claim that, *ceteris paribus*, because the vaccine is safe and effective, one *ought* to get the vaccine. I focus on this normative claim. I also restrict my analysis to the competence and wrongness conditions of trustworthiness.² Thus, even if laypersons have little rational grounds for epistemic mistrust or distrust on competence grounds, they may have grounds regarding honesty.

3. Rational epistemic trust

Wedgwood's (2017) account of rationality provides a useful way to evaluate laypersons' epistemic trust, distrust, or mistrust in scientists' vaccine recommendations. On this account, we use the concept of rationality to evaluate the ways of thinking that directly guide thinkers at specific times (Wedgwood 2017: 2). 'Ways of thinking' are the mental processes or methods by which thinkers form, revise, and maintain their mental states; in the vaccine context, they would refer to the decision-making procedures that lead an agent to accept or reject vaccination. Direct guidance means that one adopts one's way of thinking because one recognises its rational value and acts in virtue of that recognition. Whether a person is rational is evaluated by comparing (1) a way of thinking that directly guides that thinker at a specific time with (2) the alternative ways of thinking that are available to that thinker at that time (Wedgwood 2017: 236). 'Availability' is contextual (Wedgwood 2017: 150). For our purposes, I define availability in terms of epistemic demandingness and will explain what this amounts to further in Section 4.1.

Although internalist, Wedgwood imposes an external standard for evaluating available ways of thinking: correctness (2017: 209–10). Correctness is defined by the aim of the relevant mental state that is formed, maintained, or revised by a method. For example, Wedgwood accepts the common claim that belief aims at truth (Sosa 2015; Wedgwood 2017: 242; Williams 1973). An agent is rational if their actual way of thinking compares at least equally well to the available alternative methods with respect to their expected degree of incorrectness (Wedgwood 2017: 213). In other words, (for the case of beliefs) one should be directly guided by an available way of thinking that is more likely to yield true beliefs than an available way of thinking (such as random guesses) that does not. Rationality won't guarantee true beliefs. For example, Thales' view that water is the fundamental principle of generation would count as rational even though it is false, insofar as this view is what was most probable given the superficial observations in nature that were best available to Thales in his time.

²Of course, one could think that an expert is wrong because they dishonest. Insofar as this is true, I exclude such cases, focusing instead on cases where wrongness is due to some epistemic error.

Wedgwood's account resembles a common definition of rationality in the context of medical decision-making. Here, we use rationality to evaluate the decision-making procedures that help us reach our goals, the primary goal being to protect our health (Djulgovic *et al.* 2018: 656; Evans 1996; Stanovich 2011). The decision-making procedure is the analogue of Wedgwood's 'ways of thinking' by which thinkers realise their aims. Wedgwood would describe the aim here as realising the desire to protect one's health. Another similarity between this theory and Wedgwood's is that rationality does not guarantee the satisfaction of one's goals, but instead, increases the chances of doing so by making optimal medical decisions.

To know whether vaccine hesitancy is rational, we need to specify the relevant goals of laypersons. Since we are focusing on vaccine hesitancy resulting from safety and efficacy concerns, we can assume that these laypersons have the primary desire to protect their health. Protecting health includes the obvious case of protecting oneself from death from COVID or a vaccine, but also from less life-threatening illness, be it COVID symptoms or vaccine side-effects. Insofar as epistemic trust, distrust, or mistrust in different sources is used to justify one's vaccine acceptance or rejection, we can infer that laypersons have the epistemic aim of trusting epistemically trustworthy sources and distrusting untrustworthy ones. It is important to stress that the epistemic aim is in the service of the practical one; its importance is derived from the fact that to protect themselves from potentially dangerous vaccines, laypersons need to know whether a vaccine is or is likely to be dangerous.

In sum, for a method to be rational in the case of vaccine hesitancy, it must satisfy the following three conditions:

- (1) It will be available to laypersons, i.e., not too epistemically demanding.
- (2) When compared to available alternatives, it is best placed to realise laypersons' practical and epistemic aims, i.e., protecting health, and trusting trustworthy sources.
- (3) It will not sacrifice the practical goal at the expense of the epistemic goal.³

Now we have our rationality framework, we can evaluate methods for forming epistemic trust, mistrust, and distrust in scientists' vaccine recommendations. In the next section, I consider and reject the most prominent of such approaches: the second-order approach.

4. The second-order approach

Supporters of the second-order approach claim that while laypersons lack the competence to evaluate experts' claims directly, they can examine scientists (individuals as well as groups) and their testimony for features that indicate epistemic trustworthiness (Anderson 2011; Goldman 2001; Hardwig 1991). The following list contains common indicators of the competence condition of epistemic trustworthiness as adapted from Anderson (2011), Goldman (2001), and Miller (2013).

³(3) holds because, if trying to satisfy one's epistemic goal came at the expense of the practical goal, the epistemic would fail to serve its fundamental purpose, namely, to protect health. Suppose the method was radically individualistic, such that every individual person must become a virologist and do their own experiments to determine vaccine safety and efficacy of COVID-19 vaccines. Given that pandemics require a quick response to keep them under control, such a method would undermine the practical goal of protecting health in the service of striving for greater epistemic justification.

- (1) Consensus – occurs when there is majority joint acceptance among experts on a (hopefully) approximately true theory within the relevant scientific community.⁴
- (2) Quality assurance – typically refers to the publication of work in well-regarded peer-reviewed journals, but could also involve other quality controls, such as clinical trials.
- (3) Track record – in terms of honesty and competence. Has this researcher (or group/institution) been successful in the past? Do they have a good record of publishing in top journals?
- (4) Prestige – refers to the honours received by a scientist, group, or institution, where these honours signal competence and success.

Placing the second-order approach in our rationality framework, entails the following claims:

- (1) *First-order engagement with experts is unavailable to laypersons.* It is unavailable because laypersons lack the competence to understand and engage with complex science and scientific methodologies (Anderson 2011: 144; Goldman 2001: 90; Lane 2014: 97).
- (2) *The second-order approach is available to laypersons.* Anderson argues that ‘anyone of ordinary education with access to the Web’ can use it, because the relevant information is ‘discoverable within the first few entries of a simple Google search, or in prominent links from these entries’ (Anderson 2011: 144).
- (3) *The second-order approach is the best way for laypersons to rationally trust experts.* In our case, this means that those who follow the approach are more likely to satisfy the epistemic goal (trusting trustworthy sources) and the practical goal (protecting one’s health).

In Section 4.1, I demonstrate that (2) and (3) are false. The second-order approach is unavailable to most laypersons, and it is not the best available method because it is too epistemically demanding; following it is unlikely to help laypersons realise their practical and epistemic aims.⁵ I show that (1) is false in Section 4.

4.1. The unavailability of the second-order approach

In this section, I show that the second-order approach is not available in the case of COVID-19 vaccine science because it is too epistemically demanding. In defining availability in terms of epistemic demandingness, I follow Anderson’s constraint that an approach ought not impose excessive burdens of judgement on laypersons (2011: 144). Another way of framing this is in terms of Miller and Record’s practicability condition, which defines the limit of ‘what can be expected of [an agent] before he attempts

⁴This is Miller’s rough characterisation of consensus (2013: 1298). More specifically, he argues that consensus occurs when experts agree on the same fundamental assumptions, when it is based on varied lines of converging evidence, and when there is social diversity among research communities (2013: 1294). For alternative notions of consensus, see Miller’s (2019) survey. For our purposes, we only require the rough characterisation, since our focus is on the practical question of how laypersons can reasonably know when there is a consensus, which will not entail requiring them to work out directly whether something as complicated as Miller’s three consensus conditions hold as that would be too demanding. Rather, it will be through some kind of proxy such as Anderson’s, discussed in more detail in Section 4.1.

⁵Or at least, it is inferior to the approach I develop in Section 5.2.

a judgement' (2013: 126). There are two components to this; first, the competence of the person in question, which in our case, will require competence in using technology such as the internet and the devices to access it, and at least a basic grasp of relevant scientific concepts (vaccines, viruses, etc.) to understand medical advice. Second, they must have reasonable access to the relevant resources, which means possessing or being able to use the technology in question or having access to medical professionals be it in institutions, online, through other media channels, or through personal relationships. Anderson takes the second-order approach to be available to anyone with an ordinary education and access to the Web (Miller and Record 2013), but I show this to be false in our case.

Although I need not dispute the general applicability of the second-order approach, it is worth noting some general criticisms. Brennan makes two criticisms of the approach. First, he claims that laypersons often lack the necessary knowledge needed to apply the criteria (2020: 230). He argues that some of this knowledge is the kind of knowledge that only insiders in the academic community would know about. For example, while an expert may exhibit competence by working at a prestigious institution or having high-profile publications, this doesn't always give a clear indication of whether this expert is well-regarded by fellow academics, and academics may not always share their scepticism of prestigious academics publicly (Brennan 2020). Brennan notes a further knowledge gap in lay awareness of the standards of scholarly practices, 'in particular, a proper understanding of how the peer-review process works' (Brennan 2020: 231). He demonstrates this by showing how such knowledge gaps allow laypersons to be misled by propaganda (Brennan 2020). A further criticism along these lines is from Keren, who disputes the ability of laypersons to determine whether there is a consensus on a given issue (2018: 790). With Keren et al. (2018), he conducted an empirical study and found that even among students, experience in determining the extent of scientific consensus is limited.

Brennan's second criticism is that laypersons' attempts to interpret the criteria can be thwarted by bias (Brennan 2020: 231). For example, confirmation bias in favour of a vaccine's being dangerous may lead to underestimating of the salience of passing a clinical trial, and an overestimation of prestigious but dissenting experts. Given Brennan's criticisms, we might worry that the intuitive appeal of the second-order approach among philosophers comes from our academic insider status, which makes us blind to the difficulties that outsiders face in using them.

Record and Miller expand on Brennan's point, in a way that attacks Anderson's claim about the ability to easily obtain information online. Unlike Brennan, who focuses on how one's own biases and lack of knowledge can cause one to go awry in applying the second-order criteria, Miller and Record show how algorithmic biases, online filtering in search engines and social media, can result in skewed access to information even on the web. This leads them to conclude, not that online sources are totally unreliable as means to secure information, but that they must be supplemented by corroboration with other sources, either online or off (2013: 131).

A further objection from Lane suggests that current positive appraisals of the second-order approach (such as Anderson's), come from its application to climate change science, where there is a strong established consensus that is easily accessible (Lane 2014: 104). Yet Lane correctly notes that not every case of science has this level of consensus, especially cases of emerging science. I claim that COVID-19 vaccine research is a case of emerging science lacking a clear consensus, and that this

exacerbates difficulties for applying the approach.⁶ This gives us reason to think it is too epistemically demanding. To show this, let's examine vaccine hesitancy in mRNA vaccines. Vaccines using mRNA technology received FDA approval for the first time in the US with emergency approval for Pfizer in 2021 (Beyrer 2021). Concerns about the novelty of mRNA technology are one of the most common predictors of vaccine hesitancy (Wouters *et al.* 2021: 1030). Such concerns are often framed in contrast with other vaccines such as Valneva, Novavax, or Astra-Zeneca, which use traditional vaccine technologies with a well-established track record for success (Mascellino *et al.* 2021).

The status of research for mRNA vaccines is complex and emerging. Although mRNA technology has been used for decades (CDC 2022b), its history reveals no consensus on its safety or efficacy in the past (Dolgin 2021). Moreover, although it has been used for decades, mRNA has not typically been used for vaccines (Beyrer 2021). General expert opinion was that mRNA was too unstable and expensive to be used effectively (Dolgin 2021). While the success of the vaccine rollout may lead one to suppose that there is now a consensus on safety and efficacy, even this is not clear. A recent feature article in the *British Medical Journal* urges the FDA to publish follow-up studies on vaccine safety signals for safety in elderly patients (Demasi 2022). This call comes from the failure of the FDA to follow up its findings of a potential increase in four adverse events in 'elderly people who had Pfizer's COVID-19 vaccine' in July 2021 (Demasi 2022). Putting aside concerns about transparency, and the speculations of dishonesty they invite, uncertainties surrounding the long-term effects of the vaccines show how, even after FDA approval, vaccine science is still ongoing and emerging. The problem with cases of emerging science is that they exacerbate the kind of general worries that Brennan pointed out. The constant shifting status of the science, awareness of gaps in knowledge, and new scientific disputes add to the difficulty in successfully utilising the criteria, especially for laypersons. Arguably, this is made even worse when organisations adopt paternalist strategies of avoiding sharing information, since, when the information inevitably comes out, as in the FDA case above, it invites further speculation of trustworthiness.⁷

In addition to worries about the general ability of laypersons to effectively use the second-order approach, there are worries pertaining to specific groups. In the US, information surrounding the status of COVID-19 vaccine science has been shared predominantly in English, yet 21.9% of the US population speaks a language other than English at home (Treweek *et al.* 2020). Not only does this mean that a significant portion of the population may not have reasonable access to the appropriate information needed, but more broadly, it is one of the factors driving underrepresentation from non-English speaking groups in studies (Treweek *et al.* 2020).

In conclusion, the second-order approach is too epistemically demanding. We have shown how it requires knowledge and understanding of each criterion that is not reasonably available to the average layperson. Moreover, the fact that COVID-19 research is emerging means that laypersons have additional difficulties in keeping abreast with information, and dealing with conflicting authorities, mistakes, and oversights. Therefore, we have reason to doubt that this formulation of the approach is available

⁶This may be less true at present than it was in 2020; however, we can take our analysis at this point and use it to reflect on what rationality requires in similar cases of uncertainty.

⁷I have argued against epistemic paternalism on the grounds that it violates trust norms between laypersons and experts (Kelsall 2021). For views in favour of epistemic paternalism, see Ahlstrom-Vij (2013), and John (2018).

to laypersons in the required sense. In the remainder of this paper, I will propose a simpler approach that is primarily first-order. Though it still involves some second-order considerations, they are not as epistemically demanding because they are in line with laypersons' ordinary medical decision-making practices.

5. The first-order approach

In this section, I present a *prima facie* argument that Lane's (2014) first-order approach is the best available method for epistemic trust. Lane's key claim is that, in some cases, laypersons can evaluate experts' claims directly (Lane 2014: 97). In Section 5.1, I show that while this claim is correct, rationality places strong conditions on the scope of these evaluations and, in most cases, justifies only epistemic mistrust and rarely distrust. In Section 5.2, I provide a modified first-order account with some second-order considerations that avoids the pitfalls of Lane's approach.

Lane describes laypersons' capacity to evaluate expert claims by utilising Collins and Evans (2006) distinction between contributory and interactional expertise. The former describes the ability to participate in an activity and advance its objectives, while the latter describes the ability to discuss and understand talk about the activity, though one is incapable of either contributing to or teaching others to do it (Lane 2014: 102). For Lane, laypersons sometimes have interactional expertise that allows them to evaluate scientists' claims, and failures to do so in such cases can cause misplaced trust. To put this in our rationality framework, the best method for laypersons to satisfy their epistemic and practical aims is one where they use their interactional expertise to evaluate scientific testimony.

We can construct a *prima facie* case for laypersons having the interactional expertise to evaluate vaccine recommendations by looking at Goldenberg (2021) and Cassam's (2021) discussions of vaccine-hesitant parents.⁸ Cassam and Goldenberg argue that laypersons believe they have access to information that allows them to evaluate vaccine recommendations. The first thing to note about hesitant parents is that they are often neither distrustful nor mistrustful of the descriptive claim that vaccines are safe and effective. Instead, they are distrustful or mistrustful of the normative claim that *because* the vaccines are safe and effective, *their* children ought to get vaccinated. Goldenberg notes that vaccine-hesitant parents criticise studies for being too 'broad brush,' implicitly showing a recognition that these claims apply to the general person, but not specifically to *their* kids (2021: 36; Cassam 2021: 6). This makes sense; the parent's practical goal is not to settle general questions of safety and efficacy, but to know whether a vaccine will be safe and effective for individual children.

Vaccine-hesitant parents claim they have reason to doubt that the general descriptive claim is sufficient to justify accepting a vaccine recommendation for their children. According to Goldenberg and Cassam, parents base this claim on their belief that they have expertise or knowledge regarding their children's particular health needs, and that this allows them to evaluate vaccine recommendations. Putting this in Lane's terms, parents' knowledge of their children's health needs gives them the interactional expertise to evaluate whether the recommendation applies to their children.

The vaccine recommendation gets its normative force from the descriptive claim. It is *because* the descriptive claim about safety and efficacy is true, that one *ought to get*

⁸Cassam and Goldenberg focus on vaccine hesitancy for MMR vaccines, but their insights are applicable to our case, as the same kind of vaccine hesitancy persists with COVID-19 vaccines (Vidgen *et al.* 2022).

vaccinated. This descriptive claim is general: it does not mean that the vaccine will be safe and effective for every individual, only that it will be in most cases. I have stipulated that laypersons have the goal of protecting their health, or in this case, their children's health. If one has information, knowledge, or beliefs, that suggest that one may of the unlucky few for whom the vaccine is unsafe or ineffective, then it is rational to refuse a vaccine. It is rational because ignoring information that suggests that a vaccine might unsafe or ineffective is incompatible with a desire to protect one's health. This is an instance of what I termed *exception information* in the introduction. It is information relating to one's individual experiences that give reason to suspect that one may be an exception to the general rule that a vaccine is safe and effective.⁹

There are other instances of exception information. Indeed, considerations of such information often play a significant role in the attitudes of the vaccine hesitant. People who have bad reactions to COVID vaccines are more likely to be hesitant in the future (Geers *et al.* 2022; The Economist/YouGov Polls 2021), as are those concerned that they may be susceptible to an adverse effect (Rief 2021); those who have previously had COVID (Cunliffe *et al.* 2022); and those who believe they are at less risk from contracting serious illness from COVID, such as younger people (Jones *et al.* 2021). In every case, beliefs about health, vaccines, or COVID-19, are used to evaluate vaccine recommendations. Laypeople believe they have exception information that gives them a reason to believe that the vaccine may not be safe and effective *for them*, and this encourages hesitancy. According to the first-order approach, this evaluation is rational because grounded in their interactional expertise.

For a way of thinking to count as rational, it must be available to laypersons, which we understand in terms of epistemic demandingness. The first-order approach requires three things, and all three are at least implicit in the cases of vaccine hesitancy discussed in this section. First, a recognition that the scientists' claim that vaccines are safe and effective is a general rather than individuated claim. Vaccine hesitant parents complaining that science is too 'broad brush' and who require more specific information display such a recognition. Second, access to exception information, typically about one's own health history, past experiences with COVID-19, and experience of vaccines. And third, an integration of that general claim with one's exception information. These conditions are present in the above cases. People claim to have exception information, which they use to argue that the descriptive vaccine claim is insufficient to justify an individual vaccine recommendation. To be rational, however, the approach must also make it probable that laypersons will achieve their practical and epistemic goals when compared to other available alternatives. In Section 5.1, I show that this approach is unlikely to satisfy such a condition.

5.1. The limitations of the first-order approach

The problem with the first-order approach is that it fails to sufficiently restrict the extent to which laypersons can evaluate expert claims with their interactional expertise. There are some clear rational limits. For example, Goldenberg's case of vaccine-hesitant parents who undertook their own research into the connection between mumps, measles, rubella (MMR) vaccines and autism (2021: 35–6) would count as irrational because they presume to have contributory expertise when they do not. However, there are other

⁹Collins would define the kind of knowledge that laypersons have here as experiential rather than the technical knowledge possessed by science (2014: 732).

senses in which laypeople might do their own research which would count as rational. Take those who (like Ronald DeSantis, the Governor of Florida) read up on the latest findings in peer-reviewed journals on COVID vaccines and evaluate scientific recommendations based on this research. In DeSantis' case, he used his knowledge to reject vaccine recommendations for young children (Crist 2022). Or consider laypersons who ignore mainstream sources of information such as the FDA, preferring alternative sources that reject mainstream institutions and experts. Neither of these cases involves presuming contributory expertise. In fact, seeking more information is an attempt to enhance one's interactional expertise, which seems commendable. Another problem is that the approach doesn't tell us whether epistemic mistrust or distrust (or both) is warranted when laypersons make negative evaluations of recommendations. I claim that rationality puts strong constraints on lay evaluations, and that when such evaluations are permissible, they permit mistrust in most cases. Moreover, one has a rational obligation to seek medical advice rather than relying on one's own epistemic capacities, or those of other laypersons or alternative experts, to resolve one's mistrust.

My view rests on a distinction between the information that constitutes lay knowledge and the information that constitutes virologists' expertise. I am committed to the claim that scientific theories in certain disciplines, such as virology, offer a closer approximation of the truth than the folk theories or experiential knowledge of laypersons (Collins 2014; Rowbottom 2019; Sterpetti 2016). I'll motivate the claim by showing that the information constituting virologists' expertise gives them a greater predictive power than laypersons when it comes to resolving safety and efficacy concerns.¹⁰

In Section 5, I suggested that laypersons' exception information is typically about their health, and their past experiences with vaccines and COVID-19. In most cases, this knowledge will be superficial. It will be concerned with the symptoms of a medical issue rather than the underlying biological mechanisms causing them. On the other hand, the virologist has access to those underlying biological mechanisms by which they can better determine whether a vaccine is safe and effective. This information is esoteric and inaccessible to laypersons except through exoteric presentations. Virology requires a wide range of technical knowledge and skills. It requires a thorough knowledge of specialised serological and molecular techniques, for example, antigen, antibody detection, sequencing, and polymerase chain reaction (NHS 2021). It also requires the ability to perform chemical analyses on substances released by viruses when they interact with organic matter, the ability to collect and analyse samples and quantitative data, along with knowledge of how to use relevant lab equipment and tools such as air samplers, collectors, infra-red spectrometers, analysing equipment, and sterilising equipment (Betterteam 2022).

To show the greater predictive power of virology, we can look at two cases and see how both kinds of information fare in resolving safety and efficacy concerns. First, we have the case of psychosomatic symptoms following vaccination. So-called 'nocebo effects,' occur when people's negative expectations and anxieties about being vaccinated can trigger psychosomatic symptoms of the virus (Geers *et al.* 2022). Our second example is the conspiracy theory that vaccines cause rather than protect people from COVID-19 (Saleska and Choi 2021: 823). In both cases, the information accessible to laypersons is unable to resolve uncertainty about the vaccine. In the psychosomatic case, one's experience of the symptoms can appear all too real; it may be indistinguishable from having symptoms that are really triggered by the virus. In the conspiracy

¹⁰The predictive power of science is typically used to demonstrate scientific realism (Chakravartty 2017).

theory case, the inability to distinguish between cases of vaccine side-effects and COVID-19 symptoms drives the conspiracy. When the conspiracy theorist observes a vaccinated person develop COVID-19-like symptoms after vaccination, this is taken as evidence supporting the conspiracy (Saleska and Choi 2021). Both cases show how overreliance on one's lay knowledge can result in false beliefs or uncertainty about vaccines.

On the other hand, the virologist's understanding of vaccines and their ability to examine what is going on at the molecular level gives them a greater ability to distinguish the two cases. A virologist could determine whether a vaccine contains a modified version of the virus (viral vector), a subunit protein, or, as in the case of mRNA vaccines, no virus at all (Mayo Clinic 2022). They can detect whether a vaccine triggers an immune response, and whether it causes the virus to take full effect (Mayo Clinic 2022). It is in this that virologists have greater predictive power than laypersons. The virologists' access to underlying biological structures gives them the capacity to make distinctions that are indistinguishable at the surface level.

As already noted, the information accessible to virologists is *esoteric*. To the extent that it is so, this information is inaccessible to someone without adequate training. Moreover, attempts by scientists or communicators to convey esoteric findings will nearly always be *exoteric*, which is to say they will be presented less rigorously so that they can be grasped by someone without the requisite training. The esoteric nature of virology, and the exoteric understanding possessed by ordinary people, should make us doubt whether laypersons can evaluate experts claims with much veracity. Indeed, those who do their own research, in the sense of reading academic papers, like DeSantis, can be explained as irrational if we acknowledge this point. We can accept that while a layperson can become aware of the conclusions of scientific studies, their understanding will most often remain at the exoteric level and thus, we should doubt that such methods will yield a high chance that the self-researcher will satisfy their practical and epistemic aims.¹¹ Further to this point, Collins notes how reading scientific studies can encourage an overestimation of one's competence on technical scientific questions. He notes that the impersonal passive voice used to convey objectivity in studies can give readers 'the impression of being a "virtual witness" of the experiments described... but the sense of empowerment provided by being a virtual witness, or something similar, is misleading' (2014: 723). To put this in my terms, reading scientific studies can lead one to overestimate the extent to which one knows, or can engage with, scientist esoteric theories or claims.

At best, most attempts to self-educate may result in people becoming *marginal insiders* regarding virology. A marginal insider is someone with an understanding of some fundamental concepts, theories, and scientific methodologies, and minimal experience of testing and refuting their hypotheses (Feinstein 2011: 180). A paradigm marginal insider is the undergraduate student. The problem with marginal insiders is that there is little evidence showing increased competence among them (Osborne and

¹¹Epstein's study of AIDS activists provides a rare exception case where activists were able to develop an esoteric understanding of vaccine science. 'Through a wide variety of methods – including attending scientific conferences, scrutinizing research protocols, and learning from sympathetic professionals both inside and outside the movement – the core treatment activists have gained a working knowledge of the medical vocabulary' (Epstein 1995: 417). Note, however, that this required an extensive engagement with expertise that included access to scientific institutions, journals, and thorough engagement with experts, some of whom also were part of their organisations (Epstein 1995: 418). Thus, while exception cases are possible, they are very rare as it requires a lot more than reading academic papers or articles online.

Dillon 2008). Feinstien suggests that instead of becoming marginal insiders, we had better be *competent outsiders*.¹² The competent outsider is one who has “learned to recognize the moments when science has some bearing on their needs and interests and to interact with sources of scientific expertise in ways that help them achieve their goals” (Feinstein 2011: 180). Given the latter are more likely to recognise and rely on the esoteric understanding of experts when it is relevant to do so, the competent outsider will fare better than the marginal insider (relying on exoteric understanding of esoteric topics), at realising their goals and therefore count as more rational.

Given the differences in predictive power between the information that constitutes lay interactional expertise and virologists’ expertise, we can argue against the first-order approach thus. First, assume laypersons aim at protecting their health, either from COVID-19 or a dangerous vaccine. Satisfying this aim rests on the epistemic aim of having true beliefs about vaccine safety and efficacy. Therefore, if the layperson is rational then they will adopt methods that are more likely to yield true beliefs about vaccine safety and efficacy. The sources that are most likely to yield such beliefs are *esoteric* and within the domain of competence of virologists, as opposed to the exoteric and surface level information accessible to laypersons. Given that laypersons lack access to esoteric information, it is rational for them to epistemically trust those with access to it. To use their lay knowledge to settle this issue is not the best available method, because while exoteric information may be available to them, it is unlikely to satisfy their epistemic and practical goals due to its weak predictive power when contrasted with the esoteric knowledge of experts. The competent outsider, who recognises that virology has a bearing on their practical aims, and who epistemically trusts the relevant experts in virtue of that recognition, is much more likely to satisfy their goals, since they are more disposed to rely on the testimony of those with the best information available.

5.2. A hybrid approach

If the above argument holds, it seems that the first-order approach fails. However, in this section, I show that there is room for limited first-order evaluations, and even second-order evaluations, albeit of a simpler kind than those discussed in Section 3.

First, there are cases where laypeople have access to exception information that, even if exoteric, is enough to justify hesitancy and distrust of vaccine recommendations. These are cases where the layperson’s knowledge of some fact *S* that means that it would be dangerous for them to receive a vaccine. For example, suppose one knows that one is allergic to chemical *x*, and *x* is in a vaccine. *S* can rationally refuse to receive a vaccine and distrust a vaccine recommendation based purely on this knowledge. Even if *S* cannot explain the underlying biological mechanisms that cause the allergy, they know enough to know that they ought not get a vaccine containing *x*. Distrust is rational here since one’s knowledge that a vaccine is harmful is sufficient to justify the judgement that an expert, if they were to advise such a thing (which is of course unlikely), is untrustworthy.

However, most cases of exception information are inconclusive. In the examples discussed in the introduction, of pregnant women, young people, along with people with concerns relating to pre-existing health conditions, past experiences with vaccines, or testimony about others past experiences with vaccines, the exemption information will often require further investigation by experts. In such cases, the layperson does not *know* that a vaccine is unsafe or ineffective; their exemption information only

¹²Keren *et al.* (2018) also adopt this view.

gives them reason to *suspect* that it might be. Indeed, this is precisely the claim that Goldenberg's vaccine-hesitant parents made regarding their vaccine anxieties for their children (2021: 36). In these cases, when one is aware of exemption information, it is rational to mistrust, but not distrust, a vaccine recommendation in the first instance. Distrust would be irrational since the information is inconclusive; however, mistrust is rational because one would be less likely to satisfy one's practical goal of protecting one's health if one dismissed it. To see this, let's consider the example case of pregnant women (and women who plan on becoming pregnant), a group where vaccine hesitancy is more common (Skirrow *et al.* 2022). Pregnant women have two kinds of exemption information. First, they may be concerned that while vaccines have been tested generally, there has been insufficient testing specifically on pregnant women which, given that pregnant women are often excluded from certain treatments due to their condition, gives them reason to mistrust a vaccine that has not been tested on their group.¹³ Relatedly, given that pregnant women are generally supposed to be careful when taking treatments in their condition, they have additional reason to make sure that the vaccine is specifically safe and effective. Pregnant women who disregarded exemption information of either kind would be vulnerable to taking unsafe medications for themselves and their children, thus undermining their health goals. The same is also true in other cases; wherever one has exemption information, it is important that one doesn't overlook it if one is to yield a high chance of taking treatments that are safe and effective.

Although mistrust is rational in the first instance, rationality also requires laypersons to engage with medical professionals to settle their mistrust.¹⁴ It is important that engagement with medical professionals amounts to more than the layperson being simply told, without justification or reason, that the vaccines are safe and/or effective. As we have seen, laypersons consistently bring up being dismissed by scientists and doctors as reasons for distrust. Suppose a woman is concerned that she may have a miscarriage after being vaccinated because this happened to a friend of hers. She wants to know whether the vaccine poses a risk to her unborn child. If a doctor dismisses her exception information without reason or evidence, then it is reasonable for the woman to continue to mistrust, since her reason for mistrust has not been addressed. It may be enough for someone with staunch trust in the medical profession to take a doctor's word, but, as I have argued elsewhere, it is important for experts to earn, rather than presume, public trust by engaging with their concerns (Kelsall 2021). On the other hand, if the doctor provides evidence and reasons why vaccination is safe, by pointing to clinical trials on pregnancy and vaccines, this provides a counter to that exception information that a rational person ought to accept. This is because, as argued in Section 5.1, this information is exoteric and has weak predictive power when compared to the esoteric knowledge accessible to medical professionals.

Returning to Freidman's work on competent outsiders and marginal insiders, we can say that the rational person is a competent outsider, who recognises that expertise has a

¹³As an example of how COVID-19 vaccines specifically affect women, consider that the vaccines were found to cause changes to menstrual cycle length (Edelman *et al.* 2022; Royal College of Obstetricians and Gynaecologists 2021).

¹⁴It isn't contradictory to say that laypersons are required to engage with experts when we have said that mistrust is appropriate. Only a narrow form of mistrust in testimony about a vaccine recommendation is rational; therefore, it is consistent with narrow mistrust in this domain, to have epistemic trust in other domains.

bearing on their goals when it comes to vaccine decision-making and in particular understanding their exception information. This returns us to the problem of how laypersons can identify the appropriate experts. It is here that second-order considerations can play a role in determining which experts' laypersons ought to consult. I propose that rather than engaging in messy debates and research trying to find out whether there is a consensus and checking the credentials and track records of various experts and institutions, laypersons ought instead to refer to the testimony of their local medical professionals. Such a consideration is readily accessible to laypersons because it is part and parcel with people's ordinary medical practices of seeking local medical professionals when opting for treatments. For most societies, it is also part of the standard epistemic and professional division of labour. Indeed, this recognition is not only clear in the fact that most people go to their doctor's when they have medical complaints, but even clear in the actions of many vaccine-hesitant people. Although there are some fringe conspiracy theorists with pervasive distrust in all public institutions, many vaccine-hesitant people often become so only after failed attempts to consult frontline medical professionals, where they are dismissed, sometimes rudely, which results in distrust (Evans *et al.* 2001; Goldenberg 2021: 35; Kirby 2006; Leach 2005: 8). This in fact, is a familiar story in the public understanding of science literature; many laypersons begin by trusting mainstream sources, only to become distrustful when their local knowledge and concerns are disregarded as unscientific, irrelevant, or irrational (Irwin and Wynne 1996).

A second-order requirement to follow ordinary social practises is a preferable second-order method since it is closer in keeping with ordinary medical advice seeking behaviour and does not require understanding of complex academic processes such as consensus, peer-review, prestige, and track-record. Given that it falls within ordinary practice, it is an approach that is not overly epistemically demanding and thus is available to laypersons. While there may be some residual difficulties with respect to availability, for instance, if one does not speak the native language of the country in which one lives and accessing medical professionals who speak one's language is difficult or impossible, or if one has no way of reaching a medical professional either in person or virtually, these cases will be in the minority.

There is a further, more practical, reason in favour of this simplified approach. Research suggests that vaccine hesitancy is generally tied to institutional distrust, either in governments, or aloof organisations such as the WHO, while trust in local medical professionals is stronger (Rozek *et al.* 2021). This is especially true of marginalised groups, whose distrust is often targeted at the institutional level rather than at local medical practitioners (Lockyer *et al.* 2021; Reid and Mabhala 2021). Given that front line medical professionals have greater public trust and given that seeking medical advice at the local level is already part of people's ordinary health practices, frontline, local medical professionals are best placed to share this information and to address patient exception information.

Of course, there will be people in society who, for whatever reason, fail to follow these norms or even recognise them. Hardline conspiracy theorists with unshakable trust in public institutions at every level, or people trapped in echo chambers and epistemic bubbles may find it impracticable, due to their deep-seated distrust and limited access to information (if they are in a bubble or echo chamber), to follow mainstream medical advice from their doctors. I am willing to concede that such people may be beyond hope, but it is worth noting that the actual prevalence of such people is very

small. And the number of people who live in echo chambers and epistemic bubbles has been overstated.

Dubois and Blank (2018) note that much of the empirical data on the influence of echo chambers is contradictory and suffers from narrow analysis. For example, studies often focus on single media platforms, whether that is social media (Facebook, Twitter, etc.) or different news sites, or search engines. The problem is that in practice most people receive their news in multi-media contexts. Dubois and Blank show that ‘the greater the number of media a citizen uses the more opportunity to be exposed to different political opinions and news...’ and thus, ‘the less likely they are to be in an echo chamber’ (Dubois and Blank 2018: 734). They provide three reasons for this: first, they note that even people of different political persuasions, such as Republicans and Democrats, have similar media diets; they note that even those who are more restricted often incidentally come across opposing views, either on other sites or in interactions with others.¹⁵ They conclude ‘whatever may be happening on any single social media platform, when we look at the entire media environment, there is little apparent echo chamber’ (Dubois and Blank 2018: 740). They found that only 8% of people had both a low media diet and lack of interest in politics, meaning that they are more susceptible to echo chambers. However, they note that even these groups typically follow the views of so-called opinion leaders within their groups, these being people who are politically interested and who do consume a wide variety of media, meaning that even among these 8%, it isn’t necessarily the case that they will be in a dangerous echo chamber.

In light of this research, I would say that there may be very rare cases in which one is truly locked in an echo chamber or epistemic bubble in such a way that one is unaware or unable to recognise the standard division of epistemic labour in society. However, this isn’t the case for the average person who typically goes to the doctor when they have a health complaint, who consumes a wide variety of media and political content, and whose concerns about vaccines are specific, such as anxiety about mRNA vaccines, clinical trials, or any of the cases of exemption information discussed in this paper.

In conclusion, Lane’s claim that it is possible for laypeople evaluate expertise, is correct. However, the extent of these evaluations is limited in the case of vaccine science. It is rational for laypersons to mistrust vaccine recommendations if they possess exception information. However, since we assume that laypersons have the epistemic aim of forming true beliefs about safety and efficacy, and we have argued that scientific knowledge provides a closer approximation to truth in those beliefs, the layperson is rationally required to engage with medical experts to resolve this mistrust, as opposed to doing it themselves. To do one’s own research is unavailable to laypersons because the esoteric information is inaccessible to them. Moreover, it is unlikely for laypersons to satisfy their practical and epistemic goals given the inability of exoteric lay knowledge to draw reliable conclusions about vaccine safety and efficacy.

I don’t think this view is especially controversial, as I have suggested, it aligns with ordinary people’s medical decision-making practices. One objection to the approach of using local avenues to form epistemic trust and share vaccine information is that it is inappropriate in the case of a public health problem, which requires wider institutional solutions. A second objection is that frontline medical professionals are technically not

¹⁵A study by Andrew Guess (2021) also arrived at similar results, finding a 65% overlap between republican and democrat’s media diets in 2015, and a 50% overlap in 2016.

the most salient experts on vaccine safety and efficacy. In response to the first point, I acknowledge that the pandemic is a public health problem, but that does not preclude the exploitation of local channels of communication, especially when these local channels are more trusted than broader institutional communication strategies.

My response to the second objection is that while it is true that local medical professionals are not virologists, it does fall within the ordinary duties of doctors to help their patients make informed medical decisions that will have a positive impact on their health (Rizo *et al.* 2002: 711). Thus, it is not unreasonable to expect doctors to be given the relevant information to address patient concerns, and to be willing to engage in such conversations with patients when they bring up concerns. What this means, in practical terms, is that frontline medical professionals are equipped with up-to-date information on vaccine safety and efficacy by the relevant institutions, and that when laypersons raise concerns about exemption information, doctors take those concerns at face value and engage with their patients. Even if, say, no studies for a given vaccine have been done on pregnant women yet being open and transparent about that fact, or being open and transparent about uncertainties, displays honesty and therefore trustworthiness. While one might worry that such openness might increase distrust, recent psychological research into suggests that such effects are only marginal (van der Bles *et al.* 2020), and sometimes that the reverse is true (Peterson *et al.* 2021).

6. Conclusion

In this paper, I argued that there are two ways that one can lack epistemic mistrust. One can distrust, in which case one judges that one's trustee is either incompetent or dishonest. Or one can mistrust, in which case one is uncertain about the competence or honesty of the trustee.

To rationally trust a vaccine recommendation, the layperson must use a way of thinking or decision-making procedure that, when compared to the available alternative ways of thinking, is most likely to help realise their goals. I focused on cases of vaccine hesitancy which result from concerns about the safety and efficacy of vaccines and inferred from this that these laypersons have the goal of protecting their health either from a dangerous vaccine or a dangerous virus. From that, we derived an epistemic aim of laypersons, which is to have true beliefs about the safety and efficacy of vaccines.

On the assumption of those goals, I argued that it is possible for laypersons to evaluate experts' vaccine recommendations in Lane's first-order sense. They can evaluate a vaccine recommendation if they have access to exemption information, which is information that gives them a reason to believe that a vaccine may not be safe and effective for them. While this evaluation is possible, I noted that, given the exoteric status of exemption information, it is (in most cases) not strong enough to justify epistemic distrust in a vaccine recommendation, though it is strong enough to generate mistrust in that recommendation. However, I argued that laypersons with such mistrust are rationally required to engage with medical professionals (who have access to esoteric information that is 'closer' to the truth), to resolve their mistrust.¹⁶

¹⁶Thank you to Sven Bernecker, Teresa Branch, Tom Sorell, and Matthew Green for discussions and thoughtful feedback on this paper. Also thank you to participants of the 2022 Dalehead Conference, and colleagues at the University of Warwick Moral Philosophy Reading Group, for their comments on my presentation of the paper.

References

- Ahlstrom-Vij Kristoffer.** (2013) *Epistemic paternalism: a defence* (New York: Palgrave-Macmillan).
- Anderson E.** (2011). 'Democracy, Public Policy, and Lay Assessments of Scientific Testimony.' *Episteme* 8(2), 144–64.
- Asay J.** (2019). 'Going Local: A Defense of Methodological Localism about Scientific Realism.' *Synthese* 196, 587–609. <https://doi.org/10.1007/s11229-016-1072-6>.
- Betterteam.** (2022). 'Virologist Job Description.' <https://www.betterteam.com/virologist-job-description> (accessed 22 September 2022).
- Beyrer C.** (2021). 'The Long History of mRNA Vaccines.' *John Hopkins Bloomberg School of Health*. <https://publichealth.jhu.edu/2021/the-long-history-of-mrna-vaccines>.
- Brennan J.** (2020). 'Can Novices Trust Themselves to Choose Trustworthy Experts? Reasons for (Reserved) Optimism.' *Social Epistemology* 34(3), 227–40. <https://doi.org/10.1080/02691728.2019.1703056>.
- Cassam Q.** (2021). 'Misunderstanding Vaccine Hesitancy: A Case Study in Epistemic Injustice.' *Educational Philosophy and Theory*, 55(3), 315–329. <https://doi.org/10.1080/00131857.2021.2006055>.
- CDC.** (2022a). 'Benefits of Getting a Vaccine.' <https://www.cdc.gov/coronavirus/2019-ncov/vaccines/vaccine-benefits.html#:~:text=COVID%2019%2Dvaccines%20are%20effective,best%20protection%20against%20COVID%2D19> (accessed 02 October 2022).
- CDC.** (2022b). 'Understanding mRNA Vaccines.' *CDC: Centres for Disease Control and Prevention*. <https://www.cdc.gov/coronavirus/2019-ncov/vaccines/different-vaccines/mrna.html>.
- Chakravarty A.** (2017). 'Scientific Realism.' In Edward N. Zalta(ed), *The Stanford Encyclopedia of Philosophy (Summer 2017 Edition)*. Stanford: Stanford University <https://plato.stanford.edu/archives/sum2017/entries/scientific-realism>.
- Collins H.** (2014). 'Rejecting Knowledge Claims Inside and Outside Science.' *Social Studies of Science* 44(5), 722–35. <http://www.jstor.org/stable/43284247>.
- Collins H. and Evans R.** (2006). *Rethinking Expertise*. Chicago, USA: Chicago University Press.
- Crist C.** (2022). 'Florida Governor Rejects COVID Vaccines for Young Kids.' *WEBMD NEWS BRIEF*. <https://www.webmd.com/vaccines/COVID-19-vaccine/news/20220617/florida-governor-rejects-covid-vaccines-young-kids> (accessed 07 September 2022).
- Cunliffe K., Richardson Z. and Flower G.** (2022). 'Characteristics Associated With Third Vaccination Uptake.' *Office for National Statistics*. <https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/conditionsanddiseases/articles/coronaviruscovid19infectionsurveytechnicalarticlecharacteristicsassociatedwiththirdvaccinationuptake/21april2022> (accessed 01 August 2022).
- Demasi M.** (2022). 'FDA Urged to Publish Follow-Up Studies on COVID-19 Vaccine Safety Signals.' *British Medical Journal* 379, 1–3. <https://doi.org/10.1136/bmj.o2527>.
- Djulgovic B., Elqayam S. and Dale W.** (2018). 'Rational Decision Making in Medicine: Implications for Overuse and Underuse.' *Journal of Evaluation in Clinical Practice* 24, 655–65. <https://doi.org/10.1111/jep.12851>.
- Dolgin E.** (2021). 'The Tangled History of mRNA Vaccines.' *Nature News Feature*. <https://www.nature.com/articles/d41586-021-02483-w> (accessed 22 September 2022).
- Dubois E. and Blank G.** (2018). 'The Echo Chamber is Overstated: The Moderating Effect of Political Interest and Diverse Media.' *Information Communication & Society* 21(5), 729–45.
- The Economist/YouGov Polls** (2021). 'Why Won't Americans Get Vaccinated Poll Data.' *Yougov*. <https://today.yougov.com/topics/politics/articles-reports/2021/07/15/why-wont-americans-get-vaccinated-poll-data> (accessed 01 August 2022).
- Edelman A., Boniface E.R., Male V., Cameron S.T., Benhar E., Han L., Matteson K.A., Van Lamsweerde A., Pearson J.T. and Darney B.G.** (2022). 'Association Between Menstrual Cycle Length and COVID-19 Vaccination: Global, Retrospective Cohort Study of Prospectively Collected Data.' *BMJ Medicine* 1, e00029. <https://doi.org/10.1136/bmjmed-2022-000297>.
- Epstein S.** (1995). 'The construction of lay expertise: AIDS activism and the forging of credibility in the reform of clinical trials.' *Science, Technology, and Human Values* 29(4), 408–437.
- Evans J.S.B.T.** (1996). *Rationality and Reasoning*. Hove, UK: Psychology Press.
- Evans M., Soddart H., Condon L., Freeman E., Grizzell M. and Mullen R.** (2001). 'Parents' Perspectives on the MMR Immunisation: A Focus Group Study.' *British Journal of General Practice* 51(472), 904–10.

- FDA. (2022). 'COVID-19 Frequently Asked Questions.' <https://www.fda.gov/emergency-preparedness-and-response/coronavirus-disease-2019-COVID-19/COVID-19-frequently-asked-questions#general> (accessed 02 October 2022).
- Feinstein N. (2011). 'Salvaging Science Literacy.' *Science Education* **95**, 168–85. <https://doi.org/10.1002/sce.20414>.
- Firouzbakht M., Sharif N., Kazeminavaei F. and Rashidian P. (2022). 'Hesitancy About COVID-19 Vaccination Among Pregnant Women: A Cross-Sectional Study Based on the Health Belief Model.' *British Medical Journal of Pregnancy & Childbirth* **22**(611), 1–9. <https://doi.org/10.1186/s12884-022-04941-3>.
- Friedman J. (2019). 'Inquiry and Belief.' *Noûs* **53**, 296–315. <https://doi.org/10.1111/nous.12222>.
- Geers A.L., Clemens K.S., Faasse K., Colagiuri B., Webster R., Vase L., Sieg M., Jason E. and Colloca L. (2022). 'Psychosocial Factors Predict COVID-19 Vaccine Side Effects.' *Psychotherapy & Psychosomatics* **91**, 136–38. <https://doi.org/10.1159/000519853>.
- Goldenberg M. (2021). *Vaccine Hesitancy: Public Trust, Expertise, and the War on Science (Science, Values, and the Public)*. Pittsburgh, USA: University of Pittsburgh Press.
- Goldman A.I. (2001). 'Experts: Which Ones Should You Trust?' *Philosophy and Phenomenological Research* **63**(1), 85–110.
- Guess A.M. (2021). '(Almost) Everything in Moderation: New Evidence on Americans' Online Media Diets.' *American Journal of Political Science* **65**, 1007–22. <https://doi.org/10.1111/ajps.12589>.
- Hardwig J. (1991). 'The Role of Trust in Knowledge.' *Journal of Philosophy* **88**(12), 693–708.
- Hawley K. (2014). 'Trust, Distrust and Commitment.' *Noûs* **48**, 1–20. <https://doi.org/10.1111/nous.12000>.
- Irwin A. and Wynne B. (1996) *Misunderstanding Science? The public reconstruction of science and technology* (New York: Cambridge University Press).
- John S. (2018). 'Epistemic trust and the ethics of science communication: against transparency, openness, sincerity, and honesty.' *Social Epistemology*, **32**(2), 75–87.
- Jones B., Wardman L. and Tinkler L. (2021). 'Coronavirus Vaccine Hesitancy in Younger Adults: June 2021.' *Office for National Statistics*. <https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/healthandwellbeing/articles/coronavirusvaccinehesitancyinyoungeradults/june2021#main-points> (accessed 26 July 2022).
- Kelsall J. (2021). 'The Trust-Based Communicative Obligations of Expert Authorities.' *Journal of Applied Philosophy* **38**, 288–305. <https://doi.org/10.1111/japp.12465>.
- Keren A. (2007). 'Epistemic Authority, Testimony and the Transmission of Knowledge.' *Episteme* **4**, 368–81.
- Keren A. (2018). 'The Public Understanding of What? Laypersons' Epistemic Needs, the Division of Cognitive Labor, and the Demarcation of Science.' *Philosophy of Science* **85**(5), 781–92. <https://doi.org/10.1086/699690>.
- Keren A., Liviatan I. and Barzilai S. (2018). *Searching for the Scientific Consensus: A Productive Path for Belief Change about Global Warming*. Unpublished manuscript, University of Haifa.
- Kirby D. (2006). *Evidence of Harm: Mercury in Vaccines and the Autism Epidemic*. New York: St. Martin's Griffin.
- Lane M. (2014). 'When the Experts Are Uncertain: Scientific Judgement and the Ethics of Democratic Judgement.' *Episteme* **11**(1), 97–118. <https://doi.org/doi:10.1017/epi.2013.48>.
- Leach M. (2005). *MMR Mobilisation: Science and Citizens in a British Vaccine Controversy*. Brighton, UK: Institute of Development Studies. <http://www.drc-citizenship.org/system/assets/1052734466/original/1052734466-leach.2005-mmr.pdf?1289493597>.
- Lenard P.T. (2008). 'Trust Your Compatriots, but Count Your Change: The Roles of Trust, Mistrust and Distrust in Democracy.' *Political Studies* **56**, 312–32.
- Lockyer B., Islam S., Rahman A., Dickerson J., Pickett T., Wright J., McEachan R. and Sheard L. (2021). 'Understanding COVID-19 Misinformation and Vaccine Hesitancy in Context: Findings From a Qualitative Study Involving Citizens in Bradford UK.' *Health Expectations* **24**, 1158–67. <https://doi.org/10.1111/hex.13240>.
- van der Bles A.M., van der Linden S., Freeman A.L.J. and Spiegelhalter D.J. (2020). 'The Effects of Communicating Uncertainty on Public Trust in Facts and Numbers.' *Proceedings of the National Academy of Sciences* **117**(14), 7672–83.
- Mascellino M.T., Di Timoteo F., De Angelis M. and Oliva A. (2021). 'Overview of the Main Anti-SARS-CoV-2 Vaccines: Mechanism of Action, Efficacy and Safety.' *Infectious Drug Resistance* **31** (14), 3459–76. <https://doi.org/10.2147/IDR.S315727>.

- Mayo Clinic.** (2022). 'Different Types of COVID-19 Vaccines and How They Work.' <https://www.mayoclinic.org/diseases-conditions/coronavirus/in-depth/different-types-of-COVID-19-vaccines/art-20506465> (accessed 08 July 2022).
- Miller B.** (2013). 'When is Consensus Knowledge Based? Distinguishing Shared Knowledge from Mere Agreement.' *Synthese*, **190**(7), 1293–1316.
- Miller B.** (2019). The Social Epistemology of Consensus. In Henderson David, Peter Graham, Miranda Fricker, Nikolaj Pedersen (eds), *The Routledge Handbook of Social Epistemology*, pp. 228–237. New York: Routledge.
- Miller B and Record I.** (2013). 'Justified Belief in a Digital Age: On the epistemic implications of secret internet technologies.' *Synthese*, **10**(2), 117–134.
- NHS.** (2021). *Virology (Healthcare Scientist)*. <https://www.healthcareers.nhs.uk/explore-roles/healthcare-science/roles-healthcare-science/life-sciences/virology-healthcare-scientist> (accessed 02 October 2022).
- NHS.** (2022). 'Coronavirus (COVID-19) Vaccine.' *NHS Website*. <https://www.nhs.uk/conditions/coronavirus-COVID-19/coronavirus-vaccination/coronavirus-vaccine/#:~:text=Everyone%20aged%205%20and%20over,dose%20before%20any%20booster%20doses> (accessed 01 August 2022).
- Osborne J. and Dillon J.** (2008). *Science Education in Europe: Critical Reflections*. London: Nuffield Foundation.
- Peterson M., Bor A., Jorgensen F. and Lindholt M.** (2021). 'Transparent Communication About Negative Features of COVID-19 Vaccines Decreases Acceptance But Increases Trust.' *Proceedings of the National Academy of Sciences* **118**(29), 1–8.
- Pettit P.** (1997). *Republicanism: A Theory of Freedom and Government*. Oxford: Oxford University Press.
- Reid J.A. and Mahbala M.A.** (2021). 'Ethnic and Minority Group Differences in Engagement With COVID-19 Vaccination Programmes – at Pandemic Pace; When Vaccine Confidence in Mass Rollout Meets Local Vaccine Hesitancy.' *Israel Journal of Health Policy Research* **10**, 33. <https://doi.org/10.1186/s13584-021-00467-9>.
- Rief W.** (2021). 'Fear of Adverse Effects and COVID-19 Vaccine Hesitancy: Recommendations of the Treatment Expectation Expert Group.' *Journal of the American Medical Association Health Forum* **2**(4), 1–3. <https://doi.org/10.1001/jamahealthforum.2021.0804>.
- Rizo C.A., Jadad A.R. and Enkin M.** (2002). 'What's a Good Doctor and How Do You Make One? Doctors Should Be Good Companions for People.' *British Medical Journal* **325**(7366), 711.
- Rowbottom D.P.** (2019). 'Scientific Realism: What It Is, the Contemporary Debate, and New Directions.' *Synthese* **196**(2), 451–84.
- Royal College of Obstetricians and Gynaecologists.** (2021). 'RCOG/FSRH Responds to Reports of 30,000 Women's Periods Affected After COVID-19 Vaccine.' (accessed 16 June 2022).
- Rozek L.S., Jones P., Menon A., Hicken A., Apsley S. and King E.J.** (2021). Understanding vaccine hesitancy in the context of COVID-19: the role of trust and confidence in a seventeen country survey.' *International Journal of Public Health*, **66**, 636255.
- Saleska J.L. and Choi K.R.** (2021). 'A Behavioral Economics Perspective on the COVID-19 Vaccine Amid Public Mistrust.' *Translational Behavioral Medicine* **11**(3), 821–25. <https://doi.org/10.1093/tbm/ibaa147>.
- Skirrow H., Barnett S., Bell S., Riaposova L., Mounter-Jack S., Kampmann B. and Holder B.** (2022). 'Women's Views on Accepting COVID-19 Vaccination During and After Pregnancy, and For Their Babies: A Multi-methods Study in the UK.' *British Medical Journal of Pregnancy & Childbirth* **22**(33), 1–15. <https://doi.org/10.1186/s12884-021-04321-3>.
- Sosa E.** (2015). *Judgment and Agency*. Oxford: Oxford University Press.
- Stanovich K.E.** (2011). *Rationality and the Reflective Mind*. Oxford: Oxford University Press.
- Sterpetti F.** (2016). 'Scientific Realism, the Semantic View and Evolutionary Biology.' In E. Ippoliti, F. Sterpetti and T. Nickles (eds), *Models and Inferences in Science*, pp. 55–76. Springer.
- Treweek S., Forouhi N., Narayan V. and Khunti K.** (2020). 'COVID-19 and Ethnicity: Who Will Research Results Apply to?' *The Lancet* **395**(10242), 1955–57.
- Vidgen B., Taylor H., Bereskin C. and Margetts H.** (2022). 'Understanding Parental Concerns About COVID-19 Vaccination.' *The Alan Turing Institute*. https://www.turing.ac.uk/sites/default/files/2021-09/parental_concerns_about_c-19_vaccination_1.pdf (accessed 08 October 2022).
- Wedgwood R.** (2017). *The Value of Rationality*. Oxford: Oxford University Press.
- WHO.** (2022). 'COVID-19 Advice for the Public: Getting Vaccinated.' *World Health Organisation*. <https://www.who.int/emergencies/diseases/novel-coronavirus-2019/COVID-19-vaccines/advice> (accessed 01 August 2022).

Williams B. (1973). 'Deciding to Believe.' In B. Williams (ed.), *Problems of the Self*, pp. 136–51. Cambridge, USA: Cambridge University Press.

Wouters SK., Salcher-Konrad M., Pollard A., Larson H., Teerawattanon Y. and Jit M. (2021). 'Challenges in Ensuring Global Access to COVID-19 Vaccines: Production, Affordability, Allocation, and Deployment.' *The Lancet* **397**(10278), 1023–34. [https://doi.org/10.1016/S0140-6736\(21\)00306-8](https://doi.org/10.1016/S0140-6736(21)00306-8).

Joshua Kelsall is a postdoctoral research fellow at the University of Warwick. He completed his Ph.D. at the University of St Andrews. Currently, he works on the philosophy of trust, with a focus on epistemic trust. He is currently part of the ARCH-DFG funded project *Moral Obligation and Epistemology: The Case of Vaccine Hesitancy*. The project explores the rationality and reasonableness of vaccine refusal, as well as the ethics of vaccine hesitancy. Preferred Email: Josh.Kelsall@warwick.ac.uk

Cite this article: Kelsall J (2023). The Rationality of COVID-19 Vaccine Hesitancy. *Episteme* 1–20. <https://doi.org/10.1017/epi.2023.47>