



It Takes a Village to Trust Science: Towards a (Thoroughly) Social Approach to Public Trust in Science

Gabriele Contessa¹ 

Received: 18 May 2021 / Accepted: 17 October 2021

© The Author(s), under exclusive licence to Springer Nature B.V. 2022

Abstract

In this paper, I distinguish three general approaches to public trust in science, which I call the individual approach, the semi-social approach, and the social approach, and critically examine their proposed solutions to what I call the problem of harmful distrust. I argue that, despite their differences, the individual and the semi-social approaches see the solution to the problem of harmful distrust as consisting primarily in trying to persuade individual citizens to trust science and that both approaches face two general problems, which I call the problem of overidealizing science and the problem of overburdening citizens. I then argue that in order to avoid these problems we need to embrace a (thoroughly) social approach to public trust in science, which emphasizes the social dimensions of the reception, transmission, and uptake of scientific knowledge in society and the ways in which social forces influence both positively and negatively the trustworthiness of science.

✉ Gabriele Contessa
gabriele_contessa@carleton.ca

¹ Philosophy Department, Carleton University, Patterson 3A46, 1125 Colonel By Drive, Ottawa, ON, Canada

1 Introduction

Among academics and commentators, there is a growing sense that public trust in science and scientists is at a low ebb.¹ While reports of the death of expertise (Nichols, 2017) seem to be greatly exaggerated,² it is undeniable that, in the United States (as well as in other countries), there are sizeable pockets of distrust in the scientific consensus on issues ranging from the contribution of human activity to climate change³ to the safety of vaccines.⁴ Given the social and environmental ramifications of some of these issues, the question of public trust in science has gained a new urgency. If a sizeable proportion of our fellow citizens do not believe in the reality of anthropogenic climate change or take its possible consequences seriously, how can we expect them to support the sort of social, political, and economic reforms required to prevent the worst-case scenarios? If people do not trust vaccines to be safe, how can we expect them to have their children vaccinated?

To the extent to which distrust in science is both unwarranted and socially harmful (as it seems to be in the two cases just mentioned),⁵ it gives rise to what I call *the problem of harmful distrust*. In this paper, I discuss three general approaches to public trust in science and their proposed solutions to the problem of harmful distrust. In Sect. 2, I outline a general taxonomy of relationships of trust and use it to distinguish those three general approaches, which I call *the individual approach*, *the semi-social approach*, and *the social approach*. In Sects. 3 and 4, I critically examine, respectively, the individual approach and the semi-social approach. I argue that, despite their differences, both approaches embrace individualistic solutions to the problem

¹ See, e.g., Gawande (2016), Tsipursky (2018), Oreskes (2019), Jewett (2020), and Parikh (2021).

² For example, according to a 2019 Pew Research Centre survey, 86% of US adults have a great deal or fair amount of confidence in scientists acting in the best interests of the public [up from 76% in 2016 and compared to 57% for religious leaders, 47% for the news media, and 46% for business leaders (Funk et al., 2019)].

³ For example, according to a Pew Research Center survey, 20% of US adults believe that human activity does not contribute too much or at all to climate change. The number goes up to 45% for conservative Republicans (as compared to 21% for moderate Republicans, 8% for moderate Democrats, and 2% for liberal Democrats) (Funk & Hefferon, 2019).

⁴ For example, 25% of respondents to a CBS News/YouGov Poll conducted in February 2021 said that they would not be getting the COVID-19 vaccine when it became available and 24% said that they might not get it (Salvanto et al., 2021). There is a clear partisan divide among vaccine refusers, with 34% of Republicans and 30% of independents refusing to get the vaccine compared to only 10% of Democrats. Among those who might not or would not get the vaccine, some of the reasons given were that they did not trust the scientists and companies that make it (27%), that they never get vaccines (22%), and that they were not concerned about coronavirus (20%).

⁵ It is important to note that distrust in science is not always necessarily socially harmful or unwarranted. For example, flat-earthism might be a possible case of *harmless distrust*. While flat-earthers are clearly distrustful of science and scientists, their distrust might not be socially or politically harmful. A possible case of *warranted distrust* is distrust among social groups that have been (or still are) discriminated or stigmatized by the scientific establishment. For example, vaccine hesitancy seems to be high among African Americans (see, e.g., Salvanto et al., 2021). However, given the extent to which racism has shaped (and still shapes) the practice of medicine, this distrust does not seem to be unwarranted. More generally, as I argue later, the problematic relationship between certain areas of biomedical research and pharmaceutical companies seems to warrant a certain level of distrust in the latter.

of harmful distrust and that, as a result, both approaches face two general problems, which I call *the problem of overidealizing science* and *the problem of overburdening citizens*. In Sect. 5, I argue that in order to avoid these problems we need to embrace a (thoroughly) social approach to public trust in science, which emphasizes the social dimensions of the reception, transmission, and uptake of scientific knowledge in society and the ways in which social forces influence both positively and negatively the trustworthiness of science. The social approach encourages social solutions to the problem of harmful distrust—in order to improve public trust in science, we need to improve what I call the socio-epistemic infrastructure of society.

2 Three Approaches to Public Trust in Science

In this section, I sketch a general taxonomy of relationships of trust (or distrust) based on the occupants of the trustor and trustee positions and use it to distinguish three approaches to public trust in science.⁶ In what follows, I refer to the occupants of a trust relationship (i.e., ‘*x* (dis)trust *y*’) as, respectively *the trustor* and *the trustee*. So, for example, if Jane distrusts Joe, then Jane is the trustor and Joe is the trustee in that trust relationship.⁷ For the sake of simplicity, I first focus only on relationships of trust that have individuals or groups as their *relata*. This gives rise to four kinds of trust relationships—i.e.:

1. *individual-to-individual* trust relationships (e.g., Jane might trust Donald Trump but distrust Hillary Clinton),
2. *individual-to-group* trust relationships (e.g., Jane might distrust politicians as a group (even if she might still trust specific politicians)),
3. *group-to-individual* trust relationships (e.g., evangelicals might trust Donald Trump but distrust Hillary Clinton), and
4. *group-to-group* trust relationships (e.g., evangelicals might distrust politicians in general (even if not all evangelicals distrust all politicians)).

In addition to these four kinds of trust relationship, we can add trust relationship in which the trustee is an institution, which give us two additional kinds of trust relationship⁸:

5. *individual-to-institution* trust relationships (e.g., Jane might trust the Supreme Court but distrust Congress), and

⁶ In what follows, I use ‘trust relationship’ to cover both relationship of trust and relationships of distrust. Distrust, as I understand it here, is not simply the absence of trust. It is an attitude akin to trust but with opposite valence [for a more detailed discussion of the relationship between the two, see Hawley (2014)].

⁷ Since this terminology applies also to cases of distrust, it is important to note that ‘the trustee’ does not necessarily refer to some entity that is trusted by the trustor. It only refers to an entity that is in a position to be trusted or distrusted by the trustor.

⁸ While institutions can also be trustors, I ignore this possibility here.

6. *group-to-institution* trust relationships (e.g., evangelicals might trust the Supreme Court but might distrust Congress).

In what follows, I use the label '*individual trust relationships*' for trust relationships in which the trustor is an individual and '*social trust relationships*' for those in which the trustor is a group. Also, I refer to trust relationships as '*individual-oriented*,' '*group-oriented*,' or '*institution-oriented*' depending on the sort of entity that plays the role of the trustee in the relevant trust relationship. The taxonomy and its associated terminology are summarized in Table 1.

Once we look at the question of public trust in science through the lens of the above taxonomy, it becomes apparent that it is possible to understand public trust in science in a number of different ways. Here, we can focus on two crucial distinctions. The first (and, for our purposes, the most crucial) distinction is the one between approaches that understand public trust in science as, primarily, an individual trust relationship (i.e., the individual and the semi-social approaches) and approaches that conceive of it as, primarily, a social trust relationship (i.e., the social approach, which I defend in this paper). One of the main differences between these approaches is that, according to the individual and the semi-social approaches, social trust in science depends on (and is explained by) individual trust in science, while the reverse is true of the social approach. As a first approximation, the individual and the semi-social approaches maintain certain societies (and certain social groups within them) (dis)trust science (or scientists) because their members tend to (dis)trust science (or scientists), while the social approach maintains individuals tend to (dis)trust science (or scientists) because they belong to societies (or to specific social groups within them) in which (dis)trust in science is prevalent. One of the most important consequences of these two different understandings of public trust in science is that they suggest different solutions to the problem of harmful distrust. The individual approach and the semi-social approach tend to favor individualistic solutions to the problem of harmful distrust (i.e., solutions at the individual level), while the social approach tends to favor social solutions (i.e., solutions at the social level).

The second crucial distinction is the one between approaches that understand public trust in science as, primarily, individual-oriented (as the individual approach) and those that understand it as, primarily, institution-oriented (as the semi-social and the social approaches).⁹ According to the individual approach, institution-oriented trust depends on (and is explained by) individual-oriented trust, while the reverse is true for the social and the semi-social approaches. As a first approximation, the individual approach maintains people trust science only insofar as they trust individual scientists, while the semi-social and the social approaches maintain that people trust

⁹ A third distinction that is important but that will not be explored in detail in this paper is the one between accounts that conceive of public trust in science as, primarily, group-oriented and accounts that conceive of it as, primarily, institution-oriented. The question, in other words, is whether people trust scientists because they trust science or vice versa. While this is an interesting question, it is beyond the scope of this paper.

individual scientists only insofar as they trust science as an institution (and they see the individual scientist as acting on behalf of that institution).

Somewhat schematically, we can thus summarize the main differences between the three approaches as follows. The individual approach conceives of public trust in science as, primarily, an individual-to-individual trust relationship—i.e., as a trust between individual citizens and individual scientists. The semi-social approach conceives of public trust in science as, primarily, an individual-to-institution trust relationship—i.e., a trust relationship between individual citizens and the institutions of science. Finally, the social approach understands it as, primarily, a group-to-institution trust relationship—i.e., a trust relationship between society in general (and certain social groups within society in particular) and the institutions of science. On each of these approaches, other forms of trust in science depend on (and can be explained in terms of) what the account takes to be the fundamental trust relationship. In the next three sections, I discuss each of these approaches in turn.

3 The Individual Approach

As we have seen in the previous section, the individual approach understands public trust in science as, primarily, a trust relationship between individual citizens and individual scientists. Perhaps, the simplest version of the individual approach is what we might call *the naïve account* of public trust in science. According to the naïve account, ordinary citizens should trust scientists because (individual) scientists are disinterested and objective seekers of the truth who follow the scientific method, which is the most reliable way to establish the truth about that domain. While, to my knowledge, nobody nowadays explicitly accepts (let alone defends) the naïve account, something along its lines seems to be still popular among scientists and the general (trustful) public.

One of the theses of this paper is that accounts of public trust in science tend to face two general problems, which I call, respectively, *the problem of overridealizing science* and *the problem of overburdening citizens*. The naïve account faces a particularly severe version of the problem of overridealizing science. This is because the naïve account relies on a picture of science which, nowadays, is almost universally rejected by science scholars, and which we might call *the naïve view of science*. Since the challenges to the naïve view of science are well-known, I only mention the three most relevant for our purposes here. The first challenge is that the task of explaining exactly what the scientific method is has proven exceptionally difficult, so difficult that today most philosophers of science seem to have given up entirely on the project of identifying a set of principles that are both sufficiently informative and sufficiently general to deserve the honorary title of ‘scientific method.’ The second challenge is that scientists who have access to the same body of empirical evidence often seem to disagree with each other about which theory or hypothesis is best supported by the available evidence. These disagreements seem to suggest that, even if there is such a thing as the scientific method, either it is not applied correctly by all scientists, or it does not yield univocal results. The third challenge is that the picture of individual scientists as objective and disinterested truth-seekers is, at best, highly

Table 1 A taxonomy of relationships of trust

		Trustee			
		Individual	Group	Institution	
Trustor	Individual	Individual-to-Individual	Individual-to-Group	Individual-to-Institution	← <i>Individual Trust Relationships</i>
	Group	Group-to-Individual	Group-to-Group	Group-to-Institution	← <i>Social Trust Relationships</i>
		↑ <i>Individual-Oriented</i>	↑ <i>Group-Oriented</i>	↑ <i>Institution-Oriented</i>	

unrealistic. Since the 1960s, a growing body of work in the history and sociology of science has clearly established that individual scientists are susceptible to the influence of non-epistemic values (such as foibles, idiosyncrasies, social prejudices, personal, political, or economic interests, or metaphysical or religious views) and that these non-epistemic values often affect their scientific decisions.¹⁰ This third challenge seems to be particularly worrisome at a time when the direct and indirect influence of private funding on science is growing and it is increasingly common for scientists to operate under more or less open forms of conflict of interest.

In light of these challenges, supporters of the individualist approach are likely to reject the naïve account (and the naïve view of science on which it relies). However, if they do so, they seem to run into the second of the general problems I mentioned above—i.e., the problem of overburdening citizens. If one accepts the individual approach without assuming that scientists are always and everywhere disinterested and objective truth-seekers that dispassionately apply the scientific method, then the relationship between scientists and ordinary citizens is similar to the relationships between other non-disinterested experts (e.g., mechanics or lawyers) and their clients, which tend to give rise to what Alexander Guerrero (2016) calls strategic expertise contexts. Let me call the resulting version of the individual approach *the consultative account of public trust in science*. The consultative account faces a number of problems that characterize strategic expertise contexts in general. Since these problems are well-known, I only mention two of them here. The first is the problem of how non-experts can (without acquiring a significant level of expertise themselves) distinguish the genuine experts on a certain topic from pseudo-experts, quacks, and charlatans; the second is the problem of how non-experts can (again, without acquiring a significant level of expertise) decide whom to believe in cases of disagreement among genuine experts.¹¹ While it is possible to find a number of attempts to mitigate the problems that arise from strategic expertise contexts, most of these attempts seem to be based on unrealistic expectations about the (epistemic and non-epistemic) resources that ordinary citizens are both willing and able to

¹⁰ On the influence of social prejudices on science, see, e.g., Lloyd (2005) or Gould (2006). On the influence of economic interests on science, see, e.g., Angell (2005) and Oreskes and Conway (2010).

¹¹ See, e.g., Goldman (2001).

devote to ascertaining who the genuine experts are or who to believe among disagreeing experts.¹²

The individual approach, thus, seems to be caught between two sets of unrealistic expectations. On the one hand, the naïve account seems to rely too heavily on the objectivity and disinterestedness of individual scientists and on the existence of a set of unambiguous norms that deserve the honorary title of ‘the scientific method’ (which is an instance of what I have called the problem of overidealizing science). On the other hand, the consultative account seems to rely too heavily on the limited resources of ordinary citizens to settle disagreements between (genuine or self-proclaimed) experts (which is an instance of the problem of overburdening citizens).

4 The Semi-social Approach

If the individual approach understands public trust in science as, primarily, based on a trust relationship between individual citizens and individual scientists, the semi-social approach understands it as, primarily, based on a trust relationship between individual citizens and science.

The semi-social approach typically rejects the naïve view of science in favor of what we might call *a social view of science*. The social view of science meets the challenges that beset the naïve view by adopting something along the line of the following three-step strategy. The first step is to distinguish the goals of science from the goals of individual scientists. Even if individual scientists have their own personal goals (such as professional success), these goals can be harnessed in the pursuit of the goals of science. Scientists know that, insofar as they want to be professionally successful, they need to act so as to promote the goals of science as their own goals and they have to act in accordance with the ethical and epistemic norms of the scientific community.¹³ The second step is to adopt a broadly pragmatic understanding of the goals of science. As a society, we might be content if science can provide us with reliable knowledge that can be used to pursue our personal and collective goals with a reasonable degree of success (i.e., with a greater degree of success than relying on common sense, intuitions, or divination).¹⁴ The third step is to maintain that, insofar as there is anything that deserves the label ‘scientific method,’ it is best understood as a set of social norms that regulate the activities of the scientific community as a whole (rather than a set of normative principles to which individual scientists adhere). On this view, the scientific method can be understood as *a social system of epistemic checks and balances* whose function is

¹² For a critical review of some of the literature on the relationship between non-disinterested experts and their clients, see Guerrero (2016).

¹³ Note that none of the above requires denying that many or even most scientists wholeheartedly adopt the goals of science as their own, independently of their other goals. The point is rather that it is not necessary for scientists to be disinterested truth-seekers for them to pursue the goals of science (whatever these might be).

¹⁴ If, as a by-product, science can also provide us with some understanding of what the world is like or satisfy our intellectual curiosity, then all the better for it.

to safeguard the reliability of scientific results. For example, scientific communities rely on mechanisms such as peer-review or replication to deter unscrupulous scientists from trying to game the system and they apply very severe sanctions to those scientists who are found to have violated those norms.

One of the most popular versions of the social view of science is the one proposed by Helen Longino.¹⁵ According to Longino, a scientific community is objective to the extent to which it meets certain conditions, which include sharing a set of standards, displaying a level of social and epistemic diversity, granting comparable epistemic authority to all of its members in good standing (irrespective of their professional status or social identity), providing recognized avenues for criticism of theories, assumptions, and methods, and engaging with those criticisms (as opposed to merely ignoring them or side-lining their proponents) (see, in particular, Longino, 1990, chap. 4).

The semi-social approach to public trust in science builds on the social view of science by claiming that, while one cannot necessarily assume that individual scientists are disinterested and objective truth-seekers, the scientific community as a whole can be trusted to pursue the goals that society expects it to pursue (whether it is truth or some lesser goal, such as empirical adequacy). The most popular variant of the semi-social approach to public trust in science is what we might call *the consensus account of public trust in science* [which is more or less explicitly embraced by Oreskes (2019) and Anderson (2011)].¹⁶ According to the consensus account, the consensus of the relevant scientific community over a certain hypothesis is a sufficient (and, possibly, necessary) condition for ordinary citizens to be justified in accepting (or rejecting) that hypothesis. For example, Anderson writes:

¹⁵ See, in particular, Longino (1990). For other prominent versions the social view of science, see Hull (1988), Kitcher (1993), and Solomon (2001).

¹⁶ I should note that neither Anderson nor Oreskes embrace the consensus account as I outline it here. First, both authors include elements of what I called the consultative account in their respective accounts. Second, both Anderson and Oreskes identify a number of additional criteria for lay acceptance beside the consensus of the scientific community. For example, in addition to scientific consensus, Anderson's account requires scientific experts to display what she calls dialogic rationality and to have the right sort of qualifications and Oreskes's account includes four additional criteria, which she labels 'method,' 'evidence', 'values,' and 'humility.' In this paper, I focus on a pure version of the consensus account, as the focus on consensus is what makes the consensus account a semi-social account and differentiates it from the consultative account. Moreover, many of my criticisms of their accounts would apply *mutatis mutandis* to the additional criteria in Anderson's and Oreskes' respective accounts. For example, if, as I argue below, it is unreasonable to expect ordinary citizens to be willing and able to devote their epistemic and non-epistemic resources to determine whether a sufficiently broad consensus holds in the relevant community, then it is also unreasonable to expect them to be willing and able to devote their resources to determine whether certain experts display dialogic irrationality or the appropriate level of epistemic humility. Finally, I should note that the consensus account is only one version of the semi-social account. For example, as I interpret it, the idiosyncrasy-free account of public trust in science proposed in Boulicault and Schroeder (2021) is a semi-social account that does not rely on the notion of scientific consensus. As far as I can see, many of the objections to the consensus account would also apply *mutatis mutandis* to the idiosyncrasy-free account. However, for reasons of space I cannot discuss that account here.

Science needs a balance of diverse inquirers to formulate and investigate a wide range of hypotheses, uncover a wide range of relevant evidence, and check one another's biases [...]. When the vast majority of diverse inquirers converge on certain conclusions, as in evolutionary theory, a robust scientific consensus obtains. Before a consensus, the best course for laypersons is to suspend judgment. Once a consensus of trustworthy experts is consolidated, laypersons are well advised to accept the consensus even in the face of a handful of dissenting scientists, or a few instances of error or dishonesty among a few of the participants in the consensus. (Anderson, 2011, 149)

The consensus account explicitly formulates one of the implicit premises of an argument that is often deployed in the context of public debates about anthropogenic climate change [and whose use was spearheaded by Oreskes herself (Oreskes, 2004)]. The argument appeals to the fact that a vast majority of climate scientists believe in anthropogenic climate change, as testified by the fact that, according to an oft-quoted study, 97% of peer-reviewed climate science papers support anthropogenic climate change (Cook et al., 2013). That argument, however, does not provide us with any explicit reason to think that the consensus of the relevant scientific community over the anthropogenicity of climate change gives ordinary citizens reasons to accept that hypothesis and, on the face of it, it is not obvious that we should do so. After all, as Stephen John (2018, 76–77) notes, the fact that astrologers unanimously believe the hypothesis that a Jupiter-Saturn conjunction ushers in the Age of Aquarius does not seem to be a reason to accept that hypothesis. According to the consensus account, however, the missing premise is that, given the system of epistemic checks and balances implemented by the community of climate scientists, the fact that that community has reached a nearly unanimous consensus over a certain hypothesis after subjecting it (and its main alternatives) to critical examination gives ordinary citizens a strong (if defeasible) reason to accept that hypothesis. In contrast, the astrological community does not seem to have any system of epistemic checks and balances in place.¹⁷

On the face of it, the consensus account seems to avoid the problems that beset the individual approach. On the one hand, it seems to avoid the problem of overidealizing science by adopting a more realistic view of science. On the other hand, it seems to sidestep the problem of overburdening citizens by relying on a seemingly simple criterion (i.e., the consensus of the relevant scientific community) for lay acceptance of scientific hypotheses.

On closer scrutiny, however, the consensus account turns out to be vulnerable to different versions of the same problems. Like the naïve account, it, too, seems to presuppose an excessively idealized picture of the workings of real-world scientific communities (which gives rise to a version of the problem of the overidealizing science) and, like the consultative account, it too seems to rely excessively on

¹⁷ While, *pace* Popper, astrologers do seem to make falsifiable predictions, they do not seem to engage in a collective effort to keep track of and check each other's predictions for accuracy or to revise their theories in light of false predictions.

the resources of ordinary citizens (which gives rise to a version of the problem of overburdening citizens). The relevant version of the problem of overidealizing science is that it is not clear how closely real-world scientific communities approximate the ideal of an objective scientific community that underlies social views of science such as Longino's.¹⁸ Just to pick one facet of this problem, consider cases of what we might call *vicious consensus*. The history of science shows that scientific communities can achieve a significant level of consensus over false hypotheses, a consensus that seems to be based more on the shared prejudices and biases than on the empirical support for those hypotheses. A classic example of vicious consensus is the broad consensus among 19th-century physical anthropologists over the hypothesis that there are innate differences in intelligence between members of different "races."¹⁹ Cases of vicious consensus seem to show that the prevalence of prejudices and biases in a scientific community can severely inhibit the effectiveness of the self-correcting mechanisms presupposed by the social view of science. Given that all sorts of implicit or explicit prejudices and biases (including racist, sexist, and classist prejudices) are still widespread in contemporary societies, it is unlikely that our scientific communities are completely immune from their influence.²⁰

Unsurprisingly, Oreskes, who is an accomplished historian of science, anticipates this objection. Her response to it is that 'in all cases [of vicious consensus], there was significant, important and *empirically informed dissent within the scientific community*' (Oreskes, 2019, 128; emphasis in the original). Even if, for the sake of the argument, we grant the truth of Oreskes' broad claim, this response seems to give rise to a dilemma. Either the consensus account adopts unanimity as the criterion of lay acceptance, or it adopts a less stringent criterion. If it adopts unanimity, then it is likely that this criterion will be rarely met even in cases of virtuous consensus, as scientific communities rarely achieve full unanimity over any interesting scientific hypothesis. Moreover, from the perspective of social views of science, this lack of unanimity is to be encouraged, as scientific disagreement is what fuels the self-correcting mechanisms of science. Finally, it is not obvious that unanimity would actually help to foster trust in science, as outsiders might worry that the apparent consensus is based on the suppression of dissent.²¹ However, if the consensus account adopts a less stringent criterion for lay acceptance, then Oreskes' response fails to distinguish between vicious and virtuous consensus, as, in both cases, the consensus is likely to be less-than-unanimous. Advocates of the consensus account would

¹⁸ A view on which Oreskes (2019, 49–54) explicitly relies.

¹⁹ See, e.g., Tucker (1994) and Gould (2006).

²⁰ For example, as feminist science scholars have persuasively argued, sexist prejudices still underlie much contemporary scientific research (see, e.g., Lloyd, 2005; Fine 2011).

²¹ It is unclear whether and to what extent the perception of a scientific consensus over a certain hypothesis leads to lay acceptance. While there is some empirical evidence that perception of a scientific consensus on anthropogenic climate change is correlated with belief in anthropogenic climate change (see, e.g., van der Linden et al., 2015, 2018), the strength of that correlation has been questioned (Kahan 2017). More importantly, it is unclear what causal structure gives rise to that correlation, as there is some evidence that people's perception of the level of scientific consensus is affected by their beliefs and their values (Kahan et al., 2011), which suggests that people who are less likely to believe in climate change (on political or religious grounds) are also less likely to perceive a consensus among the scientists.

therefore have to identify additional criteria to distinguish cases of (less-than-unanimous) virtuous consensus from cases of (less-than-unanimous) vicious consensus and it is unclear what these criteria might be (see, e.g., de Melo-Martín & Intemann, 2018). Moreover, as Oreskes and her co-author Erick Conway have abundantly demonstrated (Oreskes & Conway, 2010), even in cases of virtuous and near-unanimous consensus, it only takes a handful of highly motivated, unscrupulous, and outspoken scientists to mislead the general public into believing that a certain hypothesis is scientifically controversial.²² Oreskes might reply that these cases are not cases of empirically informed dissent. However, the task of distinguishing cases of spurious dissent from the sort of genuine dissent that, according to the social view, drives the self-correcting mechanisms of science is far from a simple task (see, e.g., de Melo-Martín & Intemann, 2018). Moreover, even assuming, for the sake of the argument, that supporters of the consensus account could provide the general public with a set of more specific criteria to distinguish between virtuous and vicious consensus and between spurious and genuine dissent, the consensus account would still run into the problem of overburdening citizens. It seems unrealistic to expect ordinary citizens to have the (epistemic and non-epistemic) resources to apply any such hypothetical set of criteria to every single case ranging from the reality of anthropogenic climate change to the safety of vaccines.

It is important to note that these two problems are not specific to Oreskes' particular version of the consensus account—they seem to be a consequence of two general features of the semi-social approach. The semi-social approach seems to presuppose that the problem of harmful distrust is primarily the problem of persuading individual citizens that science is a trustworthy institution. However, there seem to be two general problems with this. The first is that it is far from obvious that science as it is currently practiced is fully trustworthy (which is, once again, an instance of the general problem of overidealizing science). Just to focus on one facet of this problem for illustrative purposes, it seems that science's current system of external and internal incentives is not always best suited to promoting the goals that society expects it to achieve. As far as external incentives are concerned, science is increasingly reliant on private sources of funding (whether directly or indirectly).²³ A field in which the pernicious effects of this reliance on private funding are particularly obvious is biomedical research. As Jim Brown (2017) has argued, private funding influences biomedical research negatively in two ways—it *biases* it (as it creates a conflict of interest for the researchers, who are funded by the very corporations that

²² In fact, recent work in social epistemology shows that it is not even necessary to rely on less-than-honest scientists to generate the perception of a lack of consensus. A propagandist can create the appearance of dissent just by cherry-picking the results that best fit their purposes (see, e.g., Weatherall et al., 2020).

²³ According to OECD estimates, in most countries, the business sector expenditures on research and development (R&D) account for 60% of total R&D expenditures against 20% by universities and 10% by governments (OECD 2015, 159). On the potential of conflict of interest to undermine public trust in science, see, e.g., Elliott (2014) and Almassi (2017).

stand to gain or lose financially from their research),²⁴ and it *skews* it (as it promotes certain research directions (such as, e.g., development of new patentable drugs) over others (such as, e.g., research on off-label uses of old drugs, non-pharmaceutical interventions, or lifestyle changes) on the basis of considerations of profitability). The influence of private funding on biomedical research has resulted in a number of notorious scandals that have contributed to tarnishing the public's perception of the trustworthiness of that research in particular and of science in general.²⁵

The internal incentives of scientific communities, however, can also contribute to skewing and biasing their research. Consider, for example, the replication crisis that has recently engulfed social psychology. The crisis originated with a large-scale attempt to replicate some of the most influential results in social psychology, which resulted in a failure to replicate the results of most of the original studies.²⁶ Of the many factors that are likely to have contributed to the crisis, one factor illustrates particularly clearly both the biasing and skewing effects of the internal incentives (and disincentives) adopted by the community. Scientific communities tend to have a preference for surprising results, which means that such results are more likely to get published and cited, with all that this entails in terms of professional prestige and academic career. This contributes both to biasing research (as researchers have an incentive to actively search the data for surprising results that might, in fact, be just an artefact of poor experimental design or a statistical fluke) and to skewing it (as replication studies are considered much less prestigious than original studies).²⁷ The replicability of scientific results is considered a cornerstone of the scientific method. However, replicability without replication is unlikely to serve the self-correcting function it is supposed to serve, and it is unlikely that the scientific community as it is currently organized will achieve the optimal balance between finding surprising results and probing the robustness and replicability of those results.²⁸

While the semi-social approach rightly emphasizes how the norms adopted by scientific communities can function as a system of epistemic checks and balances, it seems to largely ignore the fact that scientific communities do not operate fully

²⁴ For a qualitative review of the biasing influence of pharmaceutical funding on biomedical research, see Sismondo (2008).

²⁵ According to a Pew Research Center survey, while 74% of US adults have a mostly positive view of medical doctors only 68% have a mostly positive view of research scientists in the biomedical sciences; only 35% believe that biomedical researchers care about people's best interests (as opposed to 57% for doctors), only 32% that biomedical researchers provide fair and accurate information (as opposed to 48% for doctors), only 15% believe that biomedical researchers and doctors are transparent about conflict of interest, and, finally, only 12% believe that biomedical researchers admit and take responsibility for their mistakes (a similar percentage (13%) believe the same of doctors) (Funk et al., 2019). For a social-epistemological discussion of one of these scandals, see Biddle (2007). For popular books on the negative influence of private funding on biomedical research, see Angell (2005) and Goldacre (2012); for a discussion of how conflict of interest erodes trust in science, see, e.g., Friedman (2002) and Goldenberg (2016).

²⁶ For an overview of the replication crisis and some of the philosophical issues it raises, see Romero (2019). See also Jon Krosnick's reply to Oreskes' lectures (Krosnick 2019).

²⁷ See, e.g., Romero (2017).

²⁸ For a discussion of some of these issues, see Romero (2020). For a more general assessment of the desirability of the peer-review system, see (Heesen & Bright, 2021).

independently from society and that the norms they currently adopt are not always best suited to the achievement of the goals of science. In particular, the semi-social approach fails to openly acknowledge that the trustworthiness of science as an institution is partly shaped by external and internal social forces and that these forces often limit the ability of science to self-regulate so as to be fully trustworthy. Moreover, by discounting the fact that science as it is currently practiced is not always fully trustworthy, the semi-social account fails to address some of the legitimate concerns that fuel distrust in science in certain sectors of the public. An adequate account of public trust in science should acknowledge that, as it is currently practiced, science is not always fully trustworthy and should develop proposals to improve the trustworthiness of science rather than simply presupposing it. As philosophers working on trust often note (see, e.g., Hardin, 2006), trust is not intrinsically good, as we should not fully trust those who are not fully trustworthy. While, overall, science might be more trustworthy than many other institutions, it does not mean that an account of public trust in science should ignore the ways in which its trustworthiness can be improved.

The second general problem with the semi-social approach stems from its oddly divided epistemology. The semi-social approach seems to presuppose that it is social epistemology for scientists but individualistic epistemology for everyone else. While it correctly emphasizes the social dimensions of the production of knowledge within the scientific community, it primarily relies on an inadequate individualistic picture of the reception, transmission, and uptake of that knowledge by society, a picture that relies almost exclusively on individual citizens (which, again, leads to an instance of the problem of overburdening citizens). As a result, the semi-social approach largely ignores the social dimensions of the reception, transmission, and uptake of that knowledge.²⁹

Let me focus on one specific facet of this problem to illustrate the general problem. The problem arises from the communication channels between science and the general public. Most people do not get their scientific information directly (e.g., from scientists or scientific institutions) but indirectly through traditional media and, increasingly, through social media (see, e.g., (Funk et al., 2019)). This gives rise to a number of issues that the semi-social approach is ill-equipped to address. Let me mention one issue here, which concerns the effectiveness of trying to persuade citizens who are distrustful of science of its trustworthiness. The problem is that, to some extent, we all live in an epistemic bubble that largely determines what information (and, even more critically, *misinformation*) reaches us. This includes not only scientific (mis)information but also the very attempts to persuade distrustful citizens of the trustworthiness of science, which means that it is, at best, doubtful that the attempts to persuade distrustful citizens will even reach their targets. Even worse, many distrustful citizens seem to be part of echo chambers that actively undermine the credibility of outside sources.³⁰ In light of these considerations, it is likely that

²⁹ For an overview of the social dimensions of the reception, transmission, and uptake of scientific knowledge, see, e.g., O'Connor and Weatherall (2019).

³⁰ For the distinction between epistemic bubbles and echo chambers, see Nguyen (2020).

the sort of persuasion attempts advocated by semi-social approach are likely to be ineffective (in the case of epistemic bubbles) or counterproductive (in the case of echo chambers).³¹

In the next section, I argue that, in order to overcome these two general problems with the semi-social approach, we need to embrace a thoroughly social approach to public trust in science, one that understands both the production of scientific knowledge and its reception and uptake by society as thoroughly social and that sees science as completely embedded in society and dependent on it to improve its trustworthiness.

5 The Social Approach

In this section, I argue that, in order to avoid the problems faced by the individual and semi-social approaches, we need to adopt a (thoroughly) social approach to public trust in science. The outline of the social approach that I sketch here is admittedly incomplete and underdeveloped. The details will need to be worked out elsewhere (and, in some cases, they are already being worked out).³² However, I hope that, in spite of its limitations, the outline that I sketch here can help lay the groundwork for a unified social approach to public trust in science that weaves together a number of strands of research on the social epistemology of science broadly construed, stimulate more work in this direction, and, hopefully, promote the development of different and more effective solutions to the problem of harmful distrust.

The social approach rejects the two assumptions that underlie the semi-social approach and that, I claimed, were at the root of its problems. First, the social approach rejects the divided epistemology presupposed by the semi-social approach in favor of a unified social epistemology, one that acknowledges the social dimensions of not only the production of scientific knowledge by the scientific community but also its reception, transmission, and uptake by the rest of society. According to the social approach, it is social epistemology for all. Second, the social approach rejects the semi-social approach's implicit assumption that science is largely insulated from society and that it is already fully trustworthy as it is currently practiced.

³¹ To be fair, both Anderson and Oreskes acknowledge the influence of social factors on the beliefs of individual citizens (and, in fact, Anderson explicitly discusses the issue of news media as a source of scientific misinformation). However, in spite of this, their accounts are still not (thoroughly) social, as they still understand public trust in science as, primarily, based on a trust relationship between individual citizens and the scientific community and they regard social factors as interfering with the otherwise voluntary and rational epistemic choices of individual citizens.

³² Although, to my knowledge, the social approach to public trust in science that I sketch here has not been explicitly defended, there are many similarities between my account and the views on public trust in science developed by feminist epistemologists and philosophers of science (see, e.g., Scheman, 2001; Grasswick, 2010, 2014). Moreover, many contributions to the literature on neighbouring topics presuppose a social approach to public trust in science or develop and defend aspects of the general account sketched here (see, e.g., Moore, 2017; de Melo-Martín & Intemann 2018; Goldenberg, 2021). This paper aspires to sketch a general framework in which all these different strands of research can be weaved together.

According to the social approach, science is fully embedded in society and, as a result, society can affect both positively and negatively to the trustworthiness of science.

As we have seen in Sect. 2, the social approach conceives of public trust in science as, primarily, based on a trust relationship between society (and various social groups within it) and the scientific community. Of course, to say that trust in science is primarily based on group trust is not to deny that social groups trust science only insofar as a suitable proportion of their members trusts science. Rather, it is to deny that members of the group come to trust or distrust science independently of one another. One of the most distinctive characteristics of a social approach to public trust in science is that, while individual and semi-social approach conceive of trust in science as, primarily, the result of a rational and voluntary decision made by each individual citizen independently of others, the social approach conceives of individual trust in science as the result of a habit that is neither fully voluntary nor fully rational and that is largely the result of the attitudes towards science held by those in one's social and epistemic network.

As we have seen, according to the individual and semi-social approaches, individual citizens should trust science because science is a reliable source of knowledge and, if they do not trust science, they need to be persuaded that science is trustworthy. On these accounts, the solution to the problem of harmful distrust requires educating individual citizens about science and persuading them to trust science. Even when semi-social approach acknowledges the influence of social factors, they still see them as, primarily, interfering with the otherwise rational and voluntary decisions of individual citizens.

The social approach, on the other hand, conceives of trust in science as, primarily, social and it maintains that attempts to solve the problem distrust should focus primarily on social and structural changes rather than on interventions at the individual level. Consider, for example, the extreme case of someone (let's call him Aaron) who grew up in a traditionalist religious community that denounces science as inimical to religion and as a gateway to atheism and who, as a result, is distrustful of science. For Aaron, distrust in science is a profoundly held attitude and it is unlikely that people or institutions from outside Aaron's community could persuade Aaron to trust science. In fact, as we have seen, any efforts to persuade Aaron are likely to be ineffective if not counterproductive (if they even reach him). Of course, this does not mean that it is impossible for Aaron to bring himself to trust science. Rather, it means that bringing himself to trust science for Aaron would require a concerted and conscious effort that we cannot expect him to be motivated to make insofar as he is a full member of a community that is profoundly distrustful of science.

Admittedly, the case of Aaron might be an extreme case.³³ However, according to social approach, the difference is a matter of degree—whether individuals trust

³³ However, if, as Nguyen (2020) persuasively argues, escaping an echo chamber is as epistemically difficult as escaping a religious cult, then the difference between Aaron and other citizens who are distrustful of science might be a matter of degree. As Aaron's community as I have described it does seem to have the characteristics of an echo chamber and so do many of the communities that foster distrust in science.

(or distrust) science is, by and large, a function of the trust in science in their social environment and their epistemic network. This contention is indirectly supported by the evidence. For example, while science is a global institution, distrust in science seems to vary from country to country (see, e.g., Wellcome Global Monitor, 2018) and, within a specific country, from social group to social group (see, e.g., Funk et al., 2019). These variations suggest that public trust in science is largely a social and cultural phenomenon. Moreover, if one is part of a social group (or a society) that (dis)trusts science, a number of social and cognitive mechanisms are likely to contribute to entrench that (dis)trust,³⁴ and it is difficult for external forces to change that.

But what does it mean for a society (or a social group) to trusts science? I think that it is important to emphasize that a society that trusts science is not necessarily one whose members have a positive opinion of science or scientists. A society can trust science without most of its members having any explicit views about the trustworthiness of science or scientists—a society that trusts science is, first and foremost, a society that collectively relies on science to inform its actions and decisions (and those of its members). Public trust in science can be understood as a collective (non-doxastic) disposition (Kappel, 2014) or as a collective (unquestioning) attitude (Nguyen, forthcoming) toward science. Whenever the question of whether we should trust science arises, then, usually, trust in science has already been eroded to some extent. On this view, full trust is (typically)³⁵ implicit trust. As Annette Baier aptly put it: ‘We inhabit a climate of trust as we inhabit an atmosphere and notice it as we notice air, only when it becomes scarce or polluted.’ (Baier, 1986, 234) Consider an analogy. If a friend invites you for dinner, you implicitly trust that they will not poison you. The question of whether they have poisoned your food should not even cross your mind. If the question were to genuinely cross your mind, then that would be evidence that you don’t fully trust your friend.

On the social account, a society that trusts science is not necessarily one in which the citizens have a positive opinion of science or scientists. But, if trust in science is not necessarily exemplified by the explicit opinions of citizens what does it consist in? According to the social account, a society trusts science insofar as it exemplifies an efficient division of epistemic labor between scientists and ordinary citizens. As social epistemologists argue, a significant degree of division of epistemic labor is essential to the functioning of any society and, particularly, to that of contemporary societies. A society that exemplifies a perfectly efficient division of epistemic labor is one in which people rely on the knowledge of the relevant experts for their individual and collective decisions. They rely on the knowledge of mechanics to

³⁴ For a review of the empirical evidence of “social contagion” with regards to attitudes about vaccines, see Konstantinou et al. (2021). A variety of hypotheses has been suggested to explain this form of social contagion [for a discussion and an attempt to settle this debate, see Kahan (2013)]. One hypothesis is that the social contagion is due to cognitive biases, such as conformity bias [which has been the focus of a lot of work in social psychology since Asch (1951)]. Another hypothesis is that the contagion is due to identity-protective motivated reasoning—i.e., the tendency of individuals to form beliefs that conform with the ideological or cultural values shared with their social group(s) (see, e.g., Kahan et al., 2011).

³⁵ Although this is not always necessarily the case (see, e.g., Keren, 2019).

determine what's wrong with their cars and on the knowledge of doctors to diagnose their medical conditions. A society that trusts science is one that collectively relies on the knowledge of climate scientists to settle the question of whether human activity is causing climate change and on the knowledge of biomedical researchers to settle the question whether vaccines are safe.

Obviously, there are both theoretical and practical limits to how closely real-world societies can (or should) approximate a perfectly efficient division of epistemic labor. Let me mention two here. A first limit arises from the state of what we might call *the socio-epistemic infrastructure* of society. Many experts offer their services through the market, and this might give rise to conflicts of interest between experts and their clients (e.g., a mechanic might recommend a more expensive repair when a cheaper one would be equally effective, a lawyer might recommend litigation even when settling would be in the client's best interest). The extent to which a society can approximate a perfectly efficient division of cognitive labor is largely determined by the quality of its socio-epistemic infrastructure. The social approach has a very broad understanding of what counts as part of the socio-epistemic infrastructure. For example, the social approach would include a country's healthcare system in its socio-epistemic infrastructure, as many of the potential conflicts of interest that might arise in the relationship between doctors and their patients in the context of a private or semi-private health care systems do not arise in the context of a fully public healthcare systems. Similarly, a society in which scientific research is primarily funded privately is a society that has a poor socio-epistemic infrastructure that tends to undermine trust in science.

A second limit concerns the role played by non-epistemic values in science. There is a growing consensus among philosophers of science that even scientific decisions that seem to be purely epistemic often require scientists to take non-epistemic values into account.³⁶ A standard example of this is inductive risk—i.e., the risk involved in accepting or rejecting a scientific hypothesis (see Rudner, 1953; Douglas, 2009). The idea is that, since scientific hypotheses are never conclusively verified or falsified, whenever scientists accept or reject a hypothesis, they should consider the non-epistemic consequences of a potential error (i.e., the non-epistemic consequences of accepting a false hypothesis or rejecting a true one). For example, if scientists incorrectly accept the hypothesis that the widely used herbicide glyphosate is carcinogenic to humans, then this might result in an unnecessary ban on glyphosate the substance that might affect crop yields and food prices. If, on the other hand, they incorrectly reject the hypothesis, then this might result in a higher rate of incidence of certain cancers, which would result in unnecessary deaths and human suffering. When choosing how strict the epistemic standards for acceptance or rejection of the hypothesis that glyphosate is a human carcinogen, scientist should weigh the potential consequences of error. It seems that, as a society, we should try to ensure that the values employed in the process are representative of those held by society at large (see, e.g., Elliott, 2017; Schroeder, 2021). However, this seems to limit the extent to which we can have a perfectly efficient division of cognitive labor.

³⁶ For an introduction to this topic, see Elliott (2017).

In general, according to the social approach, the problem of harmful distrust is usually the result of the breakdown of an efficient division of epistemic labor caused by a degraded socio-epistemic infrastructure. On this view, the issue with the solutions to the problem of harmful distrust put forward by the individualistic and the semi-social account is not only that they have unrealistic expectations about the resources ordinary citizens are willing and able to devote to sorting through the scientific (mis)information they receive or that they rely on an excessively idealized picture of the trustworthiness of science, but that their proposed solutions might actually contribute to the problem. According to the social approach, to expect ordinary citizens to do more than their basic epistemic due diligence on issues such as vaccine safety or climate change (as the individual and semi-social approach suggest) is likely to have the unintended effect to further undermining the efficient division of epistemic labor. The division of epistemic labor exists exactly so that we do not have to “do our own research” on each and every topic that is relevant to our own good, or that of our family, our community, our society, or the world. It is practically impossible for each citizen to do their own research on such a vast range of topics and, even if it wasn't, it would be unlikely that most would reach warranted conclusions on the basis of such research. In fact, anecdotal evidence seems to suggest that people who “do their own research” are the ones who are more likely to reject the scientific consensus on specific topics. If a doctor prescribes you a drug, you should be able to assume that the drug is relatively effective, that it does not cause side effects, and that the specific batch you are taking has undergone a sufficient level of quality control to ensure that the right amount of the medicinal ingredient is present, etc. If you had to do your own research on each of these issues before taking a prescription drug, then you'll probably never end up taking that medicine.

On the social approach, a society that fully trusts science is one that trusts it implicitly. It is a society in which science informs public policy and public debate without becoming itself the object of public debate. Is this an unattainable ideal, such as the one of a perfectly efficient division of cognitive labor? I do not necessarily think so. While a lot of the philosophical literature on public trust in science focuses on cautionary tales, we can perhaps learn as much from comparing them to the success stories. A rare but very instructive success story is the one that led to the discontinuation of the production chlorofluorocarbons (CFCs). Once, CFCs were widely used as refrigerants in products such as refrigerators and air-conditioning units as well as propellants in spray cans. However, in the mid-1970s, scientists started to realize that the chlorine released by the decomposition of CFCs in the stratosphere was rapidly depleting the stratospheric ozone responsible for absorbing some of the harmful radiation from the sun. After the scientists sounded the alarm, political leaders of 27 countries, realizing the urgency of the problem, signed the Montreal Protocol in 1987, which heavily regulated the production of CFCs. In the meantime, industry developed a number of substitutes of CFCs, including hydrofluorocarbons (HFCs), which are now commonly used, and production of CFCs ceased entirely in 1995.

In the case of CFCs, the global community instantiated something very close to the ideal of a perfectly efficient division of epistemic labor—the scientific community was quick to raise the alarm about a serious and rapidly-developing

environmental problem, the global political community was quick to take collective political action on that problem, and the private sector was quick to find a technological solution to the problem. The difference between the case of CFCs and the case of climate change is striking. Why did the division of epistemic labor work so well in one case and so poorly in the other? On a superficial reading, the crucial difference between the two cases is the result of luck. In the case of CFCs, the affected industries were able to quickly identify a relatively inexpensive technical solution to the problem, while, in the case of global warming, there is no easy technical solution in sight.

However, to see this merely as a matter of luck seems to miss the more important lesson of comparing the success story with the cautionary tale, which is that, in the case of CFCs, private economic interests did not have a strong economic incentive to forge an alliance with broadcasters, commentators, and politicians to undermine public trust in the relevant science, as in the case of climate change.³⁷ The differences between success stories like the one of the ban on CFCs and cautionary tales like the one of climate change suggest that one of the main factors that disrupts a proper division of epistemic labor is the interference of private economic interests with the production, reception, and uptake of scientific knowledge. This hypothesis is supported by the fact that, in most of cases of distrust, economic interests played a crucial role in hindering the production, reception, and uptake of the relevant bits of scientific knowledge [as, for example, in the many other cases discussed in Oreskes and Conway (2010)]. If this hypothesis is correct, then it seems that part of the solution to the problem of harmful distrust requires breaking the anti-scientific alliance between corporate interests and broadcasters, commentators, and politicians. On the social account, the solution of the problem of harmful distrust might have more to do with public policies, such as anti-lobbying legislation and campaign finance reform, than with trying to persuade ordinary citizens of the trustworthiness of science.

One might argue that we should not overgeneralize on the basis of one success story and a handful of cautionary tales. After all, they might argue, economic interests do not seem to play a role in all the standard cautionary tales. For example, one might argue that distrust of vaccines cannot be explained in terms of economic interests, as no one stands to benefit from lower vaccination rates. However, even leaving aside the issue of whether the leading figures of the anti-vaxx movement benefit from it,³⁸ it seems plausible to claim that economic interests have played a crucial, if less direct, role in distrust of vaccines, too. In particular, it is tempting to believe that public trust in biomedical research has been eroded by its exceedingly cozy relationship with the pharmaceutical industry and by the many scandals to which this relationship has given rise. This erosion of trust has provided fertile ground for distrust in vaccines to propagate.³⁹

³⁷ Although this doesn't mean that some of the usual suspects did not try to sow doubt even in that case (see Oreskes and Conway, 2010, chap. 4). What's remarkable about this case is that strategy did not result in widespread distrust in the specific scientific finding.

³⁸ It is now well-known that Andrew Wakefield, whose now-retracted *Lancet* article was responsible for giving a veneer of scientific respectability to anti-vaccine sentiments, was paid a large sum of money by a group of lawyers who were trying to prove vaccines were unsafe (Deer 2006).

³⁹ See, e.g., Goldenberg (2021).

According to the social approach, however, economic interests are not the only social force that can negatively affect the trustworthiness of science. Scientists and scientific institutions, for example have also contributed to the distrust of science among members of specific communities and social groups (including women, people with disabilities, racial and ethnic minorities, gender non-conforming individuals, etc.) by perpetuating and validating prejudices against members of these communities and groups and by perpetrating injustices against them. Examples of this are the distrust towards the biomedical profession in sectors of the African-American community (see, e.g., Boulware et al., 2003) and the gay community (see, e.g., Hoyt et al., 2012). The distrust in these communities is largely rooted in a history of discrimination and injustice [perhaps best exemplified by the now infamous Tuskegee Study in the case of the African-American community (see, e.g., Brandt, 1978) and by the problematic response to the HIV/AIDS epidemic in the 1980s (see, e.g., France, 2016)]. In these cases, it seems that improving the social diversity of scientific communities and fostering constructive relationship with stakeholders might be a crucial step towards restoring trust in science among members of these groups and communities.

As these all-too-brief remarks suggest, according to the social approach, the problem of harmful distrust calls for primarily social solutions, not individualistic solutions. According to social approach, if we have to resort to persuading individuals to trust science, it is already too late—we already live in a society in which a proper division of epistemic labor has broken down due to a degraded socio-epistemic infrastructure and, unfortunately, it is extraordinarily difficult to unring the bell of scientific misinformation.⁴⁰ Remediating this situation requires implementing large-scale social and political reforms that are aimed at improving the socio-epistemic infrastructure as well as the (actual and perceived) trustworthiness of science rather than trying to persuade individual citizens of the trustworthiness of science.

To be clear, the social approach does not necessarily reject all interventions at the individual level. However, it emphasizes that the effective interventions at the individual level tend to be social in nature as well. For example, while the social approach insists that an adequate account of public trust in science should not rely too heavily on the epistemic and non-epistemic resources of individual citizens, it still supports efforts to improve the critical thinking skills and the scientific literacy of ordinary citizens. However, the social approach emphasizes the fact that even measures that seem to be targeting individuals should be understood as social. The

⁴⁰ The empirical evidence seems to suggest that misinformation and, in particular scientific misinformation, is difficult to correct and that attempts to correct might even backfire. For example, an influential study found that, even when communication strategies aimed at correcting the mistaken belief that the seasonal flu vaccine can cause the flu are successful, the intent to vaccinate might nevertheless decrease as a result (Nyhan & Reifler, 2015). If this is true, then the best approach is to try to develop better strategies to prevent the spread of disinformation. If, as some studies suggest (see, e.g., Germani & Biller-Andorno, 2021), a handful of influencers play a disproportionate role in spreading misinformation online, then the best approach might be to either prevent the spread of disinformation (in an act of epistemic paternalism) or find strategies to reduce the credibility of those sources.

scientific literacy and the critical thinking skills of individuals can only be improved in a society that is fully committed to providing all of its citizens with a high-quality education. The social approach understands the epistemic resources of ordinary citizens to be themselves partially a social product—i.e., they are largely a product of the resources that society is willing to invest in the education of its citizens.⁴¹ This fact seems to be better understood by those who try to undermine public trust in science than by those who try to foster it. For example, the campaign to teach creationism in schools can be seen as part of a more general effort to sabotage public science education as a means to undermining both the public understanding of science and public trust in science.

According to the social approach, however, interventions aimed at improving the scientific literacy and critical thinking skills of individual citizens can only play a limited role. This is partly because, as the social approach emphasizes, many of the issues to which scientific knowledge is relevant are not issues that can be effectively addressed through the actions of individual citizens—they are issues that can only be effectively addressed through public policy. While we have already seen a successful case of this sort of collective action in the case of the ban on CFCs, the clearest example of this in the contemporary context is the case of climate change. Even if we were to convince the vast majority of citizens that human activities are causing climate change and that, if not addressed on time, climate change is likely to have profound and irreversible consequences on the planet and on humans in particular, the challenge of climate change cannot be met solely (or, even, mainly) through individual decisions (such as the consumption choices of individuals). Climate change requires large-scale social and structural responses that can only be brought about through public policy and international agreements (as in the case of the ban on CFCs). Similar considerations apply even in cases that seem to be based mostly on individual choices, such as the case of vaccines. Since vaccines are not completely effective at the individual level, their effectiveness is partly due to herd immunity. So, even what seems to be a personal choice (whether to have one's children vaccinated) has social ramifications—vaccinating one's kids in a society in which most people do not vaccinate them is only going to provide partial protection against the diseases the vaccine is meant to protect against. Finally, even cases that seem to involve only individual choices and that have significantly less dramatic social consequences might, on closer inspection, require social solutions. Consider, for example, the case of flossing discussed by Oreskes (2019, 118–127). Oreskes notes that the best evidence so far of the effectiveness of flossing is from studies in which children received flossing by a dental hygienist in school. While this might be just an artefact of study design,⁴² it might also be that flossing is only effective (or is significantly more effective) when performed by a dental hygienist. If this is the case, then, if we want to follow the evidence where it leads us, then we should conclude that the

⁴¹ If this is true, then it might be that, contrary to what I suggested in n.5 above, all forms of distrust in science are ultimately socially harmful, as they are likely to affect negatively the resources that society is willing to invest in the education of its citizens or undermine the quality of that education (as illustrated, for example, by the movement to teach creationism in schools).

⁴² For example, it is possible that studies that are based on people self-reporting their flossing habits systematically overestimate how much people actually floss.

question is not so much how regularly individuals floss at home, but whether people should have more access to professional flossing. Oreskes implicitly concedes this when she half-jokingly suggests that flossing bars might be part of the solution to the health issues that result from poor dental hygiene. However, the truth is that, if the evidence does indeed suggest that the benefits of flossing are only limited to professional flossing, then what is required is not simply a change in individual flossing habits, but changes to the way we floss as a society. The point is that even cases that seem to involve only individual choices that appear to have no social ramifications (such as that of flossing one's teeth), the best solution to the problem might be one that requires new social norms, policies, institutions, or services.⁴³

According to the social approach, public trust in science is primarily a form of social trust that is exemplified by an efficient division of epistemic labor between scientists and non-scientists and that depends on the state of the socio-epistemic infrastructure of society. While the individual and semi-social approach focus almost exclusively on interventions at the individual level to solve the problem of harmful distrust, the social approach identifies at least four distinct (though sometimes overlapping) levels of possible intervention. The first is *the individual level*. For example, as we have seen, the social approach agrees with the individual and the semi-social approach that efforts should be made to improve the scientific literacy of ordinary citizens through high-quality public education. The second level is *the social level*. For example, as we have seen, the social approach advocates a greater inclusion of stakeholders (and especially from stakeholders from marginalized groups) in decisions about the direction of publicly funded scientific research. The third level is at *the interface between science and society*. For example, the social approach advocates for improving the quality of science communication, possibly by funding public news sources devoted to providing high-quality science journalism [as opposed to the sort of sensationalistic science journalism that is often provided by private media sources or a journalism that mistakenly pursues the journalistic norm of balance at all costs (see, e.g., Gerken, 2020)]. The fourth level is that of *scientific communities*. For example, society could protect science from the undue influence of economic interests by effectively socializing scientific research while shielding it from the interference of political interests by letting scientists and a diversity of stakeholders be primarily responsible for the allocation of the public funds.

⁴³ I should note that Anderson openly acknowledges the influence of social factors on the beliefs of individual citizens and, in particular, of some of the factors I mentioned in this section. However, in spite of this, Anderson's account is still not a (thoroughly) social account of public trust in science, as it still understands public trust in science as primarily based on a trust relationship between individual citizens and the scientific community and it regards social factors as interfering with the otherwise voluntary and rational epistemic choices of individual citizens.

6 Conclusion

In this paper, I distinguished three general approaches to public trust in science. I argued that the first two approaches, the individual approach and the semi-social approach, conceive of public trust in science as primarily based on an individual trust relationship and that, as a result, conceive of the solution to the problem of harmful distrust as trying to attempt persuade individual citizens that science is trustworthy. I argue that both approaches face two general problems, which I called the problem of overridealizing science and that of overburdening citizens. I argued that, to avoid these problems, we should adopt a (thoroughly) social approach to public trust in science instead. On this approach, public trust in science is primarily based on a social trust relationship between society and science. I also argued that a society that fully trust science is one that trusts it implicitly by adopting an efficient division of epistemic labor. The problem of harmful distrust arises from a poor socio-epistemic infrastructure that undermines the efficient division of epistemic labor. To solve the problem of harmful distrust we need to try to improve the state of society's socio-epistemic infrastructure, rather than try to persuade citizens to trust science. While this includes interventions at the individual level, these interventions need to be supplemented with interventions at the level of social groups, of the interface between science and society, and at the level of improving the (actual and perceived) trustworthiness of science.

The outline of the social approach that I provided in this is admittedly just a sketch and many of the details will need to be worked out elsewhere (or are already being worked out). However, I hope that, in spite of its shortcomings, the outline I provided here contributes to laying the groundwork for a unified social approach to public trust in science, an approach that weaves together a number of strands of research on the social epistemology of science broadly construed, that stimulates more work in that direction, and that, hopefully, promotes the adoption of a more constructive and effective approach to the problem of harmful distrust.

Acknowledgements I would like to thank Gianluca Andresani, Joe Roussos, Katherine Furman, Rachel Herdy, Sushruth Ravish, and Tony Ward, and two anonymous referees for this journal for their helpful comments on previous drafts of this paper. This paper draws on research supported by the Social Sciences and Humanities Research Council of Canada through Insight Development Grant 311412.

Funding Social Science and Humanities Research Council of Canada Insight Development Grant (311412).

Declarations

Conflict of interest The author has no affiliations with or involvement in any organization or entity with interest in the subject matter or materials discussed in this manuscript.

Ethical approval Not applicable.

Informed consent Not applicable.

References

- Almassi, B. (2017). Toxic funding? Conflicts of interest and their epistemological significance. *Journal of Applied Philosophy*, 34(2), 206–220. <https://doi.org/10.1111/japp.12180>
- Anderson, E. (2011). Democracy, public policy, and lay assessments of scientific testimony. *Episteme*, 8(02), 144–164. <https://doi.org/10.3366/epi.2011.0013>
- Angell, M. (2005). *The truth about the drug companies: How they deceive us and what to do about it* (1st ed.). Random House.
- Asch, S. E. (1951). Effects of group pressure upon the modification and distortion of judgments. In H. Guetzkow (Ed.), *Groups, leadership and men: research in human relations* (pp. 222–236). Carnegie Press.
- Baier, A. (1986). Trust and antitrust. *Ethics*, 96(2), 231–260.
- Biddle, J. (2007). Lessons from the Viiox debacle: What the privatization of science can teach us about social epistemology. *Social Epistemology*, 21(1), 21–39. <https://doi.org/10.1080/02691720601125472>
- Boullicault, M., & Andrew Schroeder, S. (2021). Public trust in science: Exploring the idiosyncrasy-free ideal. In K. Vallier & M. Weber (Eds.), *Social trust* (pp. 102–121). Routledge.
- Boulware, L. E., Cooper, L. A., Ratner, L. E., LaVeist, T. A., & Powe, N. R. (2003). Race and Trust in the Health Care System. *Public Health Reports (Washington, D.c.: 1974)*, 118(4), 358–365. <https://doi.org/10.1093/phr/118.4.358>
- Brandt, A. M. (1978). Racism and research: The case of the Tuskegee Syphilis Study. *The Hastings Center Report*, 8(6), 21–29. <https://doi.org/10.2307/3561468>
- Brown, J. R. (2017). Socializing medical research. In K. C. Elliott & D. Steel (Eds.), *Current controversies in values and science. Current controversies in philosophy* (pp. 147–160). Routledge.
- Cook, J., Nuccitelli, D., Green, S. A., Richardson, M., Winkler, B., Painting, R., Way, R., Jacobs, P., & Skuce, A. (2013). Quantifying the consensus on anthropogenic global warming in the scientific literature. *Environmental Research Letters*, 8(2), 024024. <https://doi.org/10.1088/1748-9326/8/2/024024>
- Deer, B. (2006). *Huge sums paid to Andrew Wakefield*. The Sunday Times, December 31, 2006.
- Douglas, H. (2009). *Science, policy, and the value-free ideal*. University of Pittsburgh Press.
- Elliott, K. C. (2014). Financial conflicts of interest and criteria for research credibility. *Erkenntnis*, 79(S5), 917–937. <https://doi.org/10.1007/s10670-013-9536-2>
- Elliott, K. C. (2017). *A Tapestry of Values: An Introduction to Values in Science*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780190260804.001.0001/acprof-9780190260804>
- Fine, C. (2011). *Delusions of gender: How our minds society and neurosexism create difference* (Reprint). W. W. Norton.
- France, D. (2016). *How to survive a plague: The story of how activists and scientists tamed AIDS*. Vintage. <https://www.penguinrandomhouse.com/books/209900/how-to-survive-a-plague-by-david-france/>
- Friedman, P. J. (2002). The impact of conflict of interest on trust in science. *Science and Engineering Ethics*, 8, 413–420.
- Funk, C., & Hefferon, M. (2019). U.S. public views on climate and energy. *Pew Research Center* (blog). November 25, 2019. <https://www.pewresearch.org/science/2019/11/25/u-s-public-views-on-climate-and-energy/>
- Funk, C., Hefferon, M., Kennedy, B., & Johnson, C. (2019). Trust and mistrust in Americans' views of scientific experts. *Pew Research Center Science & Society* (blog). August 2, 2019, from <https://www.pewresearch.org/science/2019/08/02/trust-and-mistrust-in-americans-views-of-scientific-experts/>
- Gawande, A. (2016). The mistrust of science | The New Yorker. *The New Yorker*, June 10, 2016. <https://www.newyorker.com/news/news-desk/the-mistrust-of-science>
- Gerken, M. (2020). How to balance balanced reporting and reliable reporting. *Philosophical Studies*, 177(10), 3117–3142. <https://doi.org/10.1007/s11098-019-01362-5>
- Germani, F., & Biller-Andorno, N. (2021). The anti-vaccination infodemic on social media: A behavioral analysis. *PLoS ONE*, 16(3), e0247642. <https://doi.org/10.1371/journal.pone.0247642>
- Goldacre, B. (2012). *Bad pharma: How drug companies mislead doctors and harm patients*. Fourth Estate.
- Goldenberg, M. J. (2016). Public misunderstanding of science? Reframing the problem of vaccine hesitancy. *Perspectives on Science*, 24(5), 552–581. https://doi.org/10.1162/POSC_empha00223

- Goldenberg, M. J. (2021). *Vaccine hesitancy: Public trust, expertise, and the war on science*. University of Pittsburgh Press.
- Goldman, A. I. (2001). Experts: Which ones should you trust? *Philosophy and Phenomenological Research*, 63(1), 85–110. <https://doi.org/10.1111/j.1933-1592.2001.tb00093.x>
- Gould, S. J. (2006). *The mismeasure of man* (Revised and Expanded). W. W. Norton.
- Grasswick, H. E. (2010). Scientific and lay communities: Earning epistemic trust through knowledge sharing. *Synthese*, 177(3), 387–409. <https://doi.org/10.1007/s11229-010-9789-0>
- Grasswick, H. E. (2014). Climate change science and responsible trust: A situated approach. *Hypatia*, 29(3), 541–557. <https://doi.org/10.1111/hypa.12090>
- Guerrero, A. A. (2016). Living with ignorance in a world of experts. In R. Peels (Ed.), *Perspectives on ignorance from moral and social philosophy* (pp. 156–185). Routledge. <https://doi.org/10.4324/9781315671246-15>
- Hardin, R. (2006). *Trust*. Polity Press.
- Hawley, K. (2014). Trust, distrust and commitment. *Noûs*, 48(1), 1–20. <https://doi.org/10.1111/nous.12000>
- Heesen, R., & Bright, L. K. (2021). Is peer review a good idea? *British Journal for the Philosophy of Science*, 72 (3), 635–663. <https://doi.org/10.1093/bjps/axz029>.
- Hoyt, M. A., Rubin, L. R., Nemeroff, C. J., Lee, J., Huebner, D. M., & Proeschold-Bell, R. J. (2012). HIV/AIDS-related institutional mistrust among multiethnic men who have sex with men: Effects on HIV testing and risk behaviors. *Health Psychology*, 31(3), 269–277. <https://doi.org/10.1037/a0025953>
- Hull, D. L. (1988). *Science as a process: An evolutionary account of the social and conceptual development of science*. University of Chicago Press.
- Jewett, A. (2020). *Science under fire*. Harvard 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.
- Kahan, D. M. (2013). Ideology, motivated reasoning, and cognitive reflection. *Judgment and Decision Making*, 8(4), 407–424.
- Kahan, D. M. (2017). The “gateway belief” illusion: Reanalyzing the results of a scientific-consensus messaging study. *Journal of Science Communication*, 16(5), 3. <https://doi.org/10.22323/2.16050203>
- Kahan, D. M., Jenkins-Smith, H., & Braman, D. (2011). Cultural cognition of scientific consensus. *Journal of Risk Research*, 14(2), 147–174. <https://doi.org/10.1080/13669877.2010.511246>
- Kappel, K. (2014). Believing on trust. *Synthese*, 191(9), 2009–2028. <https://doi.org/10.1007/s11229-013-0376-z>
- Keren, A. (2019). Trust, preemption, and knowledge. In Dormandy, K. (Ed.), *Trust in Epistemology* (pp. 114–135), Routledge.
- Kitcher, P. (1993). *The advancement of science: Science without legend, objectivity without illusions*. Oxford University Press.
- Konstantinou, P., Georgiou, K., Kumar, N., Kyprianidou, M., Nicolaidis, C., Karekla, M., & Kassianos, A. P. (2021). Transmission of vaccination attitudes and uptake based on social contagion theory: A scoping review. *Vaccines*, 9(6), 607. <https://doi.org/10.3390/vaccines9060607>
- Krosnick, J. A. (2019). Comments on the Present and Future of Science, Inspired by Naomi Oreskes. In N. Oreskes (Ed.), *Why trust science? The University Center for human values series* (pp. 202–211). Princeton University Press.
- Lloyd, E. A. (2005). *The case of the female orgasm: Bias in the science of evolution*. Harvard University Press.
- Longino, H. E. (1990). *Science as social knowledge: Values and objectivity in scientific inquiry*. Princeton University Press.
- de Melo-Martín, I., & Intemann, K. (2018). *The fight against doubt: How to bridge the gap between scientists and the public*. Oxford University Press.
- Moore, A. (2017). *Critical elitism: Deliberation, democracy, and the problem of expertise*. Cambridge University Press.
- Nguyen, C. T. (2020). Echo chambers and epistemic bubbles. *Episteme*, 17, 141–161. <https://doi.org/10.1017/epi.2018.32>
- Nguyen, C. T. (forthcoming). Trust as an unquestioning attitude. *Oxford Studies in Epistemology*.
- Nichols, T. (2017). *The death of expertise: The campaign against established knowledge and why it matters*. Oxford University Press.

- Nyhan, B., & Reifler, J. (2015). Does correcting myths about the flu vaccine work? An experimental evaluation of the effects of corrective information. *Vaccine*, 33(3), 459–464. <https://doi.org/10.1016/j.vaccine.2014.11.017>
- O'Connor, C., & Weatherall, J. O. (2019). *The misinformation age: How false beliefs spread*. Yale University Press.
- OECD. (2015). *OECD science, technology and industry scoreboard 2015: Innovation for growth and society*. Text. OECD Publishing. https://www.oecd-ilibrary.org/science-and-technology/oecd-science-technology-and-industry-scoreboard-2015_sti_scoreboard-2015-en
- Oreskes, N. (2004). The scientific consensus on climate change. *Science*, 306(5702), 1686–1686. <https://doi.org/10.1126/science.1103618>
- Oreskes, N. (2019). *Why trust science?* Princeton University Press.
- Oreskes, N., & Conway, E. M. (2010). *Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming*. Bloomsbury Press.
- Parikh, S. (2021). *Why we must rebuild trust in science*. PEW Charitable Trust. February 9, 2021. <https://pew.org/2MoiYIW>
- Romero, F. (2017). Novelty versus replicability: Virtues and vices in the reward system of science. *Philosophy of Science*, 84(5), 1031–1043. <https://doi.org/10.1086/694005>
- Romero, F. (2019). Philosophy of science and the replicability crisis. *Philosophy Compass*, 14(11), 1–14. <https://doi.org/10.1111/phc3.12633>
- Romero, F. (2020). The division of replication labor. *Philosophy of Science*, 58(5), 104–1025. <https://doi.org/10.1086/710625>
- Rudner, R. (1953). The scientist qua scientist makes value judgments. *Philosophy of Science*, 20(1), 1–6.
- Salvanto, A., De Pinto, J., Backus, F., & Khanna, K. (2021). *Many would get COVID-19 vaccine, but hesitancy remains*. CBS News. February 26, 2021. <https://www.cbsnews.com/news/covid-vaccine-hesitancy-opinion-poll/>
- Schman, N. (2001). Epistemology resuscitated: objectivity as trustworthiness. In N. Tuana & S. Morgen (Eds.), *Engendering rationalities. SUNY series in gender theory* (pp. 23–52). SUNY Press.
- Schroeder, A. (2021). Democratic values: A better foundation for public trust in science. *British Journal for the Philosophy of Science*, 72 (2), 545–562. <https://doi.org/10.1093/bjps/axz023>
- Sismondo, S. (2008). Pharmaceutical Company funding and its consequences: A qualitative systematic review. *Contemporary Clinical Trials*, 29(2), 109–113. <https://doi.org/10.1016/j.cct.2007.08.001>
- Solomon, M. (2001). *Social empiricism*. MIT Press.
- Tsipursky, G. (2018). *(Dis)Trust in science*. Scientific American Blog Network. July 5, 2018. <https://blogs.scientificamerican.com/observations/dis-trust-in-science/>
- Tucker, W. H. (1994). *The science and politics of racial research*. University Illinois Press.
- van der Linden, S., Leiserowitz, A., & Maibach, E. (2018). Scientific agreement can neutralize politicization of facts. *Nature Human Behaviour*, 2(1), 2–3. <https://doi.org/10.1038/s41562-017-0259-2>
- van der Linden, S., Leiserowitz, A. A., Feinberg, G. D., & Maibach, E. W. (2015). The scientific consensus on climate change as a gateway belief: Experimental evidence. *PLoS ONE*, 10(2), e0118489. <https://doi.org/10.1371/journal.pone.0118489>
- Weatherall, J. O., O'Connor, C., & Bruner, J. P. (2020). How to beat science and influence people: policymakers and propaganda in epistemic networks. *The British Journal for the Philosophy of Science*, 71 (4), 1157–1186. <https://doi.org/10.1093/bjps/axy062>
- Wellcome Global Monitor. (2018). *How does the world feel about science and health?* Wellcome Trust.

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