



Trusting scientific experts in an online world

Kenneth Boyd¹ 

Received: 30 October 2020 / Accepted: 25 January 2022 / Published online: 18 February 2022
© The Author(s), under exclusive licence to Springer Nature B.V. 2022

Abstract

A perennial problem in social epistemology is the problem of expert testimony, specifically expert testimony regarding scientific issues: for example, while it is important for me to know information pertaining to anthropogenic climate change, vaccine safety, Covid-19, etc., I may lack the scientific background required to determine whether the information I come across is, in fact, true. Without being able to evaluate the science itself, then, I need to find trustworthy expert testifiers to listen to. A major project in social epistemology has thus become determining what the markers of trustworthiness are that laypersons can appeal to in order to identify and acquire information from expert testifiers. At the same time, the ways in which we acquire scientific information has changed significantly, with much of it nowadays being acquired in online environments. While much has been said about the potential pitfalls of seeking information online (e.g. the prevalence of filter bubbles, echo chambers, and the overall proliferation of “fake news”), little has been said about how the nature of seeking information online should make us think about the problem of expert testimony. Indeed, it seems to be an underlying assumption that good markers of trustworthiness apply equally well when seeking information from expert testifiers in online and offline environments alike, and that the new challenges and opportunities presented by online environments merely affects the methods by which we can acquire evidence of said trustworthiness. Here I argue that in making this assumption one risks failing to account for how unique features of the ways in which we acquire information online affect how we evaluate the trustworthiness of experts. Specifically, I argue for two main claims: first, that the nature of information-seeking online is such that the extent to which information is susceptible to manipulation is a dominant marker of trustworthiness; second, as a result, one will be more likely to seek out a particular kind of expert testifier in online environments, what I call a cooperative as opposed to preemptive expert. The result is that criteria for expert trustworthiness may look significantly different when acquiring information online as opposed to offline.

✉ Kenneth Boyd
kenneth.boyd@gmail.com

¹ Department for the Study of Culture, University of Southern Denmark, Campusvej 55, 5230 Odense, Denmark

Keywords Expert testimony · Scientific testimony · Online communication · Social epistemology

A perennial problem in social epistemology is the *problem of expert testimony*, specifically expert testimony regarding scientific issues¹: for example, while it is important for me to know information about anthropogenic climate change, vaccine safety, etc., I may lack the scientific background required to determine whether the information I come across is true. One way I could address this problem would be by learning more about science, but I don't really have the time to do this, and perhaps also lack the capacity. Without being able to evaluate the science itself, I need to find trustworthy expert testifiers to listen to. We can then ask: what are the markers of trustworthiness that laypersons can appeal to in order to identify, and acquire information from, expert testifiers? The idea is that while laypersons may not be able to evaluate the content of expert testimony itself, they can at least evaluate evidence for the trustworthiness of the testifier.

At the same time, the ways in which we acquire scientific information have changed significantly, with much of it nowadays being acquired in *online environments* (Takahashi & Tandoc Jr., 2016; Cinelli et al., 2020). Online sources come in many different forms: some look like traditional sources transposed into a new setting—e.g. journals and books that are available online, news websites that report scientific information, etc.—while others are unique to online environments—e.g. blogs, message board sites, social media, etc. While much has been said about the potential pitfalls of seeking information online (e.g. the prevalence of filter bubbles, echo chambers, and the overall proliferation of “fake news”²), little has been said about how the nature of seeking information online should make us think about the problem of expert testimony. Indeed, while arguments for criteria for expert trustworthiness do not necessarily exclude discussions of how individuals seek information online, they also do not tend to address the issue head-on. What is needed, then, is an analysis of the differences in how trustworthiness is evaluated in online as opposed to offline environments.

Here I aim to make progress on such an analysis, in two ways. First, I argue that due to differences in the ways in which we acquire information online and offline, that we should consider two additional markers of trustworthiness when looking at information acquisition in online environments: the extent to which information is *endorsed*, and whether it is presented in what I call a *cooperative* as opposed to *preemptive* manner. I am not here arguing that what qualifies as a good marker of trustworthiness offline will not also be one online; rather, I argue that the nature of information-seeking online should make us reconsider which markers of trustworthiness are most salient in such environments. Second, I argue that by recognizing how it is that online users make trustworthiness evaluations that we can then make some prescriptions about how experts can more effectively communicate, namely that they should strive to be *popular* and *cooperative*. Thus, my aim here is that in addition to being a descriptive project

¹ See for example Goldman (2001), Anderson (2011), Guerrero (2016), Irzik and Kurtulmus (2019), Brennan (2020), Grundmann (forthcoming), among others.

² See e.g. Sunstein (2002), Del Vicario et al. (2016), Bramson et al. (2017), among others.

concerning the ways in which users do, in fact, make trustworthiness evaluations online, it can also be meliorative, in that it suggests ways in which recognition of these differences can lead to better communication practices online.

The paper will proceed as follows. In Sect. 1 I consider the problem of expert testimony and some of the more prominent proposed criteria for trustworthiness, specifically as they relate to trusting scientific experts. In Sect. 2 I highlight some unique characteristics of online environments that influence how users evaluate trustworthiness: first, that they contain tremendous amounts of *misinformation*, regarding both first-order scientific claims and claims about who is a trustworthy source of scientific information. The amount of misinformation online can make it more difficult to evaluate claims of expert trustworthiness while also causing a general level of skepticism in the sources one comes across. Second, that interactions with others online often occur in environments in which there are fewer *communication cues* that individuals rely on when evaluating trustworthiness in offline environments. Given the quantity of potentially misleading information and the dearth of communication cues, I argue that online information-seekers are thus more liable to treat *susceptibility to manipulation* as a key determinant of whether to accept information. As a result, online information-seekers are more likely to appeal to two additional markers of trustworthiness: in Sect. 3 I argue that the extent to which information is *endorsed* plays a key role in trustworthiness evaluations, and in Sect. 4 I argue that information-seekers will be more inclined to seek out what I call *cooperative* as opposed to *preemptive* experts. That such factors *are* appealed to in evaluations of trustworthiness does not, however, mean that they are *reliable* indicators of trustworthiness. In Sect. 5 I then argue that online information-seekers can appeal to endorsement and cooperation of a certain kind as markers of expert trustworthiness, what I call *genuine endorsement*—i.e., endorsement that is motivated by the belief that some given information is true—and *genuine cooperation*—i.e., cooperation that has the goal of producing true beliefs. This then suggests a meliorative project, in that we can consider how experts might make use of the ways individuals seek information online to communicate more effectively, namely that they should strive to be *popular* and *cooperative*. Section 6 concludes with some considerations for further areas of research.

1 Determining the trustworthiness of scientific experts

There is a lot of scientific information that is important for me to know. However, as a layperson, it can be difficult to know what information to accept, given that I am unable to evaluate it for plausibility. Say, for example, that I read a claim made by a scientist that anthropogenic climate change will result in sea levels rising 26–55 cm. in the twenty-first century with a 67% confidence interval (Kopp et al., 2014). I do not have the scientific background required to evaluate this claim, and so cannot tell on the basis of the claim itself whether I should accept it. While it is imperative that people possess a general understanding of issues pertaining to global warming, then, a problem in trying to communicate the relevant scientific information to laypeople is that they may not be able to fully make sense of it given a lack of scientific expertise (Keil, 2010).

In order to acquire scientific knowledge, one typically needs to rely on the testimony of others, preferably that of scientific experts. In doing so one needs to determine who to trust.³ There are numerous proposals concerning the kind of evidence that one can appeal to in order to determine the trustworthiness of an expert testifier. Goldman (1999, 2001), for instance, argued that a layperson can appeal to an expert's *argumentative performances; agreement from outside experts; appraisals by "meta-experts"*; evidence of their *potential biases and competing interests*; and their *track records* (93). In the years since there have been numerous refinements and additions to this list. Anderson (2011), for example, proposes a number of "principles for second-order lay assessments of scientific claims" (145), in which layperson evaluations of testifier trustworthiness depend on three kinds of evaluations: *expertise* (determined by factors including one's credentials, and whether one has published and are recognized in their field), *honesty* (determined by evidence concerning conflicts of interest, dishonesty and academic fraud, and whether one has a history of making misleading statements), and *epistemic responsibility* (determined by evidence about the ways in which one has engaged with the scientific community in general and their peers specifically). More recently, Guerrero (2016) makes a general classification of four major questions that one can ask to determine expert trustworthiness, concerning whether they possess expertise, comparative expertise, are sincere, and whether their testimony is reliably true (11).

In general, in determining whether a given criterion constitutes good evidence of testifier trustworthiness one can ask two questions: first, whether it is, in fact, reliably indicative of expert trustworthiness; and second, whether a layperson can effectively appeal to it when making an evaluation of trustworthiness. Proposed criteria have been challenged in both of these ways: for example, in terms of the first challenge, some have questioned whether indications of quality of character like transparency, openness, sincerity, and honesty are, in fact, reliable evidence for trustworthiness, given that one might be able to exhibit these qualities without actually being a reliable source of information (John, 2018). While one could debate which criteria do and do not belong on the list, I will be more concerned here with the second challenge, namely whether laypeople can effectively appeal to these criteria to determine testifier trustworthiness. There are two additional worries in this regard: whether evidence of trustworthiness is *accessible*, and whether it is *comprehensible*.

Consider the first worry, that evidence of testifier trustworthiness may not be readily accessible: while some are optimistic that such evidence is readily available to laypersons,⁴ others are worried that it may be inaccessible insofar as it is likely incomplete.

³ Establishing that the source of such information is trustworthy is particularly important if we want to account for how one can come to be justified in believing or know the content of the scientist's testimony. This is because, at least according to some theories of testimony, acquiring knowledge on the basis of testimony requires that one trust the testifier, or, minimally, have sufficient reason to believe that they know what they are talking about (see for example Baier (1986), Hardwig (1991), Jones (2002), Almassi (2012)). If this is the case, then indications of trustworthiness can not only help us make a good decision about whether to accept some testimony, but can also determine whether we can be in a position to know it.

⁴ For example, Anderson (2011) argues that all one needs is a working internet connection to find such evidence, while Irzik and Kurtulmus (2018) argue that such information is on the whole easily available, although it might require a bit of effort to acquire in some situations, namely when it comes to matters in which more in-depth research is required.

This is because such evidence can be *curated*: given that an expert is likely to highlight their accomplishments and suppress their failures, one will tend to lack a complete picture of a testifier's track record. Additionally, seniority levels will inflate the apparent trustworthiness of experts who have been working longer, as newly minted experts will not have as many accomplishments (Guerrero, 2016). The second concern is that a layperson may not be able to evaluate evidence of trustworthiness itself. For instance, while credentials like publishing in quality journals might be a good indication of trustworthiness, they are no doubt not on the radar of the majority of laypeople (Zuccala, 2009). Brennan (2020), for example, argues that laypersons will generally lack the kind of "insider" knowledge that professionals have with regards to who is well-respected, and whether the practices an expert engages in are dubious. There is thus a worry that laypersons may not only lack the knowledge required to interpret the content of scientific testimony itself, but may also lack the ability to interpret evidence about the trustworthiness of scientific experts.

Given worries about accessibility and comprehensibility, a number of proposals have been made as to how laypersons can effectively evaluate expert testifier trustworthiness (for some recent proposals, see Guerrero (2016), Irzik and Kurtulmus (2019), Brennan (2020), Grundmann (forthcoming)). There is much to be said about these debates, and I will return to some of them in Sect. 4. The focus of my argument here, however, will be how the nature of *online environments* affects the problem of identifying experts. The way in which people acquire scientific information has undergone significant changes since the advent of the internet, with much of the information we acquire about science now being acquired online.

Of course, there is no homogeneous entity known as "the internet"⁵: there are myriad ways in which people acquire information online, and there is no doubt that we can receive excellent scientific information through online sources (e.g. online libraries, journal articles, books, etc.). Nevertheless, while such sources are available to some online, they may not be available to everyone. Indeed, while laypersons once acquired much of their scientific information from vetted science reporters (Trench, 2007), they now acquire scientific information from a broader variety of sources, including blogs, social media, and personal websites (Fahy & Nisbett, 2011). As a result, online environments represent an "overlapping information and communication space" (Trench, 2009) where scientific information is being provided not only by scientists, journalists, and other kinds of science communicators, but members of the intended audience themselves, as well. That the internet brings along significant changes in the way that laypersons acquire scientific information then raises questions about the ways in which determining the trustworthiness of scientific experts might differ in this new environment. In what follows, I will be concerned with the online spaces in which there is such an overlap of communication—such as the aforementioned blogs, personal websites, and social media—and not, say, those spaces in which information is made available after having been vetted or otherwise passing certain checks for quality (e.g. online

⁵ Thanks to an anonymous reviewer for suggesting I address this.

academic journals, Wikipedia, etc.). When I refer to “online environments”, then, I will only have the former category in mind.⁶

Online environments, I argue, have unique characteristics that influence the ways in which individuals seek out and evaluate evidence of testifier trustworthiness. I focus here on two: first, many online environments are rife with misinformation, and an information-seeker may be presented not only with evidence of the trustworthiness of scientific experts, but also lots of potentially misleading evidence that can cause one to doubt an expert’s trustworthiness; and second, the nature of many online interactions are such that one lacks the kinds of *communication cues* that play a role in evaluations of trustworthiness in face-to-face interactions. As a result of these differences, I argue that laypersons may seek out different markers of trustworthiness when acquiring information online, and that as a result of these differences the problem of determining expert trustworthiness can take on a different shape in online environments.

2 Misinformation and expert trustworthiness

As we saw above, there are two concerns with identifying potential markers of trustworthiness: how easily one can acquire evidence of trustworthiness, and whether one can make sense of it once acquired. However, even if information about expert trustworthiness is available and comprehensible, there is an additional problem that is present in online environments, namely the widely recognized fact that such environments are a potential source of a tremendous amount of false or misleading information (Del Vicario et al., 2016). Frequent users of the internet are no doubt familiar with this phenomenon, and it has been well-documented that false and misleading information about scientific claims, especially ones that have become politicized (i.e. are such that one’s stance on those claims is treated as a marker of personal or social identity, viz. anthropogenic climate change, vaccine safety, appropriate Covid-19 precautions, etc.), are prevalent online. For example, some have argued that there has been an “infodemic” that accompanied the Covid-19 pandemic: with so much information coming out so quickly, and not all of it good, separating the good information from the bad becomes increasingly difficult (Eysenbach, 2020).⁷ Of course, with so much information to sift through it is perhaps more important than ever to find those testifiers who really are trustworthy. However, just as misinformation exists with regards to first-order scientific claims, it also exists with regards to who is, in fact, trustworthy.

There are many well-documented instances of individuals and groups attempting to discredit scientific information. One prominent case concerns ongoing campaigns

⁶ This distinction between spaces is admittedly still broad, and one’s mileage may vary in terms of the reliability of the information received in either space. While imperfect, this distinction will, I believe, suffice for the arguments that follow.

⁷ This is not the sole ground for skepticism when seeking information online. For instance, one might find oneself in an “echo chamber”, i.e. an environment in which certain views are mutually reinforced by other individuals, and in which dissenting views may be actively excluded, which may contribute to one’s skepticism towards certain views or evaluations of expertise (see Del Vicario et al., 2016). There is much to be said about the broader epistemological and psychological factors that might contribute to skepticism online; however, to keep the discussion as concise as possible I will not engage with these issues in detail here.

aimed at discrediting climate science, which look to “cloud the science on important issues or undercut the trustworthiness of the scientific community at large” (Slater et al., 2019). For instance, Brulle (2014) describes how the “climate change counter movement” involves a number of activities, including “political lobbying, contributions to political candidates, and a large number of communication and media efforts that aim at undermining climate science” (682). In addition to providing misinformation with regards to scientific claims, misinformation campaigns target the experts themselves: again, in the case of climate science the credentials and expertise of scientists arguing for the existence of anthropogenic climate change have been challenged (Dunlap & McCright, 2015), and contrarian scientists have been invited to numerous scientific hearings in order to present an image of scientific controversy (Koebler, 2014).

While many of these campaigns began long before the development and subsequent ubiquity of the internet, they are very much alive and well in contemporary online environments. For instance, Bloomfield and Tillery (2019) investigated the perpetuation of climate change denial on Facebook, finding that deniers employed a number of methods to attempt to make themselves appear more credible. Such methods included linking to blogs and blog communities dedicated to perpetuating climate change denialism that not only gave their claims the appearance of being well-supported, but also had the effect of populating search results with misleading information. Furthermore, these interconnected networks of sites have been used to both attempt to lend credibility to non-experts, while at the same time discrediting actual experts: as claims made by non-experts are continually reproduced and cited, said non-experts are given the appearance of expertise. Finally, these websites have also produced headlines stating that climate scientists are consistently being proven wrong, that climate science is akin to a cult, and that scientists are looking to dupe the general public, along with possessing other “nefarious motives” (28). Unfortunately, studies have also reported that people are in general not good at distinguishing trustworthy websites from those that have been created to mislead, with users often being misled by “fictitious quality seal and organization/domain names” (Wogalter & Mayhorn, 2008).

While many of the classic case studies of those who have attempted to discredit scientific experts involve conscientious efforts by groups looking to perpetuate particular interests (see e.g. Oreskes & Conway, 2011), online environments provide many more opportunities for users to challenge the trustworthiness of experts without the need for such extensive planning or organization. In general, online environments can often be what Don Fallis calls *epistemically adversarial*, in the sense that many users we interact with online are “trying to interfere with our knowledge acquisition”: they can fake credentials, and present misleading information to attempt to corroborate it; bad actors are generally difficult to identify, given that online environments can make it difficult to discern someone’s identity; and we lack the kinds of markers of trustworthiness that we use in face-to-face interactions (Fallis, 2018).

Consider again some of the proposed markers of trustworthiness mentioned in Sect. 1: while markers of expertise like track record and qualifications may often be accessible online, given the quantity of misleading information available and the overall potentially epistemically adversarial nature of online environments, it can

be difficult for one to assess whether these are, in fact, good indications of testifier trustworthiness. We have also seen that while attacks on character in attempts to discredit scientific experts is nothing new, in online environments in which there is an overlapping information and communication space of the qualified and unqualified, information about conflicts of interest, fraud, and histories of making misleading statements can be more difficult to evaluate. We thus have a problem: as we saw above, dealing with an immense amount of information regarding scientific claims requires being able to sort the wheat from the chaff, and doing so requires determining which testifiers to trust. However, in online environments there is not only an immense amount of information that pertains to first-order scientific claims, but there is potentially a significant amount of information as to who is a trustworthy testifier, as well. Being able to sort through second-order claims about testifier trustworthiness then becomes a problem in its own right. Indeed, I argue that due to this problem, internet users are likely to appeal to markers of testifier trustworthiness that are more easily identifiable in an online environment. In order to show this, I will first note one additional characteristic of online environments, namely that they exhibit a general dearth of communication cues online. We have seen mention of this in passing already: as Fallis (2018) notes, lacking communication cues can contribute to the epistemically adversarial nature of online environments. I will first present some evidence that this is the case, and then show how this can lead to users online appealing to different markers of trustworthiness.

3 Social presence and communication cues

Research in computer-mediated communication concerns how personal interactions differ between computer-mediated environments and face-to-face environments, and how the effects of online anonymity impact adherence to social norms (Marino et al., 2016). For instance, Kane et al. (2014) argue that there is reason to believe that the ways individuals interact with others online “may bear little connection to offline social relationships” (286), in that face-to-face discussions tend to be more goal-oriented, and involve more interpersonal deliberation.

A key difference between online and face-to-face communication is that there tends to be less *social presence* in online environments. Sia et al. (2002) define social presence as “the degree to which people establish warm and personal connections with each other in a communication setting”, and is characterized by three types of communication cues: verbal cues that pertain to tone, volume, and rate of speech (see also Cook & Lallijee, 1972; Daft et al., 1987; McGrath, 1984); visual cues, including facial expressions and body language; and textual cues, which pertain to the information included in written text. One reason why online environments may tend to have fewer such cues is due to the potential for anonymity: many social media websites, for example, do not require that one provide identifying information about oneself, while other message board websites can allow complete anonymity. As a result, social presence is lower in online environments since “communication cues that typically yield higher social presence are those that convey immediacy” (where “immediacy” is defined as “the psychological distance between people who are communicating”),

and the kinds of communication cues that one receives in face-to-face communication—e.g. verbal and visual cues—are ones that convey immediacy, whereas the cues available in anonymous computer-mediated communication—e.g. textual cues—do not (Sia et al., 2002, p. 74).

We have seen already that the amount of information available online with regards to expert trustworthiness can make it more difficult for internet users to determine whom to listen to. Lower social presence and fewer communication cues make this problem even more difficult, as one cannot make trustworthiness evaluations in the same way as one would in face-to-face environments, and thus have fewer available tools to corroborate trustworthiness judgments. As a result of these challenges, individuals in online environments may look to different types of markers of trustworthiness than in face-to-face interactions.

One of the most prominent bases upon which users evaluate trustworthiness in online environments comes in the form of *endorsement*: this may come in the form of explicit endorsement markers—e.g. “likes”, “hearts”, or “upvotes” (Willemsen et al., 2012)—or more indirect cues, such as the number of connections that one has within a given network—e.g. the number of “friends” one has on a social media network (Lim & Van Der Heide, 2015). In general, information that has been highly endorsed is more readily accepted, and members who are highly endorsed by other members tend to be considered more trustworthy (Willemsen et al., 2012). Reliance on endorsement as a marker of trustworthiness is the result of online users having only limited social cues upon which to evaluate others. Metzger et al. (2010), for instance, found that “participants developed strategies to assess a source’s credibility as best they could” (421), given both the overall dearth of information about other members, and the concern that what information might be available—say, in the form of information presented on user-generated profiles—could be manipulated and curated by the members themselves.

In order to evaluate the credibility of other members, individuals will tend to employ what Walther (2011) calls “warranting theory”: in determining the credibility of others, individuals will seek out credentials that are the least susceptible to manipulation. Flanagin and Metzger (2013) also appeal to “signalling theory”, which argues that,

[C]ertain signals available online about information sources—particularly those signals that are difficult to fake, are supported by the rule of law or social convention, or are costly to obtain or to mimic—are most reliable for assessing information quality and source expertise and, therefore, can be trusted. (1627)

In online environments, one such credential that is most readily available comes in the form of *aggregates*, both in the form of aggregate endorsement ratings, and aggregates of information: this is because as information aggregates it becomes more and more difficult to manipulate, and any potential subjective biases in the ratings of individuals and information will have less of an effect on endorsement overall (Flanagin & Metzger, 2013). What research in online communication suggests, then, is that without the kinds of social cues one typically relies on to help assess credibility in face-to-face interactions, one’s interactions with others online will tend to be mediated in different ways.

As we have seen, however, there are lots of ways that one can determine the trustworthiness of expert sources: an expert's credentials, publications, records of interactions with the scientific community, etc., are potentially a quick Google search away. Indeed, in the study cited above, Flanagin and Metzger (2013) argue that aggregated information serves as a marker of reliability and trustworthiness in user-generated content specifically (i.e. blogs, social media posts, etc.), and that information provided by experts "should be largely immune to enhanced credibility by virtue of increased volume" since the validity of expert-provided information "is assured by means other than the aggregation of multiple opinions", i.e. by those aforementioned indicators of expertise. However, we have also seen that users often acquire scientific information precisely from places in which there is a significant amount of user-generated content (i.e. the overlapping communication spaces under consideration here) and that given the amount of misinformation available online, that these indicators may also be interpreted by online users as being subject to manipulation. Indeed, Flanagin and Metzger note that user-generated content is generally favored under conditions of "high information volume" and, as we have seen above, in times of important scientific developments there tends to be just such a high volume of information (e.g. the Covid-19 "infodemic"). In an online environment one might then look to indicators other than the markers outlined in Sect. 1 in order to make evaluations of testifier trustworthiness.

Of course, one may not *typically* interact with experts face-to-face: reading books, articles, or parts of the internet that have been thoroughly vetted allow one to receive expert testimony, but does not involve the communication cues one can acquire offline.⁸ In this regard, then, there is not much of a difference in the way we evaluate the trustworthiness of experts online and offline. Recall, however, that the online environments discussed here are those in which there are overlapping communication spaces, where experts and non-experts may contribute information, and in which it can be difficult to tell who's who. There are, then, important differences between the ways that people will evaluate testifier trustworthiness when reading an online journal article as opposed to a post shared on social media. Given the nature of these online environments, and given that the nature of online communication in general makes different markers of trustworthiness salient, then such markers will tend to be those that are appealed to more often when evaluating trustworthiness online.

One might also be concerned that while aggregate endorsement may be *treated* as a marker of trustworthiness, that highly endorsed information online will often not, in fact, be good information. For instance, algorithms that drive social media like Twitter and Facebook make popular posts more visible to users, which are in turn likely to become more highly endorsed due to increased exposure. However, these algorithms tend to favor information that will tend to elicit the most engagement, regardless of whether it is true (Yardi & Boyd, 2010). There have been documented negative epistemic consequences: for instance, it has been shown that in many social media networks that misinformation spreads faster than mechanisms that have been put in place to help ameliorate it, such as fact-checking (Vosoughi et al., 2018). Reliance on endorsement as a marker of trustworthiness may also have negative effects in terms of

⁸ Thanks to an anonymous referee for suggesting I address this point.

the creation and sustaining of echo chambers: like-minded members of a group will tend to endorse content that conveys messages of a certain type, which in turn makes that content more visible to others and more likely to be endorsed (see Boyd, 2018).

Of course, no potential marker of trustworthiness is, in isolation, always necessarily indicative of trustworthiness, and this is no different in the case of aggregate endorsement.⁹ The purpose of this section has been to argue that, given the challenges of finding markers of trustworthiness that are not easily susceptible to manipulation in online environments, that aggregate endorsement is one such easily identifiable and interpreted marker. Indeed, the detrimental epistemic effects of endorsement mechanisms in social media environments provide evidence that endorsement is, in fact, being treated as a marker of trustworthiness, even though there are many cases in which it probably should not be. In Sect. 5 I will consider some additional implications for what reliance on endorsement online has for how experts should communicate online, as well as what users can do in order to make sure that endorsement is, in fact, a reliable marker of trustworthiness in a given context. So far, however, we have seen that the nature of online environments is such that information-seekers put a special emphasis on *resistance to manipulation* as a marker of credibility. I argue next that this emphasis indicates a second way in which users are likely to identify trustworthy experts online, namely that in online environments one is more likely to seek out experts that provide information *cooperatively* as opposed to *preemptively*.

4 Preemption and cooperation

The question of how laypeople can effectively rely on expert testimony is a species of a broader question about how people rely on testimony in general. One way in which testimony is distinguished from other epistemic sources (like perception, memory, and inference) is that it is necessarily *deferential*: when I come to believe something on the basis of testimony, I believe it because someone else told me so, and I take their word for it.¹⁰ Testimony from an expert to a layperson is a special class of testimony in that it typically involves deference of an extreme sort: as a layperson, not only do I take an expert's word for it, but I have little choice but to take their word for it given that I am likely unable to independently determine the veracity of the content. The problem of whom to listen to and how to interpret whether someone is a good source of testimony, then, is a problem of deciding whom to defer to.

There is debate, however, about what deferring to experts *ought* to involve. Zagzebski (2012), for instance, argues that when it comes to relying on experts a layperson

⁹ This, I think, is a mistake in some of the literature surrounding expert trustworthiness, namely that proposed markers of trustworthiness are evaluated in isolation and shown how they can be manipulated and exploited by bad actors. While I do not have space to defend this position here, my view is that trustworthiness evaluations should be considered as more of a holistic project, the result of weighing multiple factors that may be better or worse markers of trustworthiness in different contexts and when considered alongside each other.

¹⁰ Questions about how it is that one can acquire justification and knowledge via testimony are topics that are too involved to deal with here. I am not interested here in defending any particular theory of testimony; for my purposes I merely require the uncontroversial assumption that in relying on testimony we rely on what someone else has said.

should not only accept their testimony, but do so *preemptively*, i.e., accept it without weighing it against any beliefs or reasons that they might already possess, and accepting the expert's reasons as their own.¹¹ For example, say that I believe that properly worn facemasks are not an effective means of preventing the transmission of the Covid-19 virus. An expert epidemiologist then provides testimony that facemasks are, in fact, effective in this way, and provides reasons in the form of results of empirical studies. Although I possess beliefs and reasons that support my own view, we might think that, given my status as a layperson, that I should replace those beliefs and reasons with those provided by the expert. One of Zagzebski's primary motivations in defending this view concerns the track record of an expert versus that of a layperson: since I do not possess any expertise myself, I will do better in the long run if I accept an expert's reasons in place of my own, given that any reasons I do possess are more likely to lead me astray. However, others have argued that when accepting testimony from an expert it is not *always* the case that one ought to accept it preemptively: instead, there are cases in which one ought to accept testimony in such a way that takes into account one's preexisting beliefs and reasons, especially when said reasons cohere with the new information received (Wright, 2016). In such cases, then, it can be rational to accept an expert's testimony only after weighing their reasons against those one already possesses.

My aim here is not to adjudicate this debate. Instead of saying that one always *ought* to defer preemptively, I argue instead that there simply *are* circumstances in which we treat some experts in a preemptive manner—that is, we *do* accept their reasons in place of our own—and others in which we do not. For example, consider a case in which I travel to a city I have never been to, in order to meet up with some colleagues for dinner. Not knowing anything about the city or the quality of any nearby restaurants, I rely on the local gourmand to decide where we should go. In this case I defer preemptively insofar as I have no reasons of my own to replace, so I am happy to accept those of the testifier. I might also look to defer preemptively when I possess some relevant reasons, but take them to be too weak to warrant action: for instance, say that I have some reason to think that while the restaurant around the corner is decent, the one much further away is significantly better, although I don't know enough about either to decide whether the extra time it takes to get to the better one is worth it. Here I am also happy to defer preemptively: if you as the local expert possess knowledge that I lack, then I will gladly accept what you have to say without weighing my reasons against yours. As a layperson with regards to scientific matters there are also many instances in which I look to defer preemptively. For example, when it comes to information about new discoveries in astronomy, I typically accept whatever it is that NASA tells me, and adopt their reasons as my own.

There are, then, cases in which experts about some subject matters play a certain type of role in the way one acquires knowledge, namely that one takes their word for it in such a way that one accepts what they say while replacing one's reasons with those of the expert. Call the experts who play such a role for a given agent a

¹¹ Here, I will take there to be two kinds of preemptive acceptance: cases in which one already possesses relevant reasons, which they replace with those of the expert; and cases in which one does not possess any relevant reasons, and merely adopts the expert's reasons as one's own. Thanks to an anonymous referee for suggesting I clarify this point.

preemptive expert. The same individual may occupy this role for some people and not others: for example, while ardent fans of popular scientist Neil DeGrasse Tyson may preemptively accept whatever he has to say about philosophy because of his expertise in astrophysics, philosophers will likely not accept much of his testimony on these matters, or, if they do, will only do so after consideration of his reasons against their own. To say that we treat some experts preemptively is again not to make a claim about whether we *ought* to do so, but instead that there are simply some experts that play this role in our processes of inquiry.

There are other cases, however, in which one will seek out testimony, but is not willing to preempt one's reasons. For instance, say that I am seeking advice concerning a matter about which I am indecisive. I may very well accept what you tell me, and accept what you tell me because you tell it to me, but it may not be something that I will accept without considering how my own reasons weigh against yours. I might also seek out testimony to reassure myself when I have doubts, or to double-check something, in which case I will accept your testimony, but not in a way that will preempt reasons that I may possess already. In such cases, then, we will tend to seek testifiers who are able to not only provide true information, but to provide it in such a way that addresses, in some way or another, the beliefs and reasons we possess. Call an expert who plays this role in our processes of inquiry a *cooperative expert*. Again, the same expert can play different roles for different people, and some may be willing to cooperatively defer to those whom others would preemptively defer.¹²

Given the differences in the roles that preemptive and cooperative experts play when seeking out and acquiring knowledge, one may evaluate the trustworthiness of experts who play different roles according to different standards. For instance, when seeking information from someone who occupies the role of a preemptive expert, one will be concerned solely with being provided true information, and may have less interest in explanations of the reasons the testifier possesses or their ability to address any potential questions or concerns. In other words, when seeking information from preemptive experts one simply wants to be told the right answer, and thus a good preemptive expert is one who can provide those answers. As such, when evaluating the trustworthiness of those who occupy a role of a preemptive expert, track record, qualifications, engagement with peers, and consensus with the greater scientific community may be the most reliable markers: these criteria indicate that an expert knows what they are talking about, but do not generally indicate that they will present information in such a way that takes the reasons and beliefs that a layperson may possess into account.

When seeking information from someone who occupies the role of a cooperative expert, on the other hand, one will still be concerned with receiving true information, but in such a way that it addresses one's prior beliefs and reasons. As such, when evaluating whether one is a trustworthy cooperative expert one may look to factors

¹² There has been a good amount of discussion recently concerning how "expert" should be defined, and how experts can be taxonomized. For example, some have defined experts in terms of their abilities to provide "cognitive help" to novices (Quast, 2018), while others focus on an expert's ability to make progress in a given domain (Croce, 2019). Grundmann (forthcoming) rejects these views, arguing that scientific experts in particular are those who are identified by the scientific institution. Here I do not take a stance on the necessary and sufficient conditions of expertise. Instead, in this section I am indicating the ways in which laypersons can rely on experts in different ways, not demarcating distinct classes of experts.

such as whether they make it easier to weigh their testimony against a recipient's prior beliefs, potentially by addressing them and why they should be supplemented or outweighed. Factors like argumentative and dialectical skills, honesty, and engagement with the public, then, are potentially more reliable indicators of the trustworthiness of a cooperative expert.

If there are some testifiers for some subject matters that we treat as preemptive experts, and others that we treat as cooperative experts, then lists of markers of expert trustworthiness risk conflating criteria for evaluating experts that play different roles in one's inquiries. For instance, we saw above that some have argued that factors like quality of personal character do not belong on a list of expert trustworthiness, as they do not necessarily correlate with the reliability of one's testimony (John, 2018). However, that this may be the case speaks most prominently against employing such evidence when identifying preemptive experts; when it comes to cooperative experts, on the other hand, such evidence may indeed be a good indication of a testifier's propensity to provide information in such a way as to consider the beliefs and reasons possessed by the recipient. Furthermore, while agreement with the scientific consensus is typically taken to be a reliable marker of expert trustworthiness, evidence that one's statements are in line with the scientific consensus is sometimes interpreted as a reason to *distrust* a given expert. For instance, Levy (2019) argues that consensus reports can be ineffective as a way to convey scientific information as they can be interpreted by laypeople as a marker of a failure of benevolence on behalf of a testifier:

While we are apt to accept testimony—to defer to others—we reject testimony from sources that signal unreliability by evincing cues of incompetence or lack of benevolence. When science becomes politicized, expression of the scientific consensus may itself come to serve as a signal of lack of benevolence to those on one side of the issue, leading to rejection of the testimony. On all sides, filtering mechanisms may be working as designed, but for reasons beyond the purview of the individuals involved, warrant may accrue to one side alone. (314)

Here, then, the problem with mere reporting of consensus information is not that it is inaccurate or misleading, but rather that it is perhaps best indicative of the quality of someone who occupies the role of a preemptive expert: for example, the statement that 97% of scientists agree that climate change is a real phenomenon presents the information as something to be accepted uncritically. However, if a recipient already possesses beliefs and reasons in support of a view, the issue of climate change may be something about which one seeks information from a cooperative expert instead: in this case, a presentation of mere consensus can represent a failure to engage with them, and can lead to distrust.

Most importantly for my purposes here, however, is how the distinction between the different roles that experts can play has consequences for information-seeking online. The idea is as follows: given the amount of misinformation available online, and the epistemically adversarial nature of many online environments, combined with the lack of communication cues that can serve as bases for trustworthiness evaluations in face-to-face environments, individuals online may generally be more inclined to seek information from cooperative as opposed to preemptive experts. This is because in seeking cooperative experts one looks to receive information in such a way that

can address doubts that the recipient might possess, doubts that one is more likely to possess in epistemically adversarial environments. Furthermore, one may find it easier to evaluate markers of trustworthiness that are more indicative of cooperative experts in online environments: as we saw above, information-seekers online may be particularly concerned with the susceptibility of manipulation when evaluating information, and in this regard factors more closely aligned with cooperative experts—such as an expert’s argumentative and dialectical skills, and the extent to which an expert engages with the public—may be seen as factors that are less easily manipulated.

Additional support for this argument can be found in some recent work from science communications that concerns how to effectively communicate scientific information to those who may be reluctant to accept it because of social and political reasons, e.g. communicating information about climate change to those who identify socially and politically as a type of conservative who is skeptical of such views. One approach that has seen success is a *value-based* approach: information is presented in such a way that takes into account the values of one’s intended recipient, e.g. climate change information is presented to a conservative audience in a way that emphasizes the potentially catastrophic economic consequences (as opposed to the environmental ones) (see Zhou, 2016; Dixon et al., 2017). Information presented in this way is more likely to be accepted than when it is simply presented as scientific consensus. This is a case, then, in which presenting information in a cooperative, as opposed to deferential manner, can be a more effective way of getting one’s audience to accept it.

The idea, then, is just as those who are skeptical of certain types of information are going to accept information more readily from cooperative experts, given that online information-seeking environments will tend to be ones in which information-seekers have more reason to be skeptical due to a lack of communication cues and the presence of misinformation, that they will overall be more likely to seek out cooperative experts. This is not to say that individuals *always* seek such experts; furthermore, it is not to say that individuals are always successful in identifying trustworthy experts, nor that they always acquire true beliefs as a result. Rather, my argument thus far has been that the nature of seeking information online is such that those looking to identify trustworthy experts in online environments are more likely to consider different markers of trustworthiness as more relevant, and are more likely to seek out experts that bear the markers of trustworthiness of a certain type. These differences then suggest some ways in which the social epistemologist can amend lists of expert trustworthiness, and how experts can more effectively communicate online.

5 How to be a trustworthy expert online

I have suggested that the traditional criteria of expert testifier trustworthiness outlined in Sect. 1 do not specifically take into consideration how characteristics of information seeking in online environments impact the ways in which individuals make trustworthiness evaluations. I have identified two additional factors that impact trustworthiness assessments online: endorsement and cooperation. However, while my argument thus far has been that users online *will* appeal to such factors to identify trustworthy sources,

it is not clear the extent to which such factors *are*, in fact, reliable indicators of trustworthiness. We can, then, refine these factors in order to identify two markers of trustworthiness that are particularly salient in online environments, what I will call *genuine endorsement* and *genuine cooperation*.

Consider first questions of endorsement. We can ask: under what conditions and on the basis of what factors is endorsement a good indicator of testifier trustworthiness? There is much to be said here with regards to the bases upon which endorsement is made, how it can be inflated by the use of algorithms, etc. I will not engage with all of these debates here. Instead, I will make the following general proposal as an additional marker of trustworthiness in online environments:

Genuine Endorsement: The extent to which information presented by an expert is genuinely endorsed – i.e. endorsed by individuals because they think it is true – the more trustworthy the expert presenting it.

Genuine endorsement parallels one of the traditional markers of trustworthiness, namely whether one is providing information that agrees with the relevant *consensus*. The difference between the markers comes down to who is doing the endorsing, and how: while scientists may reach consensus after extended periods of research and debate, endorsement online can often be swift, and reached as a result of inputs from many non-experts in the form of various endorsement mechanisms. That being said, a failure for some information to be genuinely endorsed online should cause concern. For example, if a well-connected scientist makes a social media post in which they make a scientific claim which is not genuinely endorsed—i.e. there is no endorsement or engagement from other members of the scientific community—then one has reason to approach the claim or source with skepticism.¹³

We saw above in the discussion of the traditional markers of expertise that we can ask whether they are usable by laypersons (whether they are accessible and whether they are comprehensible) and whether they are, in fact, good markers of trustworthiness. How does a criterion like *genuine endorsement* fare in these regards? It seems that endorsement of information is easily comprehensible, insofar as it is easy to identify when information is highly endorsed or not. Indeed, some research suggests that aggregate endorsement is sometimes taken to be a marker of expertise, with individuals who are highly endorsed by others being deemed more credible than those who are self-proclaimed experts (Willemsen et al., 2012). However, it may also be unclear as to whether endorsement is *genuine*. Again, we can consider different environments in which endorsement may be more likely to be genuine than others: for example, politically-motivated comments on opinion pieces may be highly endorsed, but may tend to be such that they are not endorsed on the basis of an evaluation of truth, but instead merely cohesion with one's political values. Other situations may very well be different: for example, if one is acquiring information from a discussion board that reports the results of new developments in science, then one may have much better reason to think that endorsement represents the sincere belief in the truth of the relevant content.

¹³ Thanks to an anonymous referee for suggesting this example.

One might worry, however, that endorsement can be relied upon *too* much: detrimental epistemic effects may be amplified if one relies *solely* on endorsement when determining which information to accept, and which sources one deems trustworthy online. We have also seen that reliance on endorsement is perhaps particularly troublesome when one relies on it uncritically in environments that are likely to result in increased propagation of misleading information, such as echo chambers. Nevertheless, given that endorsement is typically resistant to manipulation, and given that such concerns are significant in online environments, genuine endorsement should be seen as a good but fallible marker of trustworthiness when combined with relevant additional markers.¹⁴

If endorsement is a salient marker of trustworthiness online, then we can make the following prescription for those wanting to be trusted more in online environments:

Be popular: Experts in online environments should, whenever possible, strive to present their information in such a way that it can be endorsed, and to establish connections to other experts and users.

What this advice amounts to on a practical level will depend on the specific nature of the environment in which one is presenting information; in general, however, one might look to disseminate information via social media or other websites that employ endorsement mechanisms, as well as to establish connections with other users on such websites.

The second lesson I drew from the nature of information seeking in online environments is that users online are, in general, more likely to seek out cooperative as opposed to preemptive experts. Again, we can identify a set of questions concerning the relationship between cooperation and trustworthiness online. Here again I will limit myself to a general proposal for a second criterion of trustworthiness that is salient in online environments:

Genuine Cooperation: The greater the extent to which the information provided by an expert takes into consideration the relevant beliefs and reasons likely possessed by an audience with the aim of having them believe truths, the more trustworthy the expert providing it.

Here I propose that cooperation is *genuine* in the sense that an expert does not merely attempt to convince a recipient by appealing to their beliefs and reasons, but that they do so in a way that is motivated by intending for the recipient to believe the truth. Thus, one who is willing to engage in conversation but has ulterior motives, say, will not be genuinely cooperative. As was the case for genuine endorsement, a lack of genuine cooperation in online environments can give one reason to be skeptical about a given claim or source, especially when it is made in environments in which cooperation is expected or encouraged. Again, we can consider a scientific expert who posts information on social media or a message board site, but who dismisses or ignores the questions and concerns of other users attempting to engage on the platform. These failures of genuine cooperation online will then give users reason to question the trustworthiness of the expert.

¹⁴ Thanks to an anonymous referee for urging me to make this clearer.

Again, as was the case with *genuine endorsement*, *genuine cooperation* is not an infallible marker of trustworthiness. Similarly, while it may be easy to identify when someone is being cooperative, it can be more difficult to determine when they are being *genuinely* cooperative. Again, such markers must be evaluated against the background of a relevant context: social media posts in which someone is trying to convince you of the benefits of a new fad diet with a link to an online retailer may be presenting information in a cooperative manner, but is likely not doing so genuinely; fora in which laypersons can engage in Q&A with experts, on the other hand, are likely to be ones in which information is being presented in a genuinely cooperative manner. That being said, we can then generate an additional prescription for those looking to be more trustworthy in online environments:

Be cooperative: Experts in online environments should, whenever possible, strive to present their information in such a way that addresses the potential reasons and beliefs possessed by information-seekers.

Some proposals concerning the acceptance of scientific testimony provide similar advice. For instance, Leefman and Lesle (2018) argue that one way to facilitate layperson acceptance of scientific testimony is to establish more extensive relationships between layperson and scientist so that the latter can help dispel doubts of the former; and Irzik and Kurtulmus (2019) recommend that scientists ought to seek to establish “enhanced epistemic trust”, wherein layperson recipients trust that scientists have taken the way that a layperson might assess risk into account when presenting scientific information. Again, while neither of these proposals specifically address the unique characteristics of seeking information online, they are particularly apt when considering such environments.

As was the case for aggregate endorsement, one might have concerns about the extent to which cooperation should be considered a reliable marker of testifier trustworthiness. For instance, Leefman and Lesle’s (2018) concern with their own approach is that layperson doubts might be dispelled by “epistemically irrelevant” arguments; similarly, one might worry that while establishing a relationship of enhanced epistemic trust may make it more likely that individuals accept scientific testimony, it remains the case that such trust could be misplaced or manipulated. However, we again need to remember that markers of trustworthiness cannot be evaluated in isolation: it is certainly always possible for bad actors to be more cooperative in order to attempt to spread false or misleading information more effectively, but this does not mean that cooperation should not be considered to generally be a good marker of trustworthiness online.¹⁵

6 Conclusion

Debates about the problem of expert testimony have typically involved attempts to determine which markers of expert trustworthiness are those that ought to be appealed

¹⁵ It is also important to note that there is no marker of trustworthiness that one cannot attempt to fake. Genuine cooperation and genuine endorsement are therefore not exceptions in this regard, but also should not be held to a higher standard of scrutiny as a result.

to by laypersons, often resulting in lists of criteria of various lengths. Here I have argued that such discussions do not tend to take into account characteristics of what have arguably become the most prominent environments in which individuals acquire their scientific information: online environments. The result of this discussion has been that certain factors—such as endorsement and cooperation—may play much more of a role in determining trustworthiness online, and that as a result those looking to communicate scientific information to laypersons should strive to be more *popular* and *cooperative*.

I have argued that social epistemologists should then at least *supplement* existing criteria for expert trustworthiness with those which may be more prominent online. However, given the differences in the ways that we acquire information in different environments, one might instead posit that there ought not be one universal list of criteria, but instead different lists relativized to different environments. Similarly, given the distinction between preemptive and cooperative experts that I have introduced here, one might posit that there is no list of criteria of expert trustworthiness that applies to all of the ways in which we rely on experts; indeed, we have seen some indication that this might be the case, in that presenting scientific information as that which is supported by the scientific consensus can be interpreted as either reason to trust or distrust the testifier, depending on the circumstances. Again, here I have made the more modest call to supplementing lists of criteria for expert trustworthiness, but it is worth considering whether differences in the way information is acquired online call for a more radical reinterpretation of what it means to be a trustworthy expert.

Funding Funding was provided by Danmarks Frie Forskningsfond (Grant No. 8018-00053B).

References

- Almassi, B. (2012). Climate change, epistemic trust, and expert trustworthiness. *Ethics and the Environment*, 17(2), 29–49.
- Anderson, E. (2011). Democracy, public policy, and lay assessments of scientific testimony. *Episteme*, 8(2), 144–164.
- Baier, A. (1986). Trust and antitrust. *Ethics*, 96(2), 231–260.
- Bloomfield, E. F., & Tillery, D. (2019). The circulation of climate change denial online: Rhetorical and networking strategies on Facebook. *Environmental Communication*, 13(1), 23–34.
- Boyd, K. (2018). Epistemically pernicious groups and the groupstrapping problem. *Social Epistemology*, 33(1), 61–73.
- Bramson, A., Grim, P., Singer, D. J., Berger, W. J., Sack, G., Fisher, S., Flocken, C., & Holman, B. (2017). Understanding polarization: Meanings, measures, and model evaluation. *Philosophy of Science*, 84, 115–159.
- Brennan, J. (2020). Can novices trust themselves to choose trustworthy experts? Reasons for (reserved) optimism. *Social Epistemology*, 34(3), 227–240.
- Brulle, R. J. (2014). Institutionalizing delay: Foundation funding and the creation of US climate change counter-movement organizations. *Climatic Change*, 122(4), 681–694.
- Cinelli, M., Quattrocioni, W., Galeazzi, A., Valensise, C.M., Brugnoli, E., Schmidt, A.L., Zola, P., Zollo, F., & Scala, A. (2020). The covid-19 social media infodemic. arXiv preprint arXiv:2003.05004.
- Cook, M., & Lallijee, M. (1972). Verbal substitutes for visual signals in interaction. *Semiotica*, 6, 212–221.
- Croce, M. (2019). On what it takes to be an expert. *The Philosophical Quarterly*, 69, 1–21.

- Daft, R. L., Lengel, R. H., & Trevino, L. K. (1987). Message equivocality, media selection, and manager performance: Implications for information systems. *MIS Quarterly*, *11*(3), 355–366.
- Del Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., Stanley, H. E., & Quattrociocchi, W. (2016). The spreading of misinformation online. *Proceedings of the National Academy of Sciences*, *113*(3), 554–559.
- Dixon, G., Hmielowski, J., & Ma, Y. (2017). Improving climate change acceptance among US conservatives through value-based message targeting. *Science Communication*, *39*(4), 520–534.
- Dunlap, R. E., & McCright, A. M. (2015). Challenging climate change. In R. E. Dunlap & R. J. Brulle (Eds.), *Climate change and society: Sociological perspectives* (pp. 300–332). Oxford University Press.
- Eysenbach, G. (2020). How to fight an infodemic: The four pillars of infodemic management. *Journal of Medical Internet Research*, *22*(6), e21820.
- Fahy, D., & Nisbet, M. C. (2011). The science journalist online: Shifting roles and emerging practices. *Journalism*, *12*(7), 778–793.
- Fallis, D. (2018). Adversarial epistemology on the internet. In D. Coady & J. Chase (Eds.), *The Routledge handbook of applied epistemology*. Routledge.
- Flanagin, A. J., & Metzger, M. J. (2013). Trusting expert- versus user-generated ratings online: The role of information volume, valence, and consumer characteristics. *Computers in Human Behavior*, *29*, 1626–1634.
- Goldman, A. (1999). *Knowledge in a social world*. Oxford University Press.
- Goldman, A. (2001). Experts: Which ones should you trust? *Philosophy and Phenomenological Research*, *63*(1), 85–110.
- Grundmann, T. (Forthcoming) Experts: What are they and how can laypeople identify them? In J. Lackey & A. McGlynn (Eds.), *Oxford handbook of social epistemology*. Oxford University Press.
- Guerrero, A. (2016). Living with ignorance in a world of experts. In R. Peels (Ed.), *Perspectives on ignorance from moral and social philosophy* (pp. 1–21).
- Hardwig, J. (1991). The role of trust in knowledge. *Journal of Philosophy*, *88*(12), 693–708.
- Irzik, G., & Kurtulmus, F. (2018). What is epistemic public trust in science? *British Journal for the Philosophy of Science*, *70*, 1145.
- Irzik, G., & Kurtulmus, F. (2019). What is epistemic public trust in science? *The British Journal for the Philosophy of Science*, *70*(4), 1145–1166.
- John, S. (2018). Epistemic trust and the ethics of science communication: Against transparency, openness, sincerity and honesty. *Social Epistemology*, *32*(2), 75–87.
- Jones, W. (2002). Dissident vs. loyalist: Which scientists should we trust? *Journal of Value Inquiry*, *36*(4), 511–520.
- Kane, G. C., Alavi, M., Labianca, G., & Borgatti, S. P. (2014). What's different about social media networks? A framework and research agenda. *MIS Quarterly*, *38*(1), 274–304.
- Keil, F. (2010). The feasibility of folk science. *Cognitive Science*, *34*(5), 826–862.
- Koebler, J. (2014). The house science committee spent today in a climate change denial echo chamber. *Motherboard* <https://www.vice.com/en/article/9akkqv/the-house-science-committee-spent-all-day-proudly-denying-climate-change>
- Kopp, R. E., Horton, R. M., Little, C. M., Mitrovica, J. X., Oppenheimer, M., Rasmussen, D. J., Strauss, B. H., & Tebaldi, C. (2014). Probabilistic 21st and 22nd century sea-level projections at a global network of tide-gauge sites. *Earth's Future*, *2*(8), 383–406.
- Leeffmann, J., & Lesle, S. (2018). Knowledge from scientific expert testimony without epistemic trust. *Synthese*, *197*, 3611–3641.
- Levy, N. (2019). Due deference to denialism: Explaining ordinary people's rejection of established scientific findings. *Synthese*, *196*, 313–327.
- Marino, C., Vieno, A., Pastore, M., Albery, I. P., Frings, D., & Spada, M. M. (2016). Modeling the contribution of personality, social identity and social norms to problematic Facebook use in adolescents. *Addictive Behaviors*, *63*, 51–56.
- McGrath, J. E. (1984). *Groups: Interaction and performance*. Prentice Hall.
- Metzger, M. J., Flanagin, A. J., & Medders, R. B. (2010). Social and heuristic approaches to credibility evaluation online. *Journal of Communication*, *60*, 413–439.
- Oreskes, N., & Conway, E. M. (2011). *Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming*. Bloomsbury Publishing.
- Quast, C. (2018). Expertise: A practical explication. *Topoi*, *37*, 11–27.

- Sia, C. L., Tan, B. C. Y., & Wei, K. K. (2002). Group polarization and computer-mediated communication: EFFECTS of communication cues, social presence, and anonymity. *Information Systems Research*, 13(1), 70–90.
- Slater, M. H., Huxster, J. K., & Bresticker, J. E. (2019). Understanding and trusting science. *Journal for General Philosophy of Science*, 50(2), 247–261.
- Sunstein, C. (2002). The law of group polarization. *The Journal of Political Philosophy*, 10(2), 175–195.
- Takahashi, B., & Tandoc, E. C., Jr. (2016). Media sources, credibility, and perceptions of science: Learning about how people learn about science. *Public Understanding of Science*, 25(6), 674–690.
- Trench, B. (2007). How the internet changed science journalism. In M. Bauer & M. Bucchi (Eds.), *Journalism, science and society: Science communication: between news and public relations* (pp. 133–141). Routledge.
- Trench, B. (2009). Science reporting in the electronic embrace of the internet. In R. Holliman, E. Whitelegg, E. Scanlon, S. Smidt, & J. Thomas (Eds.), *Investigating science communication in the information age: Implications for public engagement and popular media* (pp. 166–179). Oxford University Press.
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151.
- Walther, J. B. (2011). Theories of computer-mediated communication and interpersonal relations. *The Handbook of Interpersonal Communication*, 4, 443–479.
- Willemsen, L. M., Neijens, P. C., & Bronner, F. (2012). The ironic effect of source identification on the perceived credibility of online product reviewers. *Journal of Computer-Mediated Communication*, 18, 16–31.
- Wogalter, M. S., & Mayhorn, C. B. (2008). Trusting the internet: Cues affecting perceived credibility. *International Journal of Technology and Human Interaction (IJTHI)*, 4(1), 75–93.
- Wright, S. (2016). Epistemic authority, epistemic preemption, and the intellectual virtues. *Episteme*, 13(4), 555–570.
- Yardi, S., & Boyd, D. (2010). Dynamic debates: An analysis of group polarization over time on Twitter. *Bulletin of Science, Technology & Society*, 30(5), 316–327.
- Zagzebski, L. (2012). *Epistemic authority: A theory of trust, authority, and autonomy in belief*. Oxford University Press.
- Zhou, J. (2016). Boomerangs versus javelins: How polarization constrains communication on climate change. *Environmental Politics*, 25, 788–811.
- Zuccala, A. (2009). The lay person and Open Access. *Annual Review of Information Science and Technology*, 43(1), 1.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.