



When should one be open-minded?

Hein Duijf¹

Accepted: 16 March 2024 / Published online: 10 May 2024
© The Author(s) 2024

Abstract

It is widely believed among philosophers and educated people that it is virtuous to be open-minded. Instead of thinking of open-mindedness as universally or unconditionally epistemically valuable, I argue that it is vital to explicate the conditions that must obtain if open-mindedness is to be epistemically valuable. This paper critically evaluates open-mindedness given certain realistic cognitive limitations. I present and analyse a simple mathematical model of open-minded decision-making that incorporates these limitations. The results are mixed. The bad news is that the circumstances where open-mindedness is epistemically valuable may be more restricted than previously thought especially if individuals are incapable of properly evaluating the trustworthiness of sources or the content of received information. The good news is that, if individuals are mildly competent in evaluating the trustworthiness of sources and the content of received information, then there are many circumstances where open-mindedness is epistemically valuable.

Keywords Open-mindedness · Bounded cognition · Epistemic value · Virtue · Truth-conduciveness

1 Introduction

It is widely believed among philosophers and educated people that it is virtuous to be open-minded (as opposed to being close-minded or narrow-minded).¹ Although the universal validity of the belief has been questioned (Carter & Gordon, 2014; Levy, 2006; Madison, 2019), this claim is often thought of as a platitude that requires no

¹ On close-mindedness, Cassam (2019, 38) writes: “it is reasonable to conclude that closedmindedness systematically obstructs knowledge both in virtue of the nature of closed-mindedness and the nature of knowledge.”

✉ Hein Duijf
duijf.hein@gmail.com

¹ Munich Center for Mathematical Philosophy, LMU Munich, Munich, Germany

defence. Indeed, it seems intuitive and compelling that conscientious agents ought to be adequately informed of all sides of a debate before they can form justified beliefs on the topic of debate.

This paper studies the epistemic value of open-mindedness. It seems plausible that intellectual virtues need to be reliable at generating true beliefs (and avoiding false ones). That is, if it turns out that an intellectual virtue predominantly produces false beliefs, then it is questionable whether it should qualify as an intellectual virtue (as opposed to an intellectual vice).² To be sure, we could say that open-mindedness is a virtue in a broad sense where it constitutes fairness and respect for other agents and their perspectives. However, it would be odd to call open-mindedness an *intellectual* virtue if it does not reliably contribute to achieving basic epistemic goals.³

Does open-mindedness reliably generate true beliefs (and avoid false ones)? It seems fairly uncontroversial that open-mindedness is not *necessarily* truth-conducive (Carter & Gordon, 2014; Levy, 2006; Madison, 2019) (*pace* (Kwong, 2017)). Arguments typically take the form of counterexamples that demonstrate the contingent connection between open-mindedness and truth. This paper goes beyond these case-based considerations and seeks to explicate general conditions under which open-mindedness is truth-conducive and those where it is not. Why is this important? Of course, we need not think that open-mindedness is a vice if it produces bad epistemic effects in extreme, untypical, or far-fetched scenarios; it might still make sense to call open-mindedness a virtue if open-mindedness typically or normally helps us achieve basic epistemic goods. On the flip side, if open-mindedness produces good epistemic effects in unrealistically optimistic scenarios, then that also need not imply that open-mindedness is an intellectual virtue. After all, if it turns out that open-mindedness only produces good epistemic effects (and avoids bad ones) in very limited circumstances, then this would put pressure on the idea that open-mindedness should qualify as a virtue. I thus propose going beyond mere case-based considerations and, instead, characterize the general conditions under which open-mindedness produces good epistemic effects (and avoids bad ones).

This paper focuses on investigating the conditions under which open-mindedness will be epistemically valuable. One of the novel features of my analysis is that it starts with *three realistic assumptions*: (1) people are not perfectly competent in forming their own opinions, (2) people are not perfectly capable of evaluating the

² Following Battaly's (2018b, 2020), it is common to distinguish between effects-virtues, responsibility-virtues and personalist-virtues. Effect virtues are cognitive dispositions that generate good epistemic effects (or avoid generating bad epistemic effects); responsibility virtues are character traits for which the agent is responsible (praiseworthy or, at least, not blameworthy); and personalist virtues are character traits constituted by epistemically good motives and values (for which the agent need not be responsible). Since I focus on whether open-mindedness reliably generates epistemic goods, my analysis concerns the question of whether and when open-mindedness is an effect virtue.

Nevertheless, some virtue-responsibilists think that intellectual virtues need to also be reliable. For example, Zagzebski (1996, 185) writes: "Thus, not only is an open-minded person motivated to consider the ideas of others without prejudice, including those that conflict with her own, and is reliably successful in doing so, but her reliable success in being open-minded is truth conducive. Therefore, the beliefs she forms out of open-mindedness are truth conducive."

³ In contrast, Song (2018) argues that open-mindedness is a *moral* virtue.

claims put forward by others, and (3) people can only give consideration to a limited number of arguments or opinions. Another central feature is that it goes beyond simplistic examples and represents the agents' reliability and evaluative capacities using *gradual scales*. Is open-mindedness epistemically reliable when these realistic constraints are taken seriously? To emphasize, I believe it is important to study the exact conditions under which open-mindedness would be epistemically beneficial—as opposed to investigating whether open-mindedness is epistemically beneficial *tout court*. To do so, I will put forward a simple mathematical model that includes, among other things, one's capacity to evaluate the source and the content of received information. Then, I will use this model to explicate the general conditions where open-mindedness is epistemically valuable and those where it is not.

My analysis brings good news and bad news. The bad news may be more alarming than previously thought: open-mindedness only leads to epistemic benefits in restricted circumstances. Open-mindedness may produce bad epistemic effects for reasonably competent agents and in reasonably friendly epistemic environments when agents are incapable of properly evaluating the trustworthiness of the sources or the content of received information. This result may be disheartening in light of abundant psychological research showing that: “Illusory truth can persist months after first exposure, regardless of cognitive ability and despite contradictory advice from an accurate source or accurate prior knowledge” (Ecker et al., 2022). The circumstances where open-mindedness increases one's epistemic reliability may thus be more restricted than philosophers and educated people previously recognized.

The good news is that there are many circumstances where open-mindedness is epistemically valuable *if (and only if)* individuals are mildly competent in evaluating the trustworthiness of sources *and* the content of received information. This strongly suggests that open-mindedness will be epistemically fruitful if the participants are not only capable of identifying trustworthy sources but also capable of evaluating the truth-conduciveness of the received information.

Before moving on, I would like to end this introduction by positioning my study in relation to other important works on open-mindedness. While my analysis focuses on the truth-conduciveness of open-mindedness, there is another important debate on the tension between open-mindedness and knowledge (see, among others, Harman, 1973; Sharon & Spectre, 2010; Baumann, 2013; Fantl, 2018) or strong belief (see, for instance, Adler, 2004; Riggs, 2010). The tension can be illustrated by the dogmatism paradox (introduced by Kripke (see Kripke, 2011, chap. 2)) which roughly says that knowing that X (resp., strongly believe X) entails that all counterevidence is misleading and, therefore, that one should (or is permitted to) ignore all counterevidence. In short: knowledge and strong belief seem to entail or licence dogmatism. Since these debates typically focus not on truth-conduciveness but on valid principles for knowledge or strong belief (such as the closure principle), my analysis of the truth-conduciveness of open-mindedness seems orthogonal to the dogmatism paradox on knowledge and strong belief.

The paper is organized as follows. In Sect. 2, I elaborate on the background literature on open-mindedness, social epistemology and the abovementioned three realistic constraints. In Sect. 3, I introduce the mathematical model and discuss a stylized example. In Sect. 4, I demonstrate how the model operationalizes epistemic

accuracy. In Sect. 5, I present some analytical results that characterize the circumstances where open-mindedness could, in principle, be epistemically beneficial and those where it would be epistemically harmful. I go on to show that there are circumstances where open-mindedness cannot even in principle be epistemically valuable—even though agents are reasonably competent. In Sect. 6, I take seriously the realistic assumption that people can only consider a limited number of arguments or views, which implies that the conditions under which open-mindedness is epistemically valuable are more constrained than those given by the analytical results of Sect. 5. The upshot is the bad news that, when agents only practise source evaluation, there are many circumstances where open-mindedness is not epistemically beneficial—even if agents are reasonably competent and are reasonably good at source evaluation. In Sect. 7, I update the model to include not only source evaluation but also content evaluation. This leads to some bad news: when agents only practise content evaluation, there are several conditions where open-mindedness is not epistemically valuable—even if agents are reasonably competent and are reasonably good at content evaluation. We also derive the good news: if people are mildly competent in evaluating the content of received information and the trustworthiness of the sources, then open-mindedness will most likely be epistemically valuable. I summarize the results and provide an outlook in the concluding section.

2 Open mind and cognitive limits

Social epistemologists have recognized that much of our knowledge comes from others and even scientific knowledge is largely based on epistemic trust (Coady, 1992; Hardwig, 1991). Although this leads to philosophical problems relating to testimony and knowledge (Adler, 2017), almost everybody agrees that ordinary people and also scientists epistemically depend on others in their practices of knowledge gathering, production, and evaluation (Hardwig, 1985). In the natural sciences, research is often pursued by a team of scientists, as opposed to individual scientists. But epistemic dependence should also be clear in other disciplines, like philosophy, given the abundant references and footnotes that defer to the authority and expertise of others.

Given the central role of epistemic dependence in our ordinary and scientific knowledge, it becomes clear that it is a central problem to figure out who can be trusted epistemically and which experts one should defer to (Goldman, 2001). This problem is particularly striking in circumstances where several experts disagree. Argumentative practices often depend, at least partly, on epistemic trust in your argumentative partner (Mercier & Sperber, 2017; Dutilh Novaes, 2020). Epistemic trust is thus one of the central topics in social epistemology (Duijf, 2021; Irzik & Kurtulmus, 2019; Wilholt, 2013).

Virtue epistemologists often consider several intellectual virtues including, but not limited to: “inquisitiveness, attentiveness, carefulness and thoroughness in inquiry, fair-mindedness, open-mindedness, and *intellectual* patience, honesty, courage, humility, and rigor” (Baehr, 2011, 3). According to Riggs (2010), open-mindedness often appears at the top of lists of intellectual virtues. These

intellectual virtues are thought to be constitutive of a person's epistemic excellence: Linda Zagzebski concurs that intellectual virtues are

a deep and enduring acquired excellence of a person, involving a characteristic motivation to produce a certain desired end and reliable success in bringing about that end. (Zagzebski, 1996, 137)

Let me proceed with a few definitions of open-mindedness:

An open-minded person is characteristically (a) willing and (within limits) able (b) to transcend a default cognitive standpoint (c) in order to take up or take seriously the merits of (d) a distinct cognitive standpoint. (Baehr, 2011, 153)

[I construe] open-mindedness as *engagement*, that is, a willingness to make room for novel ideas in one's cognitive space and to give them serious consideration. (Kwong, 2016, 71)

What makes one open-minded is the disposition to engage seriously with (relevant) intellectual options. (Battaly, 2018a, 266)

As indicated in the introduction, this paper investigates the conditions under which open-mindedness is epistemically valuable, in the sense of reliably producing true beliefs (and avoiding false ones). Many philosophers (including the aforementioned authors) agree that open-mindedness is a willingness, ability or disposition to give relevant evidence and arguments serious consideration. Levy (2006, 56) concurs: "Evidence gathering is almost the only practical way in which one can put open-mindedness in effect". Open-mindedness includes not only the openness to relevant evidence and arguments but also the disposition to evaluate the relevant information to monitor for possible errors. In the next section, I will present a mathematical model that operationalizes open-mindedness as evidence and argument gathering and evaluation. My central goal is to characterize the conditions under which open-mindedness produces good epistemic effects (and avoids bad ones). By considering some realistic assumptions about people's bounded cognition (see below), I hope to make progress in the study of the epistemic value of open-mindedness.

Before proceeding, it will be helpful to discuss a few simple cases that speak in favour of open-mindedness. In this way, I would like to highlight some of my central assumptions. First, consider an agent who is to form her opinion on whether it will rain next Saturday. She looks at the blue sky and thinks that rain will be unlikely. However, she also decides to look up the weather report, which informs her that it will likely rain. In response to the news report, she revises her belief. The agent's belief is responsive to the report by experts (or institutions) who can be assumed to be both neutral on the issue and more knowledgeable in weather forecasting.

Second, consider another agent who is to form his opinion on whether sugar is bad for one's health. The agent starts by collecting arguments and evidence from a variety of sources. Then, he evaluates the collected arguments and the evidence critically, thereby discerning the good arguments and evidence from the poor.

Finally, he forms his opinion based on his accepted sources. It seems like this situation illustrates that open-mindedness is surely a good way to lead one's epistemic life. After all, by considering and carefully scrutinizing various sources, the agent can arrive at an informed opinion on whether sugar is bad for one's health.

Third, let me also mention a version of the infamous evil demon problem. Imagine that, unbeknownst to the agent, his cognitive processes involved in interpreting and evaluating arguments and evidence are rendered unreliable by a powerful evil demon. If the agent were to form his beliefs open-mindedly, the generated beliefs would often, if not always, turn out to be false. As a result, open-mindedness would not be a reliable way to arrive at true beliefs (and avoid false beliefs) for this agent, so open-mindedness is not epistemically virtuous in these circumstances.

In the first example, there is no uncertainty regarding the competence and reliability of the weather forecasters. Indeed, if the agent can be certain that all sources are competent and reliable, then it seems very plausible that open-mindedness will generate epistemic benefits. In this paper, however, I will assume that the information received from other sources may be unreliable. The unreliability may be due to diverging background beliefs, epistemic principles, cognitive heuristics, interests or values. One type of example is where one's environment possibly includes sources of misinformation (O'Connor & Weatherall 2019). But, it has also been argued that factual claims by scientists at least partly depend on value judgments (Douglas, 2000). Another example is where a health practitioner might be misaligned even though they care about the interests of the patient, because they tried and failed to identify the interests of the patient. The study of open-mindedness thus needs to consider the possibility of unreliable agents.

In the second example, one of the central assumptions seems to be that the agent has all the resources and cognitive capacities to search for a variety of arguments and evidence and to also scrutinize these sources carefully and charitably. After all, if the agent can consider all the available arguments and evidence and carefully scrutinize this information, then it seems highly likely that open-mindedness leads to epistemic benefits. Typically, however, people lack the time to collect various arguments and evidence. This leads people to only consider a couple of pieces of information. Moreover, people may lack the cognitive skills necessary to evaluate the information especially when they are not knowledgeable in that domain. As a result, there is a chance that people's opinions are influenced by false information even if they have carefully, thoroughly, and charitably evaluated the arguments and evidence.

The third example is one where the agent's evaluative skills are so terrible, due to the intervention of the evil demon, that beliefs formed in an open-minded way will surely be unreliable. Although I would submit that agents in these situations should not form their beliefs in an open-minded way, this leaves open the question of the epistemic value of open-mindedness in less hostile (and more mundane) circumstances. The fact that open-mindedness is epistemically harmful in this extreme and untypical case does not necessarily entail that open-mindedness is epistemically harmful in typical cases.

My work goes beyond existing challenges to the truth-conduciveness or epistemic reliability of open-mindedness (Carter & Gordon, 2014; Levy, 2006; Madison, 2019). Levy argues that open-mindedness will lead to irrationality and unjustified beliefs on controversial moral and political questions. Carter and Gordon and Madison argue that the truth conduciveness of open-mindedness is highly conditional, for instance, conditional on the truth of one's existing beliefs and the hospitality of the epistemic environment. These representative analyses of the epistemic value of open-mindedness are typically carried out by considering examples that are simplistic in the sense that they consider one (or few) beliefs and/or fully friendly or fully hostile epistemic environments (e.g., the aforementioned evil demon scenario). My analysis goes beyond these analyses in, at least, three ways: (1) my analysis of open-mindedness goes beyond simplistic examples and incorporates *gradual scales* to represent agents' competences and evaluative capacities, (2) I will develop a mathematical model to study *the general conditions under which* open-mindedness produces good epistemic effects, and those under which it produces bad epistemic effects, and (3) my model incorporates *three realistic constraints* to model bounded cognition.

Let me end this section by discussing the three realistic assumptions. First, some agents may be unreliable, and even reliable agents may not be perfectly competent. Although defining competence is a tricky issue, in the present context it can be taken to represent the agent's disposition to produce evidence or arguments that support the right alternative. The assumption that agents are not perfectly competent then boils down to the assumption that agents do not always succeed in producing evidence and arguments supporting the right alternative. In other words, reliable agents may occasionally support the wrong alternative.

Second, agents are not perfectly capable of evaluating the arguments or claims put forward by others (Mercier, 2016; Mercier & Sperber, 2017; Sperber et al., 2010). Work by Daniel Gilbert and his colleagues even seems to show that people automatically start by accepting communicated information (Gilbert et al., 1990, 1993), but this is not uncontroversial (Hasson et al., 2005). One might think that people should be open-minded and consider all arguments by others and then simply throw away the bad or wrong arguments. Open-minded people would have a greater pool of good or decent arguments and should therefore be able to draw better conclusions. The problem with this thought is that people may not be perfectly able to identify the truth-conducive sources and arguments. The assumption of imperfect evaluative skills is operationalized by assuming that agents may take into consideration information from unreliable agents and also information from reliable agents which supports the wrong alternative.

Third, agents can only take into consideration a limited number of arguments or opinions. Empirical research on human cognition suggests that people have limited capacities for remembering and processing information and typically can process and recall between 3 and 5 meaningful pieces of information (Cowan, 2001, 2010).⁴

⁴ Miller (1956) is often cited as one of the first contributions to the empirical work on short-term memory capacities. He summarizes evidence that people can recall seven chunks of information in short-term memory tasks. Cowan (2001) brings together a wide variety of data suggesting that the smaller capacity limit, between 3 and 5, is real.

So, even if one were open-minded and collected various arguments, it is highly likely that the final decision is only based on a limited number of them. This puts a severe constraint on the number of arguments or opinions that people can take into consideration.

3 A model of the open mind

In this section, I will introduce a mathematical model that operationalizes open-mindedness as evidence and argument gathering and evaluation. For simplicity's sake, the rest of the paper will concern cases where an agent is to decide between two alternatives, depicted by A_0 and A_1 , respectively. It will be helpful to illustrate the components and assumptions of my model by discussing a stylized example. Let us imagine that an agent called Aretha is to make up her mind on a certain question, let us say whether Q is the case. The question might, for instance, be factual (how much money does Germany spend on foreign aid?), about forecasting (what will the weather be like in Indonesia next month?), value-laden (is it in the public interest to decrease inequality?), instrumental (are social welfare schemes an effective way to enhance equality?), or a combination (are bicycles safe?). By convention, let us assume that A_1 is the right alternative if Q is true and A_0 is the right alternative if Q is false instead.⁵ The most natural way to think about these alternatives is that A_1 represents 'believing that Q is the case' and A_0 represents 'believing that Q is not the case'. Depending on the nature of the inquiry, non-epistemic factors such as Aretha's standards of evaluation, interests or values may partially determine which of these alternatives is *the right one* (for Aretha). To give two simple examples, whether bicycles are safe (partially) depends on Aretha's risk attitude and whether it is in the public interest to decrease inequality (partially) depends on Aretha's interests and values regarding inequality.

In line with our earlier assumption, Aretha is part of a diverse community of agents with possibly different background beliefs, epistemic principles, cognitive heuristics, interests, and values. Although it would make sense to adopt a gradual scale representing the degree of overlap in background beliefs, epistemic principles, cognitive heuristics, interests and/or values,⁶ I simplify matters here and assume that the community can be divided into two groups: the *reliable group* consisting of agents with the same cognitive and evaluative makeup as Aretha; and the *unreliable group* consisting of agents with conflicting cognitive and evaluative makeup. Accordingly, Aretha is a member of the reliable group. I will call members of the reliable group *reliable agents* and members of the unreliable group *unreliable agents*.

Aretha's decision-making can be summarized as follows. Aretha can gather evidence by consulting various sources, including the opinions of others but also the evidence and arguments found in academic and other outlets. Before making up her mind, she collects various pieces of evidence and arguments and, then, evaluates

⁵ The following may be helpful for remembering these notational conventions: A_0 is the right alternative if the truth value of Q is 0, and A_1 is the right alternative if the truth value of Q is 1.

⁶ In (Duijf 2021), I introduce such a gradual scale of the degree of interest alignment to study the conditions under which one should trust experts.

their validity and/or truth-conduciveness. Consequently, she will base her final conclusion on the accepted pieces of evidence and arguments, i.e., those that were collected and then evaluated as valid. Aretha's decision process is visualized in Fig. 1. Let us go over each of these steps in turn.

Before Aretha reaches a final decision, she also considers the opinions of others. The available sources can be divided into reliable and unreliable agents. A reliable agent's *individual competence* represents their disposition to produce or share evidence or arguments that *support* the right alternative; an unreliable agent's individual competence represents their disposition to produce information that *opposes* the right alternative. Agents are not assumed to be intellectual geniuses, so these competences are modelled by a probability which can be taken to represent the chance that they produce or share evidence or arguments supporting the right alternative, in the case of reliable agents, or supporting the wrong alternative, in the case of unreliable agents.

In other words, reliable and unreliable agents are both imperfect: reliable agents can occasionally produce evidence or arguments opposing the right alternative, and unreliable agents can sometimes produce evidence and arguments supporting the right alternative. Given our terminology of reliable and unreliable agents, we assume that both competences are above chance level. Hence, a reliable agent is more likely than not to produce evidence and arguments in favour of the right alternative. By analogy, an unreliable agent is more likely than not to produce evidence and arguments in favour of the wrong alternative.

It may be helpful to point out that information that supports the wrong alternative need not be false or deceptive. After all, the accumulation of evidence can be non-monotonic. Consider the following stylized example from the domain of non-monotonic reasoning. Suppose an agent is to make up her mind on whether a given animal, Tweety, is able to fly. First, she gets evidence that Tweety is a bird. From this evidence, she (legitimately) draws the conclusion that Tweety is able to fly. Second, she gets evidence that Tweety is a penguin. As a result, she updates her conclusion that Tweety is not able to fly. The moral of this stylized example is that one might obtain true evidence (i.e., that Tweety is a bird) that leads one to legitimately endorse false conclusions.

It is a feature of my model that it is compatible with different assumptions on what grounds the reliability or unreliability of agents. These assumptions are external to my model, and these grounds can be rational or arational and epistemic or non-epistemic. Let me give some examples. First, the unreliability of an agent might be grounded in different epistemic principles and background beliefs (e.g., Ranalli & Lagewaard, 2022). The difference in epistemic principles and background beliefs could be the basis for the unreliable agent's disposition to be more likely to produce evidence and arguments opposing the right alternative. Second, realistic agents do not always reason perfectly rationally but also employ cognitive heuristics and biases (e.g., Stanovich, 2003; Tversky & Kahneman, 1974). These cognitive heuristics and biases may ground an agent's unreliability. Third, non-epistemic values, such as evaluative standards and inductive risks, play a role in epistemology and science (e.g., Douglas, 2000). For example, agents with different risk attitudes might justifiably also differ in their assessment of whether bicycles are safe and, hence, be

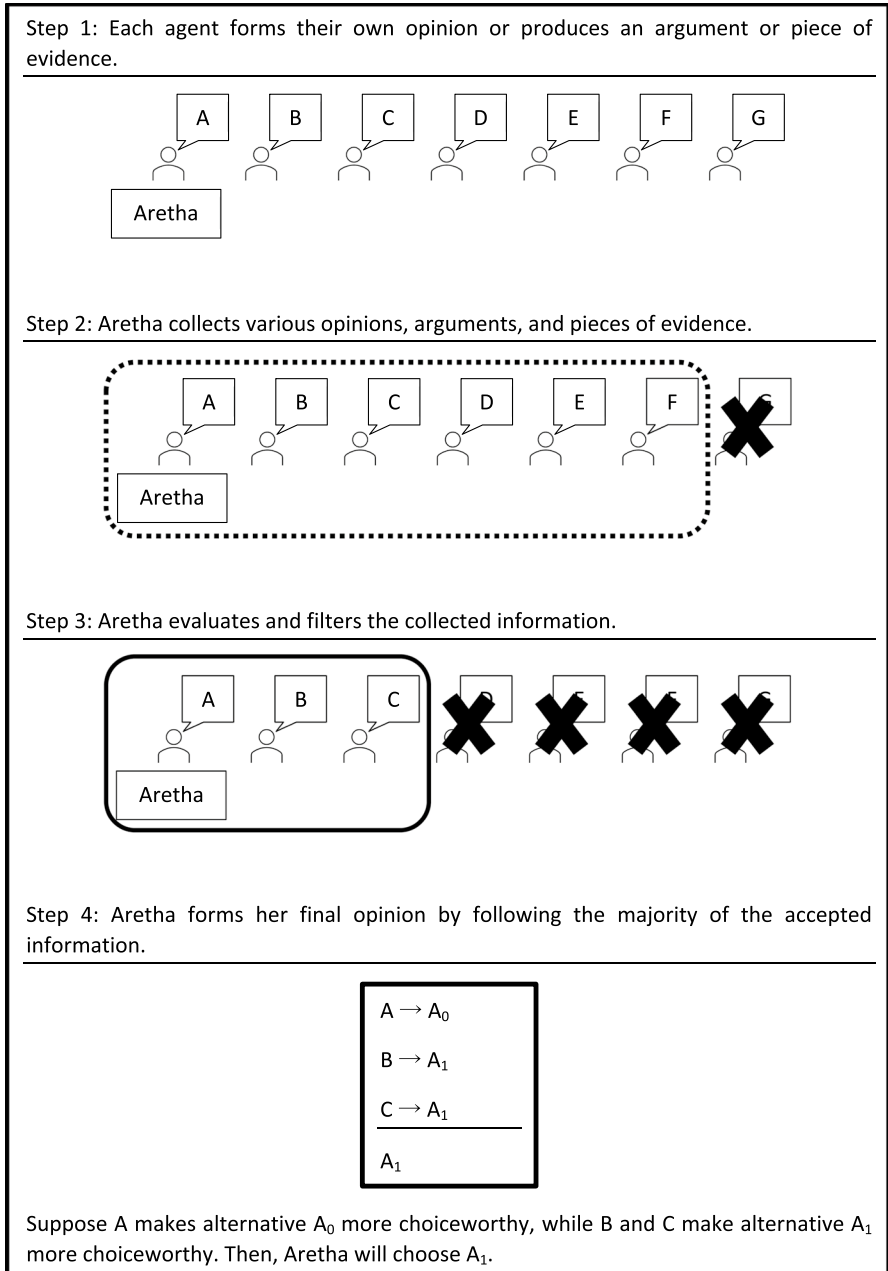


Fig. 1 A visualization of Aretha’s decision process

unreliable (for Aretha). Fourth, the interests or values of the agents may determine the (agent-relative) right alternative (e.g., de Ridder, 2021; Duijf, 2021), especially in the political domain (for example, concerning taxation and welfare programs).

Conflicting interests may ground an agent's unreliability. Note that conflicting interests are only contingently connected to malice: unreliable agents might have Aretha's best interests in mind and still have conflicting interests because they try and fail to identify Aretha's interests. Nonetheless, conflicting interests can ground the fact that unreliable agents are more likely to find and share information opposing the alternative that is in Aretha's best interest. Finally, some unreliable agents might be malicious (or, at least, not be purely epistemically motivated) and consciously produce misinformation (O'Connor & Weatherall 2019).

Let us proceed with the simplifying assumption that all reliable sources are equally competent and let p_R denote their competence. Similarly, I assume that all unreliable sources are equally competent and let p_U represent their competence.⁷ To incorporate the possibility that the competences of the reliable group and the unreliable group differ, each group is assigned a distinct competence. We might, for instance, model the differences between elites and masses, and adopt the assumption that the elites are more competent than the masses (for instance, due to their extensive resources and higher levels of education) (Spiekermann, 2020). These extra assumptions on what grounds the distinction in competences are external to my model; for my purposes, it suffices that my model does not preclude the possibility that the competences of these groups may differ.

My model also includes the assumption that the probabilities associated with the agents' competences are unconditionally independent. This means that the probability that one agent produces information supporting the right alternative is (unconditionally) independent of whether another agent does so, and vice versa.⁸ Although there might be cases where the agents' competences are not probabilistically independent (for example, if they use the same evidential sources or rely on the same opinion leaders (Dietrich & Spiekermann, 2023, sec. 2)), the assumption of unconditional independence is optimistic, leads to a diverse reliable group, and enhances tractability. Let me explain. First, this is an optimistic assumption. Consider the contrasting case where two dependent reliable agents tell Aretha that it will probably rain tomorrow but their beliefs are both based on the same weather report. In this case, the report by the second reliable agent does not provide extra epistemic goods. By contrast, if the reliable agents were unconditionally independent, the second reliable agent would provide epistemically valuable information. Second, although reliable agents are characterized by having a similar cognitive and evaluative make-up, the independence assumption implies that their individual judgments are diverse. Dietrich and Spiekermann (2023, sec. 3.3) concur: "Diversity manifests itself as probabilistic independence between individual judgments". Third, the independence assumption enhances tractability by avoiding non-trivial modelling choices. One

⁷ Competence homogeneity does not rule out cognitive diversity (Dietrich and Spiekermann 2023, sec. 3).

⁸ This is different from the assumption that the agents' *beliefs* are probabilistically independent. If all agents are more competent than chance level, then the fact that one reliable agent believes that A is not probabilistically independent of the fact that another reliable agent believes that A. After all, under this assumption, the fact that one reliable agent believes that A is evidence for the fact that A is the case which, in turn, is conducive to another reliable agent believing A.

alternative would be to specify the exact nature (and structure) of the dependence between the individual judgments of the agents (for example, in the form of a causal network). Another alternative would be to represent the dependencies using correlation coefficients. Both of these options seem to involve highly non-trivial or context-sensitive modelling choices: the structural dependence and correlation coefficients plausibly vary on a case-by-case basis. I opt for sidestepping these non-trivial issues and enhancing tractability by endorsing the independence assumption.

After gathering evidence and/or arguments, Aretha will evaluate the information and arguments as depicted by the third step in Fig. 1. At this stage, I will assume that Aretha only evaluates the sources for their epistemic trustworthiness and discards any information from sources she deems epistemically untrustworthy; in Sect. 7, I will expand the model by allowing for the possibility that Aretha also evaluates the content of the information and arguments. In light of source evaluation, I assume that agents are not perfect in assessing others' reliability, and therefore are imperfect in determining whether a given agent is reliable. Instead of elaborating on the process of information gathering and evaluation, my model focuses on the outcome of such a process. The outcome of the process can be represented by the number of pieces of accepted information and the likelihood that any accepted piece of information comes from a reliable source. To incorporate the fact that people are imperfect in evaluating the trustworthiness of sources, my model includes a variable called the *source evaluative capacity* (denoted by probability p_{ES}) which represents the likelihood that any accepted piece of evidence and/or argument comes from a reliable source. I must quickly add that although 'evaluative capacity' has an individual and cognitive connotation, a person's evaluative skills may depend on factors that are outside the person's control or responsibility, such as the informational ecology. In particular, it would be interesting to investigate whether social factors may unjustly influence evaluative skills. In any case, the other agents that influence Aretha after this evaluation phase are called her *trustees*.

Note that source evaluation can be viewed as an effort to find a diverse set of reliable sources. After all, the assumption of unconditional independence implies that the reliable group is diverse.⁹

Since Aretha is assumed to base her final decision partly on the testimony of others, the source evaluative capacity operationalizes how likely it is that she is informed by the opinions of (a diverse set of) reliable sources. For instance, imagine that Aretha bases her decision on her own opinion and those of five other agents. Then, the source evaluative capacity can be used to calculate the probability that, say, at least three out of five trustees are reliable sources. After all, the likelihood of a certain composition is given by a binomial distribution. For example, if the source evaluative capacity equals 80%, then the chance that at least three out of five trustees are reliable sources equals 94%.

I would like to acknowledge that there are other ways to model source evaluative capacities. For instance, one could elaborate on the process of information gathering

⁹ If we were to drop the independence assumption, the principal agent faces two tasks: determining which sources are trustworthy, and determining which sources are (relatively) independent.

and evaluation by assuming that Aretha considers pieces of information one-by-one. For each piece of information, Aretha needs to make up her mind on whether to epistemically trust the source and accept the piece of information. The evaluative capacity could then be thought of as conditional probabilities such as the likelihood that Aretha correctly identifies a reliable source when exposed to information from a reliable source. However, it would then also be important to consider the likelihood that Aretha correctly identifies an unreliable source, when exposed to information from an unreliable source. To determine the composition of her accepted information, one would, furthermore, need to determine the likelihood that Aretha considers pieces of information from reliable versus unreliable sources. It should be clear that this alternative way of modelling source evaluative capacities further complicates the model and brings in several new parameters. However, to keep things simple, I decided to abstract away from the process of information gathering and evaluation and, instead, simply model the outcome of the process by the number of pieces of accepted information and the source evaluative capacity.

Let me add one comment on this more elaborate process before proceeding. Consider the following pair of assumptions: (a) the assumption that the conditional probability to correctly identify reliable sources equals the conditional probability to correctly identify unreliable ones; and (b) the assumption that Aretha is equally likely to consider pieces of information from reliable sources as from unreliable ones. Let us use p_α to denote the conditional probability to correctly identify reliable sources. Then, under these simplifying assumptions, the source evaluative capacity p_{ES} , which represents the likelihood that an accepted source is a reliable one, equals the conditional probability p_α . So, another way to interpret the source evaluative capacity p_{ES} is to think of it as representing the conditional probability that Aretha correctly identifies a reliable source when exposed to information from that source, under the simplifying assumptions stated above.

After gathering and evaluating evidence and arguments, I assume that Aretha forms her final choice based on her own initial opinion and on those communicated by her trustees as depicted in the fourth step in Fig. 1. I propose to operationalize this final decision by assuming that Aretha's final choice follows the majority of those communicated by accepted sources—including her own opinion.¹⁰ One way to interpret this simplifying assumption is to view the agent as putting equal weight on everyone's opinion—none is given special privilege. This is an attitude that is largely in agreement with open-mindedness and the idea that everyone should be treated with equal respect. The topic of judgment transformation has structural similarities to judgment aggregation. Indeed, in light of May's Theorem (May, 1952), the majority rule is the only judgment transformation procedure that satisfies four plausible demands (in a binary choice scenario): universal domain, which says that

¹⁰ This is a common assumption in a large class of agent-based models of information flow and social influence. However, this is not the only option that has been considered. For example, Lehrer and Wagner (1981) assume weighted averaging where the weights represent the respect that a given individual assigns to others. Another example, in bounded-confidence models with discrete opinions, Fortunato (2004) studied an indeterministic rule: the *probability* that the agent picks a given alternative is proportional to the proportion of neighbours that pick that alternative.

all combinations of opinions are admissible; anonymity, which states that all trustees are treated equally; neutrality, which requires that the two options are treated equally; and positive responsiveness, which requires that the final choice is positively influenced by the trustees' opinions.¹¹

It should be clear that this decision-making procedure is different from intellectual gullibility (Fricker, 1994). A person is gullible if they unreflectively and/or uncritically believe others' testimony. In my model, however, the agent will first critically assess the trustworthiness of others before deciding to accept their arguments and/or evidence.

Finally, besides individual competences, reliable and unreliable groups, and source evaluative capacity, I propose to include a parameter for the *degree of open-mindedness* represented by the number of accepted pieces of information that influence Aretha (denoted by n). In other words, if Aretha accepted the opinions of many other agents, she would be considered more open-minded than if she were to consider the opinions of few other agents.

My model represents the decision-making of an agent who is open-minded and critical: the agent gathers information from multiple sources, evaluates the sources and forms her final opinion based on accepted pieces of information. In sum, the mathematical model is specified by four parameters:

1. The number of trustees: n
 - a. The number of accepted sources (i.e., trustees) that influence Aretha's final decision.
2. The competence of reliable agents: p_R
 - a. The probability that a reliable source (including Aretha) will produce or share evidence or arguments supporting the right alternative.
3. The competence of unreliable agents: p_U
 - a. The probability that an unreliable source will produce or share evidence or arguments opposing the right alternative.
4. The source evaluative capacity: p_{ES}
 - a. The probability that a trustee is a reliable agent.

The agent makes up her mind by following the majority of the accepted sources (including herself). In the next section, I will demonstrate how we can calculate the agent's accuracy given this model.

¹¹ It may be interesting to note that in non-binary choice situations, there are impossibility results on judgment transformation that are similar to the impossibility results on judgment aggregation (List 2011).

4 Accuracy of the open mind

As mentioned in the introduction, I focus on whether open-mindedness improves one's accuracy. In other words, for a given scenario, the question is whether higher degrees of open-mindedness would improve the likelihood that the agent identifies the *right alternative*. After all, if it turns out that higher degrees of open-mindedness are associated with low degrees of accuracy in a given scenario, i.e., if they yield more false beliefs and fewer true beliefs, then open-mindedness does not seem to produce epistemic goods in that scenario.

In light of my model, the *individual accuracy* of an agent depends on the aforementioned four model parameters:

1. The number of trustees: n
2. The competence of reliable agents: p_R
3. The competence of unreliable agents: p_U
4. The source evaluative capacity: p_{ES}

Given values for these parameters, the individual accuracy of a given agent can be calculated. Let me give some insight into how this works.

Let me start with a special case. First, suppose a given agent is influenced by n other agents where each of them is reliable.¹² Then, the agent's accuracy is given by the likelihood that the majority of these n reliable agents plus herself support the right alternative. The number of reliable agents that support the right alternative can be thought of as a binomially distributed variable X where $n + 1$ represents the total number of experiments and p_R represents the probability of each experiment (individually) yielding a successful result. In short, $X \sim \text{Binom}(n + 1, p_R)$. Accordingly, the probability that exactly k of these $n + 1$ agents support the right alternative, notation: $P(X = k)$, is given by:

$$P(X = k) = \binom{n + 1}{k} p_R^k (1 - p_R)^{n+1-k}.$$

Since it is assumed that the agent adopts the majority rule, in these circumstances and for even n , the individual accuracy of the agent is given by summing up the probabilities that more than half of these agents support the right alternative.¹³

$$P(X > n/2) = \sum_{k > \frac{n}{2}} P(X = k) = \sum_{k > \frac{n}{2}} \binom{n + 1}{k} p_R^k (1 - p_R)^{n+1-k}.$$

¹² Note that this can be viewed as the special case where $p_{ES} = 1$.

¹³ Since the agent also takes her own opinion in consideration, the formula includes an occurrence of $n + 1$ instead of n . Notice that since n is assumed to be even, the majority rule always produces a decision, i.e., there are no ties.

Let us move on to the general case where a given agent is influenced by n other agents, where each of them is a member of the reliable group with probability p_{ES} . In this general case, the *expected trustee's accuracy*, notation: p_T , is the probability that a trustee supports the right alternative. It is given by the sum of (i) the probability that the trustee is a reliable agent and she, as a reliable agent, successfully produces information supporting the right alternative *plus* (ii) the probability that the trustee is an unreliable agent and she, as an unreliable agent, fails to produce information opposing the right alternative. More formally:

$$p_T := p_{ES} \cdot p_R + (1 - p_{ES})(1 - p_U). \tag{1}$$

Let me refer to this as the *trustee's accuracy* for short.

Of course, if $p_T \leq 0.5$, then open-mindedness will lead to bad epistemic outcomes. Let me use a numerical example to illustrate this possibility for imperfect agents. Consider a situation where the reliable group has a competence of 0.6, the source evaluative capacity equals 0.6, and the unreliable group has a competence of 0.7. In these circumstances, the trustee's accuracy equals 0.48, i.e., $p_T = 0.6 \cdot 0.6 + 0.4 \cdot 0.3 = 0.48$. Hence, in these circumstances, the principal agent should not be open-minded since it leads to epistemic unreliability. Notice that this is an example where open-mindedness is epistemically harmful even though the reliable agents are more competent than chance level and the source evaluative capacity exceeds chance level.

For this general case, the number of trustees that produce information that supports the right alternative can be thought of as a binomially distributed variable Y where n represents the total number of experiments and p_T represents the probability of each experiment (individually) yielding a successful result. In short: $Y \sim \text{Binom}(n, p_T)$. Hence, the probability that exactly k of these n trustees support the right alternative, notation: $P(Y = k)$, is given by:

$$P(Y = k) = \binom{n}{k} p_T^k (1 - p_T)^{n-k}.$$

Accordingly, in this general case, the individual accuracy of an agent can be calculated by considering two events (for even n): (i) she herself and at least half of the trustees support the right alternative, and (ii) she herself fails to support the right alternative, but more than half of the trustees support the right alternative. Hence, more formally, her expected individual accuracy, notation: $p_{EIA}(n, p_R, p_U, p_{ES})$, is given by:

$$\begin{aligned} p_{EIA}(n, p_R, p_U, p_{ES}) &= p_R \cdot P(Y \geq n/2) + (1 - p_R) \cdot P(Y > n/2) \\ &= p_R \cdot \sum_{k \geq \frac{n}{2}} \binom{n}{k} p_T^k (1 - p_T)^{n-k} + (1 - p_R) \cdot \sum_{k > \frac{n}{2}} \binom{n}{k} p_T^k (1 - p_T)^{n-k}. \end{aligned} \tag{2}$$

These equations can hence be used to calculate the expected individual accuracy p_{EIA} on the basis of the four model parameters: n , p_R , p_U , and p_{ES} .

5 Analytical results on the accuracy of the open mind

Before moving to the main analytical results that characterize the conditions under which open-mindedness can, in principle, be epistemically beneficial, I would like to present a few stereotypical scenarios. My goal here is to illustrate some of the analytical findings that will be presented later.

To investigate the benefit or risk of open-mindedness, I propose to compare the agent's individual competence with her expected individual accuracy. The former is the probability that she herself succeeds in identifying the right alternative when left to her own devices, and the latter is her accuracy after she has taken into consideration the opinions of others. The epistemic benefit is given by the difference between the agent's expected individual accuracy and her individual competence. In sum, the epistemic benefit in a given situation (parametrized by n , p_R , p_U , and p_{ES}) is given by the following formula:

$$p_{EIA}(n, p_R, p_U, p_{ES}) - p_R.$$

Let us start with considering the hyper-optimistic case where the source evaluative capacity equals 100%. This means that whenever the agent takes another agent's opinion into consideration then they are guaranteed to be reliable. In other words, the agent only accepts reliable agents and rejects all unreliable agents. Given that reliable agents are not perfectly competent, the opinions of trustees may still fail to support the right alternative. To illustrate the benefit of open-mindedness in this hyper-optimistic case, consider Fig. 2, where the x-axis represents different degrees of open-mindedness and the y-axis represents the epistemic benefit. The figure contains several graphs, where each is associated with a homogeneous community where every agent has the same competence level.

Let me sum up a few observations. First, note that every line is bounded above since the expected individual accuracy cannot exceed 1.0. This entails, for instance, that the line associated with the competence of 90% cannot go beyond 0.10 because a greater epistemic benefit would mean that the agent's expected individual accuracy exceeds 100% (which is conceptually impossible). So, there is more room for individual improvement if the agent's individual competence is lower.

Second, notice that every line is monotonically increasing. This means that, in these hyper-optimistic circumstances, it is always somewhat beneficial to be more open-minded. However, this gain decreases as one becomes more open-minded. For example, the epistemic benefit of increasing the degree of open-mindedness from 1 to 2 is much greater than the benefit of increasing the degree of open-mindedness from 10 to 11. This should not come as a surprise, but it does indicate that there might be limits to the virtue of being open-minded *to the greatest possible extent*. After all, gathering more evidence and considering more arguments takes effort and attention, so if the expected epistemic improvement is tiny, it makes sense to refrain from pursuing the greatest degree of open-mindedness.

Let us proceed with a slightly pessimistic case where the source evaluative capacity equals 30%. To illustrate the benefit of open-mindedness in this slightly pessimistic case, consider Fig. 3, where the x-axis represents different degrees of




Fig. 2 Epistemic benefits of open-mindedness in the hyper-optimistic case where the source evaluative capacity equals 100%

open-mindedness and the y-axis represents the epistemic benefit. The figure contains several graphs, where each is associated with a homogeneous community where every agent has the same competence level.

Notice that, since all graphs take negative values for all degrees of open-mindedness, it is never beneficial to be open-minded in these circumstances. Rather, one should make one's own decision in isolation and not be influenced by others. The merit of this finding is that the formal model gives us a precise specification of some conditions where open-mindedness can be epistemically harmful. In circumstances where the source evaluative capacity is below chance level, i.e., where one is incompetent in evaluating arguments or pieces of evidence, it would be epistemically disastrous to be open-minded.

These two figures give rise to two general conjectures. To formulate these conjectures, recall that variable p_T is used to denote the trustee's accuracy, i.e., the chance that a trustee supports the right alternative (see (Eq. 1)). The two conjectures are as follows. First, whenever the trustee's accuracy is above chance level, then increasing the degree of open-mindedness leads to higher levels of individual accuracy (Result 1). Second, whenever the trustee's accuracy is below

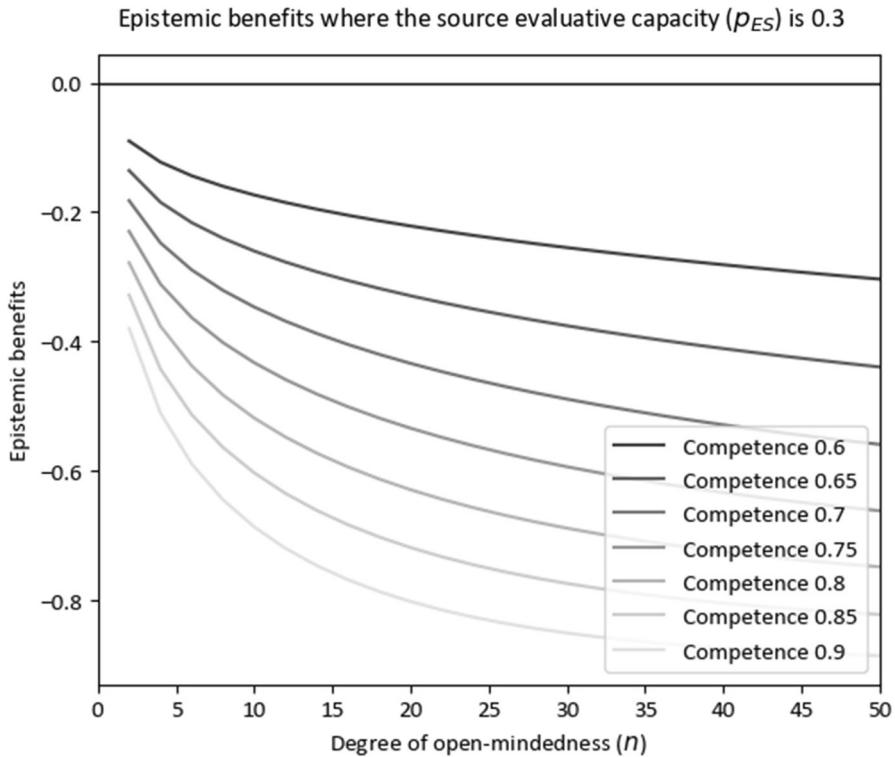


Fig. 3 Epistemic benefits of open-mindedness in the slightly pessimistic case where the source evaluative capacity equals 30%

chance level, then increasing the degree of open-mindedness leads to lower levels of individual accuracy (Result 2).

Result 1 (Good Conditions)

Let $n, m \in \mathbb{N}$ and let $p_R, p_U, p_{ES} \in [0, 1]$. Assume that $p_T > 0.5$. Then,

- (a) $n > m > 0$ implies $p_{EIA}(n, p_R, p_U, p_{ES}) > p_{EIA}(m, p_R, p_U, p_{ES})$,
- (b) $\lim_{n \rightarrow \infty} p_{EIA}(n, p_R, p_U, p_{ES}) = 1$.

Result 2 (Bad conditions)

Let $n, m \in \mathbb{N}$ and let $p_R, p_U, p_{ES} \in [0, 1]$. Assume that $p_T < 0.5$. Then,

- (a) $n > m > 0$ implies $p_{EIA}(n, p_R, p_U, p_{ES}) < p_{EIA}(m, p_R, p_U, p_{ES})$,
- (b) $\lim_{n \rightarrow \infty} p_{EIA}(n, p_R, p_U, p_{ES}) = 0$.

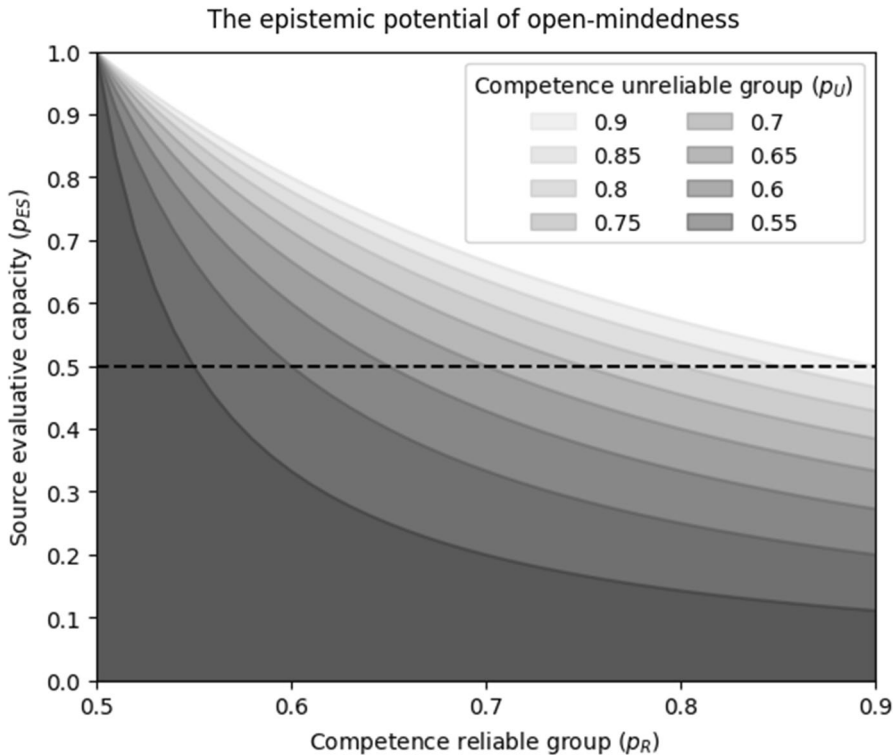


Fig. 4 Minimum source evaluative capacity needed for open-mindedness to have epistemic potential

These results obtain because the mathematical assumptions are structurally similar to the framework and conditions of Condorcet's Jury Theorem. Proofs are omitted.

These analytical results illustrate that the epistemic potential of open-mindedness depends on the trustee's accuracy, which should not come as a surprise. In particular, open-mindedness has epistemic potential if and only if the trustee's accuracy exceeds chance level. To further investigate the exact circumstances where open-mindedness has epistemic potential, consider Fig. 4, where the x-axis represents different levels of competence for the reliable group and the y-axis represents the minimum level of source evaluative capacities needed for the trustee's accuracy to exceed chance level. The figure contains several graphs, where each is associated with a specific level of competence for the unreliable group.

Let me highlight some observations. First, to understand the visualization it may be helpful to add that, for any given level of competence for the unreliable group, points above the associated line represent circumstances where the trustee's accuracy is above chance level. That is, such points represent scenarios where open-mindedness has epistemic potential. By contrast, for any given level




Fig. 5 Epistemic benefits of open-mindedness in the optimistic case where the source evaluative capacity equals 70%

of competence for the unreliable group, the coloured area below the associated line represents circumstances where open-mindedness has no epistemic potential.

Second, as expected, notice that higher levels of competence of the reliable group imply that the required level of source evaluative capacities is lower. Moreover, higher levels of source evaluative capacities are required for higher levels of competence of the unreliable group.

Third, although most instances are such that open-mindedness has epistemic potential if the source evaluative capacities exceed chance level (i.e., all the sections of the graphs below the horizontal dashed line representing $p_{ES} = 0.5$), there are also many circumstances where open-mindedness has no epistemic potential. The coloured areas depict such epistemically hostile circumstances. For example, open-mindedness has no epistemic potential in any scenario where the competence of the reliable agents does not exceed 0.6, the competence of the unreliable agents exceeds 0.7, and the source evaluative capacity does not exceed 0.65. As a consequence, the circumstances where open-mindedness has epistemic potential are more restricted than commonly thought: there are circumstances where open-mindedness has no

epistemic potential even though reliable agents are reasonably competent and are reasonably good at source evaluation.

In any case, Results 1 and 2 precisely state the exact conditions under which open-mindedness could, in principle, bring about epistemic benefits. That is, if the trustee's accuracy is above chance level, then there exists a degree of open-mindedness n^* such that degrees of open-mindedness above n^* will always lead to epistemic benefits, in the sense that the agent's epistemic accuracy would exceed her individual competence. In contrast, if the trustee's accuracy is below chance level, then any degree of open-mindedness will decrease the agent's epistemic accuracy. In short, in these bad conditions, open-mindedness will always be epistemically risky; in the good conditions, open-mindedness may lead to epistemic benefits depending on whether the degree of open-mindedness is sufficiently high.

6 The bad news

Until this point, my analysis has neglected that people can only consider a limited number of arguments or opinions. In good conditions where the trustee's accuracy does exceed chance level, open-mindedness will lead to epistemic benefits when the degree of open-mindedness is sufficiently high. In this section, I will further explore the conditions under which open-mindedness is epistemically beneficial for *feasible* degrees of open-mindedness, under the assumption that the trustee's accuracy is above chance level. I will demonstrate that the tipping point for the degree of open-mindedness to become epistemically beneficial depends on the competence of the reliable group, the competence of the unreliable group, and the source evaluative capacity. Limited cognitive capacities considerably restrict the circumstances where open-mindedness is epistemically beneficial.

This section is divided into three subsections. In Sect. 6.1, I will argue that low degrees of open-mindedness are virtually never epistemically beneficial. In Sect. 6.2, I will present an example to illustrate that feasible degrees of open-mindedness can be epistemically harmful even if the trustee's accuracy exceeds chance level. In Sect. 6.3, I identify the conditions under which feasible degrees of open-mindedness are epistemically beneficial by exploring the model's parameter space: it turns out that the conditions under which open-mindedness generates basic epistemic goods are considerably more limited than previously thought.

6.1 Low degrees of open-mindedness are not epistemically beneficial

Let us start with the observation that low degrees of open-mindedness will not be epistemically beneficial when the agent's individual competence exceeds the trustee's accuracy. That is, in the circumstances where $p_T < p_R$. After all, if an agent is more likely to support the right alternative than her trustees, then it is questionable whether she should base her final decision on the opinions of her trustees.

Under which conditions does an agent's individual competence exceed the trustee's accuracy? It is easy to verify that this obtains whenever the probability that a reliable agent succeeds in supporting the right alternative exceeds the probability that an unreliable agent fails to oppose the right alternative. That is, $p_T < p_R$ holds if $(1-p_U) < p_R$. After all, under the assumption that $(1-p_U) < p_R$, we get

$$p_T = p_R \cdot p_{ES} + (1-p_{ES}) \cdot (1-p_U) < p_{ES} \cdot p_R + (1-p_{ES}) \cdot p_R = p_R.$$

So, if the competence of reliable agents exceeds the probability that an unreliable agent fails to oppose the right alternative, then the trustee's accuracy is lower than the agent's individual competence. Note that this condition is very weak and that it is, for instance, met when the competences of each group are better than chance level.¹⁴ This observation entails that low degrees of open-mindedness are virtually always epistemically harmful: when an agent only considers the opinion of a single other agent, then her individual accuracy will be lower than it would have been if she had been fully close-minded.

6.2 The epistemic risk of feasible degrees of open-mindedness: an example

Although low degrees of open-mindedness are not epistemically beneficial, Result 1 shows that, as long as the trustee's accuracy exceeds chance level (i.e., if $p_T > 0.5$), then higher degrees of open-mindedness will eventually make open-mindedness epistemically beneficial. This brings us to the open question of where the tipping point in the degree of open-mindedness lies. For which degrees of open-mindedness does it become epistemically beneficial rather than epistemically disadvantageous? And, is open-mindedness epistemically beneficial for *feasible* degrees of open-mindedness?

To start addressing these questions, I propose to consider another stylized optimistic numerical example, where the source evaluative capacity equals 70%. Note that, in these circumstances, it is more likely than not that a given trustee is a reliable agent. By analogy, in these conditions, it is more likely than not that an accepted argument or piece of evidence is truth-conducive. In these circumstances, if the agent were to consider five pieces of information, then there is an 84% chance that most pieces of information came from reliable agents. In particular, under the assumption that everyone's individual competence is higher than chance level and everyone is equally competent, the trustee's accuracy also exceeds chance level.¹⁵ Hence, in these circumstances, there exists a tipping point such that degrees of open-mindedness exceeding that tipping point are epistemically beneficial.

One might think that it is therefore always a good idea to be open-minded and to base one's final decision on the opinions of others. To investigate the benefits and

¹⁴ Note that $p_R, p_U \geq 0.5$ implies $(1-p_U) \leq 0.5 \leq p_R$.

¹⁵ After all, $p_T = p_{ES} \cdot p_R + (1-p_{ES}) \cdot (1-p_U) = p_{ES} \cdot p_R + (1-p_{ES}) \cdot (1-p_R) > 0.5$ if $p_{ES} \cdot p_R > 0.5$ and $p_U = p_R$.

risks of open-mindedness in these circumstances, let us consider Fig. 5, where, once again, the x-axis represents different degrees of open-mindedness and the y-axis represents the epistemic benefit. The figure contains several graphs, where each is associated with a homogeneous community where every agent has the same competence level.

First, note that the shape is similar to the hyper-optimistic case (see Fig. 2): higher degrees of open-mindedness always yield higher levels of individual accuracy; and the benefit of increasing one's degree of open-mindedness diminishes for higher degrees of open-mindedness.

Second, in contrast with the hyper-optimistic case, in these circumstances, low degrees of open-mindedness are epistemically disadvantageous. That is, for degrees of open-mindedness below or equal to 4, the individual accuracy is lower than the individual competence. For these degrees of open-mindedness, the individual agent would be epistemically better off not basing her final decision on the opinions of others (at least, insofar as accuracy is concerned). This means that, in these circumstances, open-mindedness is not epistemically beneficial.

Third, notice that increasing one's accuracy by 5% (compared to one's individual competence) requires a degree of open-mindedness exceeding 10. Given that there are various costs involved in collecting and evaluating arguments and evidence, it is plausible that an increase of 5% accuracy will many times not be valuable enough to warrant these investments.

In a nutshell, this numerical example proves that, regardless of the exact competences, in case the source evaluative capacity equals 70%, being open-minded leads to lower levels of individual accuracy *unless the degree of open-mindedness exceeds 4*. And, for the benefits to be noticeable or valuable the degree of open-mindedness needs to exceed beyond 10.

In light of these observations, it is important to recall that empirical research on human cognitive capabilities suggests that people have limited capacities for remembering and processing information and typically can process and recall between 3 and 5 pieces of information (Cowan, 2001, 2010). Even if one were to collect and scrutinize various arguments, it is highly likely that the final decision is only based on a limited number of them. The analytical result on good conditions (Result 1) proves that open-mindedness will lead to epistemic benefits when the degree of open-mindedness is sufficiently high. The cognitive limitation emphasizes that, for open-mindedness to be epistemically valuable, the tipping point for the degree of open-mindedness needs to be below or equal to 4 (because, at best, the agent takes her own opinion plus 4 other pieces of information into consideration). If the tipping point lies beyond 4, then the potential epistemic benefits of open-mindedness are likely to be beyond our human capacities—even though open-mindedness could, in principle, be epistemically beneficial. Although Fig. 5 presents relatively optimistic conditions where the source evaluative capacity equals 70%, this demonstrates that, given our limited cognitive capacities, open-mindedness will produce bad epistemic effects in these circumstances. If all of this is correct, our cognitive limitations seem to put strong restrictions on the potential of open-mindedness to be epistemic valuable.

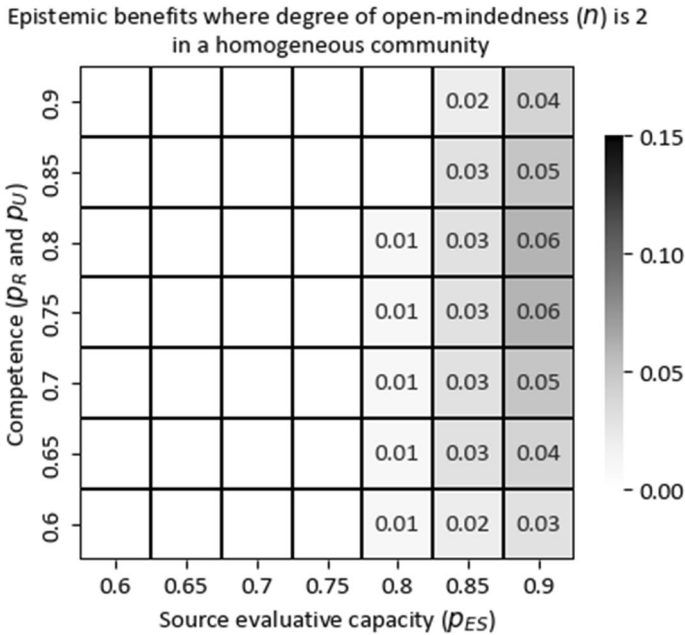


Fig. 6 Epistemic benefits where the degree of open-mindedness is 2 in a homogeneous community

So, one of the conclusions is that, to vindicate the epistemic benefits of open-mindedness, we need to determine which levels of source evaluative capacity are to be expected. At this point, one plausible way to proceed would be to survey the empirical literature for findings on this aspect. In any case, my analysis demonstrates that the conditions under which open-mindedness leads to higher accuracy are considerably more restricted than previously thought: open-mindedness will not be epistemically valuable even for some reasonably competent agents with imperfect but reasonably good source evaluative capacities.

6.3 The epistemic value of feasible degrees of open-mindedness

Are *feasible* degrees of open-mindedness epistemically beneficial for other levels of the source evaluative capacity? To address this question, I propose to explore the model’s parameter space. In particular, let us investigate the following parameter settings:

- The community is homogeneous in the sense that it consists of equally competent agents;
- The individual competence is varied from 0.6 to 0.9;
- The source evaluative capacity is varied from 0.6 to 0.9; and
- In virtue of the cognitive limits to process between 3 and 5 pieces of information, the degree of open-mindedness is set to 2 or 4.

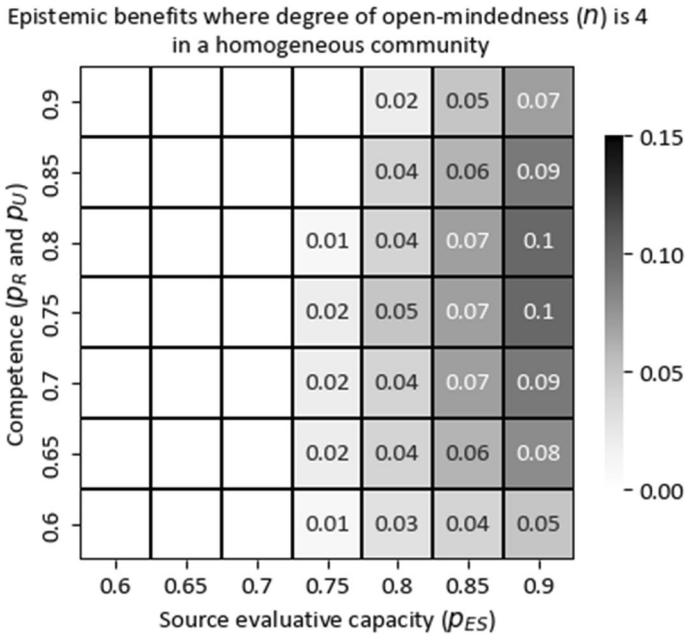


Fig. 7 Epistemic benefits where the degree of open-mindedness is 4 in a homogeneous community

The results are illustrated in Figs. 6 and 7, which are associated with degrees of open-mindedness 2 and 4, respectively. These figures represent the epistemic benefit of open-mindedness by way of a heatmap. The y-axis depicts different individual competences and the x-axis depicts different source evaluative capacities. Each cell in these figures depicts the epistemic benefit of open-mindedness for the associated individual competence and source evaluative capacity. Blank cells depict scenarios where open-mindedness is not epistemically beneficial.

Let me highlight a few observations. First, in general, higher source evaluative capacities will increase the epistemic benefits of open-mindedness. This should not come as a surprise. Furthermore, one might think that unilaterally increasing individual competence will also increase the epistemic benefit of open-mindedness. This is, however, not the case: note that the epistemic benefit does not always increase if we compare any given cell to the one directly above it in these figures. The reason is that the results depicted in these figures are working under the assumption that the community of agents is equally competent, which entails that unilaterally increasing the individual competence of reliable agents will also increase the individual competence of unreliable agents.

Second, these figures can be used to characterize some conditions where open-mindedness is epistemically beneficial and some conditions where it is not. Feasible degrees of open-mindedness are not epistemically beneficial when the source evaluative capacity is below or equal to 70%, irrespective of the individual competence. Furthermore, feasible degrees of open-mindedness are epistemically beneficial when the source evaluative capacity is above or equal to 85%, irrespective of the

individual competence. When the source evaluative capacity is between 70 and 85%, the result is mixed.

Third, for each level of individual competence, it is not straightforward whether feasible degrees of open-mindedness are epistemically beneficial. It all depends on the exact degree of open-mindedness and the exact level of the source evaluative capacity. In other words, from the fact that a community consists of highly competent individuals, one cannot simply conclude that open-mindedness is epistemically beneficial. In particular, this may mean that in the community of (presumably highly competent) scientists, one cannot simply conclude that open-mindedness is epistemically beneficial.

Finally, these figures demonstrate that the epistemic benefit of open-mindedness is likely to be small: in a homogeneous community, the epistemic benefit of feasible degrees of open-mindedness will never exceed 10%. Moreover, in many circumstances the epistemic benefits will not even exceed 5%: for instance, this holds when the source evaluative capacity does not exceed 80%, and it holds virtually always when the degree of open-mindedness does not exceed 2.

All in all, these results prove that it is highly context-dependent whether open-mindedness is epistemically valuable. After all, there are many circumstances where feasible degrees of open-mindedness are not epistemically beneficial. In these circumstances, open-mindedness does not generate basic epistemic goods. The widely held belief that open-mindedness is intellectually virtuous should thus either be qualified or further elaborated. In any case, the conditions under which open-mindedness generates basic epistemic goods are considerably more restricted than previously thought: in particular, Figs. 6 and 7 demonstrate that open-mindedness may produce bad epistemic effects for reasonably competent agents and in reasonably friendly epistemic environments when agents cannot or do not accurately evaluate the content of received information.

7 The good news

Until this point, my analysis did not include the possibility that people are capable of critically evaluating the content of the information they gather. Some content is less believable than others: examples include logical contradictions, obvious falsehoods, weak or invalid arguments, and arguments from false premises.¹⁶ Are feasible degrees of open-mindedness epistemically beneficial when people critically evaluate the content? To address this question, I propose to extend the evaluative step of the model. After the agent has gathered information and evaluated the sources (i.e., in step 3 of Fig. 1), the agent examines the content of the communicated information and filters the previously accepted information before forming her final opinion. Following Sperber et al. (2010) one could interpret the previous evaluation procedure as modelling epistemic vigilance towards the source. Information from distrusted

¹⁶ See Goldman (1997) for an analysis of good argumentation and interpersonal justification.

sources is discarded. After the sources that are deemed untrustworthy are eliminated, the agent practices epistemic vigilance towards the content.

This section is divided into three subsections. In Sect. 7.1, I will expand the mathematical model to include content evaluation. In Sect. 7.2, I will consider cases where agents only practise content evaluation and no source evaluation and demonstrate that this yields mixed results: although there are many circumstances where open-mindedness is epistemically valuable, there also exist reasonably good conditions where open-mindedness is not epistemically valuable. In Sect. 7.3, I present the good news regarding scenarios where agents practise both content evaluation and source evaluation: under these conditions, open-mindedness is epistemically valuable for reasonably competent agents.

7.1 Modelling content evaluation

To model this additional step of content evaluation, I propose to include a new parameter called *content evaluative capacity*. As before, the source evaluative capacity is represented by a probability p_{ES} , while *content* evaluative capacity is represented by a new probability p_{EC} . Intuitively, the content evaluative capacity p_{EC} represents the conditional probability that the agent accepts a piece of information or argument, given that the piece of information or argument supports the right alternative. For simplicity's sake, I assume that this conditional probability equals the conditional probability that the agent rejects a piece of information or argument, given that the piece of information or argument supports the wrong alternative.¹⁷ So, the agent is assumed to be equally good at evaluating truth-conducive and non-truth-conducive information and arguments.

The idea that people are imperfect in evaluating the content of arguments and information is not only supported by the empirical fact that people are, at most, boundedly rational but also by the observation from non-monotonic logic that truthful information can lead people to legitimately endorse false conclusions. The non-monotonic nature of evidence illustrates that there may be conceptual reasons to think that the content evaluative capacities are imperfect. In other words, even the most ideal, rational and intelligent agents will most likely not have perfect content evaluative capacities.

Given this updated model, let me explain how to calculate the probability that an accepted argument or piece of information (i.e., one that survives both the source evaluation and the content evaluation) supports the right alternative. Let us use the phrase *accepted information accuracy* (or, information accuracy for short), notation: p_I , to refer to the probability that an accepted piece of information or argument

¹⁷ In fact, it might make sense to distinguish four cases: (i) right information from reliable sources, (ii) wrong information from reliable sources, (iii) right information from unreliable sources, and (iv) wrong information from unreliable sources. It could be plausible to assume that one's content evaluative capacities are different for each of these cases because it may be harder to expose wrong information from reliable sources than wrong information from unreliable sources. Alas, for reasons of scope and simplicity, I adopt the simplifying assumption that one's content evaluative capacities are the same in each of these four cases.

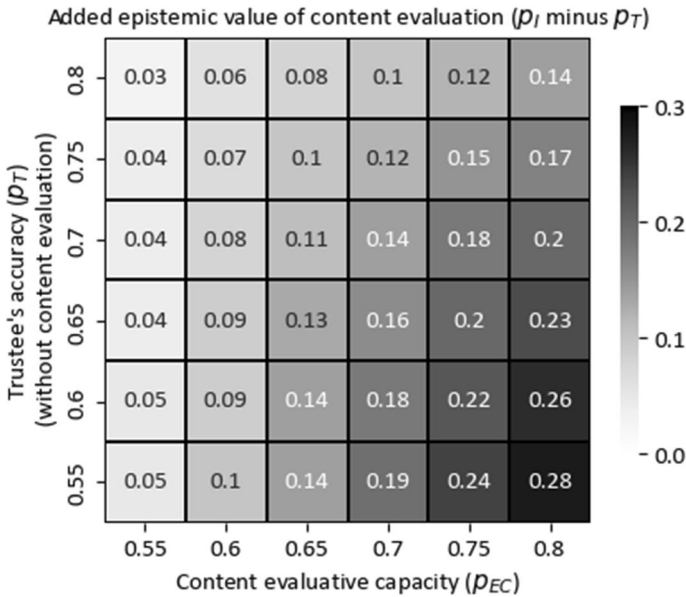


Fig. 8 The added epistemic accuracy of accepted information when practising content evaluation

supports the right alternative. Suppose an accepted source communicates a piece of information or an argument. Recall that the trustee’s accuracy (p_T) represents the probability that information from an accepted source supports the right alternative. The probability that the information from an accepted source supports the right alternative *and* also survives content evaluation is thus given by $p_T \cdot p_{EC}$. By analogy, the probability that the information from such an accepted source supports the wrong alternative and survives content evaluation is given by $(1 - p_T) \cdot (1 - p_{EC})$. Thus, the probability that an accepted piece of information supports the right alternative is given by quotient:

$$p_I = \frac{p_T \cdot p_{EC}}{p_T \cdot p_{EC} + (1 - p_T) \cdot (1 - p_{EC})},$$

where the denominator expresses the probability that information from an accepted source is accepted after content evaluation (regardless of whether it supports the right alternative) and the numerator expresses the probability that an accepted piece of information supports the right alternative.

Based on this updated model, it follows that practising vigilance towards the content typically has an epistemic payoff: practising both source and content evaluation typically improves one’s accuracy compared with only practising source evaluation. More precisely, from the perspective of my formalism, the information accuracy is higher than the trustee’s accuracy if (and only if) the agent is better than chance at evaluating the content, i.e., if $p_{EC} > 0.5$. More formally (the proof can be found in the Appendix):

Epistemic benefits of only content evaluation where degree of open-mindedness (n) is 2 in a homogeneous community




Fig. 9 The epistemic value of open-mindedness when practising only content evaluation and the degree of open-mindedness is 2 in a homogeneous community

Result 3

Let $p_{EC}, p_T \in [0, 1]$ such that $p_T \neq 0$ and $p_T \neq 1$. Then $p_I > p_T$ if and only if $p_{EC} > 0.5$.

Although practising vigilance towards the content has an epistemic payoff, to investigate whether open-mindedness together with practising vigilance towards the content produces basic epistemic goods, it is important to quantify the magnitude of this epistemic payoff in more detail. To make a start, consider Fig. 8, which represents the difference between the information accuracy and the trustee's accuracy for different levels of content evaluative capacity. The y-axis depicts different accuracy levels for trustees and the x-axis depicts different levels of the content evaluative capacity. Each cell in these figures depicts the difference between the information accuracy (with content evaluation) and the trustee's accuracy (without content evaluation).

In short, if one is mildly competent in evaluating the content of information, doing so increases the epistemic accuracy of the accepted information. In turn, this improves the epistemic prospects of open-mindedness. To investigate the epistemic value of open-mindedness more directly, I propose to assess the epistemic benefit of open-mindedness for different levels of the source evaluative and content evaluative capacity. In this way, we can get a fuller picture of the conditions under which open-mindedness is epistemically valuable.

Epistemic benefits of only content evaluation where degree of open-mindedness (n) is 4 in a homogeneous community




Fig. 10 The epistemic value of open-mindedness when practising only content evaluation and the degree of open-mindedness is 4 in a homogeneous community

7.2 Only content evaluation: mixed news

As a first step, let us consider simplified scenarios where the agent practices only content evaluation and no source evaluation. Under the assumption that the absence of source evaluation means that accepted sources are equally likely to be trustworthy than not, this assumption can be represented by setting the source evaluative capacity to 50%. Figures 9 and 10 represent the epistemic benefit of open-mindedness by way of a heatmap, which are associated with the degree of open-mindedness 2 and 4, respectively. The y-axis depicts different individual competences, and the x-axis depicts different levels of the content evaluative capacity. Each cell in these figures depicts the epistemic benefit of open-mindedness for the associated individual competence and content evaluative capacity. Each blank cell depicts a scenario where open-mindedness is not epistemically beneficial.

Let me highlight some conclusions. First, it should not come as a surprise that increasing the content evaluative capacities always improves the epistemic prospects of open-mindedness. Second, in many cases, open-mindedness is epistemically beneficial. For example, open-mindedness is epistemically valuable if the competences are below or equal to 0.75 and the content evaluative capacity exceeds 0.65.

Third, my model presupposes that the evaluative capacities and competences of individuals are independent. Nonetheless, my model is compatible with the common-sensical idea that people cannot have high content evaluative capacities if they are incompetent in a given field. This would mean that not all parameter settings are

equally plausible or likely. For example, even though open-mindedness is epistemically very beneficial when the competences equal 0.6 and the content evaluative capacity exceeds 0.7, these cases are perhaps not realistic.

Fourth, it may be surprising to see that for open-mindedness to become epistemically beneficial higher levels of individual competence require higher levels of content evaluative capacities. This might entail that (presumably highly competent) scientists need to practice high levels of content evaluation. This is particularly important given worries that most scientific results are false (Ioannidis, 2005). After all, in communities of highly competent individuals, open-mindedness is not epistemically valuable for low levels of content evaluation.

Finally, nevertheless, it is important to highlight that there are several circumstances where open-mindedness is not epistemically valuable even though the agents are fairly competent and have reasonably good content evaluative capacities; Figs. 9 and 10 demonstrate that this is, for example, the case when the competences are above or equal to 0.85 and the content evaluative capacity is below or equal to 0.65, and also when the competences are above or equal to 0.9 and the content evaluative capacity is below or equal to 0.7.

7.3 Both source and content evaluation

Let me proceed with investigating the epistemic value of open-mindedness for agents who practice both source and content evaluation. It should not come as a surprise that (mildly) competent and capable agents who exercise both source and content evaluation will epistemically benefit most from open-mindedness. To vindicate this conclusion, let us determine the tipping point where a certain content evaluative capacity would yield epistemic benefits, for different competence levels and levels of source evaluative capacity.

Figures 11 and 12 depict the tipping points where a given content evaluative capacity would start to yield epistemic benefits by way of a heatmap, which are associated with the degree of open-mindedness 2 and 4, respectively. The y-axis depicts different individual competences, and the x-axis depicts different levels of the source evaluative capacity. Each cell in these figures depicts a value such that content evaluative capacities equal to or exceeding that value would be such that open-mindedness is epistemically beneficial (for the associated individual competence and source evaluative capacity).

Let me emphasize some features and conclusions. First, notice that, in line with Fig. 7, the cells with values below 0.5 (i.e., cells in white) depict scenarios where open-mindedness is epistemically beneficial when one only exercises source evaluation.

Second, many cases where practising only source evaluation is epistemically beneficial are fragile in the sense that practising slightly negative content evaluation would make open-mindedness epistemically disadvantageous in those cases. This observation cuts both ways: many cases where open-mindedness is epistemically disadvantageous are fragile in the same way. More generally, given that many




Fig. 11 Tipping points where higher levels of content evaluative capacities yield epistemic benefits of open-mindedness, where the degree of open-mindedness is 2 and the community is homogeneous

cells depict values between 0.45 and 0.55, whether open-mindedness is epistemically valuable or disadvantageous highly depends on the exact level of one’s content evaluative capacities.

Third, and finally, notice that in many cases open-mindedness is epistemically beneficial if agents have some (positive) content evaluative capacities. Indeed, observe that if the content evaluative capacities exceed 0.55, then open-mindedness is almost always epistemically beneficial. This is good news: in circumstances where agents have some reasonably good content and source evaluative capacities, open-mindedness will be epistemically valuable under many realistic circumstances. More importantly, given the realistic assumptions of my model regarding imperfect competence, imperfect evaluative capacities and cognitive limitations on memory, the good news demonstrates that it is highly likely that open-mindedness generates basic epistemic goods for cognitively bounded agents with reasonably good content and source evaluative capacities.

8 Conclusion

In this paper, I have studied the general conditions under which open-mindedness is epistemically valuable. Instead of pursuing a case-based approach, I have presented a simple mathematical model to examine the conditions under which

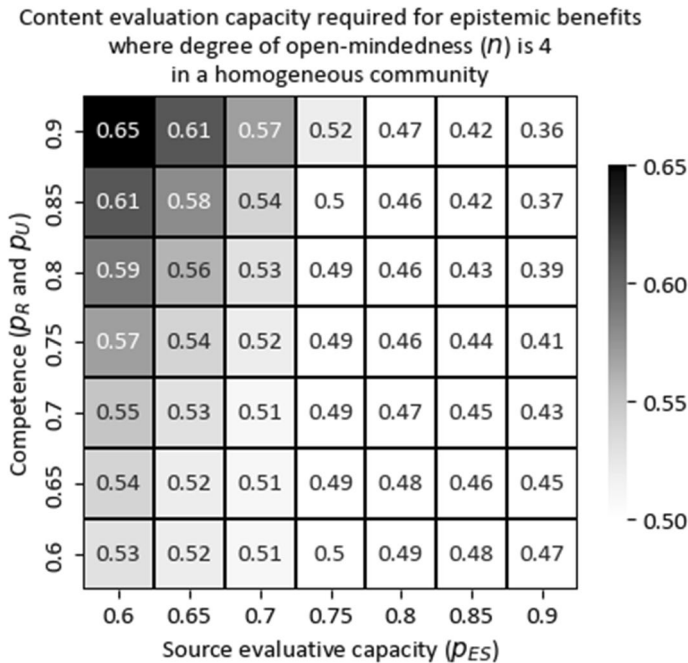


Fig. 12 Tipping points where higher levels of content evaluative capacities yield epistemic benefits of open-mindedness, where the degree of open-mindedness is 4 and the community is homogeneous

open-mindedness is epistemically valuable under some realistic assumptions for bounded agents: (1) people are not perfectly competent in forming their own opinion, (2) people are not perfectly capable of evaluating claims put forward by others, and (3) people can only give consideration to a limited number of arguments or opinions. It turns out that, under these realistic assumptions, open-mindedness does not always produce basic epistemic goods. My analysis shows that open-mindedness may even be epistemically harmful to reasonably competent agents and in reasonably friendly epistemic environments when agents are incapable of accurately evaluating the trustworthiness of the sources or the content of received information (see Figs. 4, 6, 7, 9 and 10, respectively). Hence, the circumstances where open-mindedness is epistemically valuable may be more restricted than previously thought.

My analysis can be understood as setting the limits to open-mindedness being virtuous. Most importantly, my investigation suggests that intellectual virtues should not be thought of as virtues *tout court*. If this is correct, then we urgently need a clear and compelling theory of *when* intellectual virtues are in fact epistemically virtuous. My approach could be viewed as a possible first step to address this issue. My analysis shows that it is highly likely that open-mindedness generates basic epistemic goods for cognitively bounded agents with reasonably good content and source evaluative capacities.

In contrast to a case-based analysis, my aim has been to explicate *the general conditions under which* open-mindedness is epistemically valuable. My modelling

approach makes it possible to explore and characterize the scenarios where open-mindedness is epistemically valuable by representing scenarios within a five-dimensional parameter space. In my analysis, the parameter space is given by the following five dimensions: the degree of open-mindedness (n), the competence of reliable agents (p_R), the competence of unreliable agents (p_U), the source evaluative capacity (p_{ES}) and the content evaluative capacity (p_{EC}). Of course, modelling does not, by itself, lead to the conclusion that open-mindedness is (or is not) epistemically valuable. Rather, modelling helps to delineate those areas of the parameter space where open-mindedness is epistemically valuable and those areas where it is not. Modelling can hence play a key role in explicating the conditions under which open-mindedness is epistemically valuable. For example, if one wishes to argue that open-mindedness is *typically* epistemically valuable, then my model could be used instrumentally to argue that *typical scenarios* are to be found in the areas of the parameter space where open-mindedness is epistemically valuable.

One limitation of my analysis is that it ignores the interplay between different intellectual virtues. However, epistemic virtue theorists typically endorse a more holistic view that the total body of intellectual virtues is together epistemically valuable even though a single intellectual virtue, taken in isolation, may not be. In particular, there may be situations where intellectual virtues such as courage and humility conflict. In these cases, the ability to mediate between different intellectual virtues must also be an intellectual virtue. What is the importance of my analysis for such a holistic view of intellectual virtues? Note that my analysis focuses on cognitively bounded agents who consider and evaluate the source and the content of received information before they make up their minds. Although my analysis shows that there are situations where open-mindedness is not epistemically valuable for these bounded agents, it leaves open whether other intellectual virtues could patch up these situations.

The cognitive limitations in working memory played an important role in my analysis and severally constrain the circumstances where open-mindedness is epistemically valuable, especially when the agent only practices source evaluation (see Sect. 6). In response to this, one may ask whether people can transcend these cognitive limits.¹⁸ People might circumvent these limitations by outsourcing their memory tasks (i.e., making notes or using other tools) or by sequential updating (Bayesian or otherwise). Indeed, on the one hand, one could expand the number of pieces of information to take into consideration by not solely relying on one's working memory. Instead, people (perhaps even virtuous people) may use various tools to complement their limited memory capacities. On the other hand, one could avoid being overly reliant on one's memory by updating one's opinion sequentially. In this way, people may update their original opinion at intermediate points in the face of the currently acceptable information and then later only take their updated opinion into account (i.e., without considering the information that the previous update was based on). If these strategies manage to expand one's working memory, then I believe that my analysis could be extended to degrees of open-mindedness beyond

¹⁸ I thank Naomi Kloosterboer and Kai Spiekermann for encouraging me to think more about this aspect.

4 and would likely suggest that open-mindedness together with enhanced cognitive capacities would lead to epistemic benefits.

My analysis focused on individual accuracy in a single decision problem. But, although open-mindedness may not lead to epistemic benefits in the short run (for a specific decision problem), open-mindedness may still be epistemically fruitful *in the long run*.¹⁹ After all, it seems plausible that open-minded individuals have the opportunity to develop their cognitive and argumentative skills by engaging with others' perspectives, whereas close-minded individuals will have fewer opportunities to do so. I am sympathetic to this view, and it would be interesting to investigate the general conditions under which open-mindedness leads to these long-term epistemic benefits. However, in any case, I would like to note that if it turns out that open-mindedness has these long-term benefits, it could lead to the peculiar situation where open-mindedness is epistemically beneficial in the long run even if it is epistemically harmful at each stage. This seems counterintuitive. In other words, the challenge would be to show how it is possible for open-minded individuals to improve their cognitive and argumentative skills in circumstances where open-mindedness is epistemically harmful in the short run.

Similarly, there might be circumstances where open-mindedness is epistemically valuable *for collectives* despite not being epistemically valuable for individuals. I am also sympathetic to this view, and it would be an interesting project to investigate the general conditions under which open-mindedness leads to these collective epistemic benefits. In like manner, some vice epistemologists have argued that, in some circumstances, close-mindedness may be collectively epistemically beneficial (Battaly, 2018b; Bland, 2022). Another influential example of this sort is the interactionist account of reason by Mercier, Sperber and collaborators (Mercier & Sperber, 2017; Sperber et al., 2010), according to which certain cognitive biases are epistemically beneficial for collectives despite being epistemically harmful to individuals.

I would like to end by pointing out three avenues for future research. First, although my analysis most straightforwardly applies to reliabilist virtue theory, I would like to highlight that the analysis might be generalizable to other views. After all, my main argument is formal in nature and applies an 'instrumental' assessment of open-mindedness. It is instrumental in the sense that it studies the conditions under which open-mindedness will likely bring about certain (epistemic) goods. My model and argument might be adaptable to questions like: When does open-mindedness succeed in bringing about understanding? And, more specifically, is open-mindedness conducive to understanding under some realistic assumptions: (1) people are not perfectly competent in achieving understanding (or those other goods) by themselves, (2) people are not perfectly capable of evaluating the information put forward by others in terms of whether it improves or impoverishes understanding (or for achieving some other goods), and (3) people can only give consideration to a limited number of arguments or opinions. Of course, considerable effort is needed to spell out and assess the extent to which the analysis can be reformulated for other epistemic goods such as understanding.

¹⁹ I thank Nora Kindermann, Naomi Kloosterboer, and Kai Spiekermann for this suggestion.

Second, there is a growing philosophical literature on epistemic injustices. My analysis and model may be instrumental in explicating certain mechanisms that lead to some specific epistemic injustices. As my analysis illustrated, one's individual accuracy is determined by several factors including one's own competence, the competences of reliable sources, the competences of unreliable sources, one's evaluative skills, and one's degree of open-mindedness. Let me pick out one of these factors to illustrate the prospect of using my model to explicate one of these mechanisms: one's source evaluative skills. Although one's source evaluative skills are partly determined by one's own capacities, they also depend on the information available in one's environment. Compare two agents where one is often faced with opinions of unreliable agents (for example, perhaps malicious or biased agents) and the other is often faced with opinions of reliable agents. It seems like the first agent would need better evaluative capacities for the simple reason that they encounter more information that is not truth-conducive. If a certain group of people often find themselves in the first situation, then it is to be expected that they are more suspicious of others' opinions and that their individual accuracy is low. Hence, this argument would illustrate that it is plausible that environmental features and social positions can impact one's source evaluative skills and thereby diminish one's epistemic accuracy.

Lastly, there has been a rising interest in political epistemology, in general, and epistemic theories of democracy, in particular. Philosophers debate whether democracy can be expected to lead to epistemically valuable outcomes (Landmore, 2013), or whether democracy can be expected to lead to suboptimal decision-making (Brennan, 2016). Some have argued that open-mindedness does not lead to an enlightened electorate (Kruglanski & Boyatzki, 2012). One set of models that are often used to argue for the epistemic value of democracy are Condorcet Jury Theorems (Goodin & Spiekermann, 2018). These models often do not include mechanisms for small-scale deliberation and social influence. If my model can be understood as approximating small-scale deliberation and social influence, it would be interesting to see the extent to which my present study impacts the Condorcetian arguments for the epistemic value of democracy given interpersonal social influence.

Appendix

Proof of result 3

Let $p_{EC}, p_T \in [0, 1]$ such that $p_T \neq 0$ and $p_T \neq 1$. Then:

$$p_I > p_T \text{ if and only if } \frac{p_{EC} p_T}{p_{EC} p_T + (1 - p_{EC}) \cdot (1 - p_T)} > p_T.$$

Note that $p_{EC} \cdot p_T + (1 - p_{EC}) \cdot (1 - p_T)$ is non-zero. Because $p_T \neq 0$, the previous inequality is equivalent to:

$$p_{EC} > p_{EC} \cdot p_T + (1 - p_{EC}) \cdot (1 - p_T),$$

which is equivalent to:

$$p_{EC} > 2p_{EC} \cdot p_T + 1 - p_T - p_{EC},$$

which is equivalent to:

$$2p_{EC} \cdot (1 - p_T) > 1 - p_T.$$

Given the assumption that $p_T \neq 1$, this is equivalent to $p_{EC} > 0.5$, as desired.

Acknowledgments I thank audiences in Amsterdam, Munich, and Utrecht for fruitful discussions and helpful suggestions, comments, and objections. Special thanks go to Nora Kinderman, Naomi Kloosterboer, Chris Ranalli, Colin Rittberg, and Kai Spiekermann for written comments on previous drafts.

Funding This work was funded by the European Research Council (ERC-2017-CoG project SEA, No. 771074).

Data Availability All scripts necessary to reproduce the results are available at <https://doi.org/10.5281/zenodo.10955837>.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adler, J. (2004). Reconciling Open-mindedness and belief. *Theory and Research in Education*, 2(2), 127–142. <https://doi.org/10.1177/1477878504043440>
- Adler, J. (2017). Epistemological Problems of Testimony. In E. N. Zalta & Winter (Eds.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab: Stanford University.
- Baehr, J. (2011). *The inquiring mind: On intellectual virtues and virtue epistemology. The inquiring mind*. Oxford University Press.
- Battaly, H. (2018a). Closed-mindedness and dogmatism. *Episteme*, 15(3), 261–282. <https://doi.org/10.1017/epi.2018.22>
- Battaly, H. (2018b). Can closed-mindedness be an intellectual virtue? *Royal Institute of Philosophy Supplements*, 84(November), 23–45. <https://doi.org/10.1017/S135824611800053X>
- Battaly, H. (2020). Closed-mindedness as an intellectual vice. In C. Kelp & J. Greco (Eds.), *Virtue-theoretic epistemology: New methods and approaches* (pp. 15–41). Cambridge University Press.
- Baumann, P. (2013). Knowledge and dogmatism. *The Philosophical Quarterly*, 63(250), 1–19. <https://doi.org/10.1111/j.1467-9213.2012.00104.x>
- Bland, S. (2022). In defence of epistemic vices. *Synthese*, 200(1), 59. <https://doi.org/10.1007/s11229-022-03572-7>
- Brennan, J. (2016). *Against democracy*. Princeton University Press.
- Carter, J. A., & Gordon, E. C. (2014). Open-mindedness and truth. *Canadian Journal of Philosophy*, 44(2), 207–224. <https://doi.org/10.1080/00455091.2014.923247>
- Cassam, Q. (2019). *Vices of the mind: From the intellectual to the political* (1st ed.). Oxford University Press.
- Coady, C. A. J. (1992). *Testimony: A philosophical study*. Oxford University Press.

- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24(1), 87–114. <https://doi.org/10.1017/S0140525X01003922>
- Cowan, N. (2010). The magical mystery four: How is working memory capacity limited, and why? *Current Directions in Psychological Science*, 19(1), 51–57. <https://doi.org/10.1177/0963721409359277>
- de Ridder, J. (2021). Deep disagreements and political polarization. In E. Edebenburg & H. Michael (Eds.), *Political epistemology* (pp. 226–243). Oxford University Press.
- Dietrich, F., & Spiekermann, K. (2023). Jury theorems. In Edward N. Zalta & Uri Nodelman (Eds.), *The stanford encyclopedia of philosophy*. Stanford University, Metaphysics Research Lab.
- Douglas, H. (2000). Inductive risk and values in science. *Philosophy of Science*, 67(4), 559–579.
- Duijf, H. (2021). Should one trust experts? *Synthese*, 199(3), 9289–9312. <https://doi.org/10.1007/s11229-021-03203-7>
- Dutilh Novaes, C. (2020). The Role of trust in argumentation. *Informal Logic*, 40(2), 205–36. <https://doi.org/10.22329/il.v40i2.6328>
- Ecker, U. K. H., Lewandowsky, S., Cook, J., Schmid, P., Fazio, L. K., Brashier, N., Kendeou, P., Vraga, E. K., & Amazeen, M. A. (2022). The psychological drivers of misinformation belief and its resistance to correction. *Nature Reviews Psychology*, 1(1), 13–29. <https://doi.org/10.1038/s44159-021-00006-y>
- Fantl, J. (2018). The limitations of the open mind. *Oxford University Press*. <https://doi.org/10.1093/oso/9780198807957.001.0001/oso-9780198807957>
- Fortunato, S. (2004). The Krause-Hegselmann consensus model with discrete opinions. *International Journal of Modern Physics C*, 15(07), 1021–1029. <https://doi.org/10.1142/S0129183104006479>
- Fricker, E. (1994). Against gullibility. In B. K. Matilal & A. Chakrabarti (Eds.), *Knowing from Words*. Kluwer Academic Publishers.
- Gilbert, D. T., Krull, D. S., & Malone, P. S. (1990). Unbelieving the unbelievable: Some problems in the rejection of false information. *Journal of Personality and Social Psychology*, 59(4), 601–613. <https://doi.org/10.1037/0022-3514.59.4.601>
- Gilbert, D. T., Tafarodi, R. W., & Malone, P. S. (1993). You can't not believe everything you read. *Journal of Personality and Social Psychology*, 65(2), 221–233. <https://doi.org/10.1037/0022-3514.65.2.221>
- Goldman, A. I. (1997). Argumentation and Interpersonal Justification. *Argumentation*, 11(2), 155–164. <https://doi.org/10.1023/A:1007786207614>
- 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>
- Goodin, R. E., & Spiekermann, K. (2018). *An epistemic theory of democracy*. Oxford University Press.
- Hardwig, J. (1985). Epistemic dependence. *The Journal of Philosophy*, 82(7), 335–349. <https://doi.org/10.2307/2026523>
- Hardwig, J. (1991). The role of trust in knowledge. *The Journal of Philosophy*, 88(12), 693–708. <https://doi.org/10.2307/2027007>
- Harman, G. H. (1973). *Thought*. Princeton University Press.
- Hasson, U., Simmons, J. P., & Todorov, A. (2005). Believe it or not: On the possibility of suspending belief. *Psychological Science*, 16(7), 566–571. <https://doi.org/10.1111/j.0956-7976.2005.01576.x>
- Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLOS Medicine*, 2(8), e124. <https://doi.org/10.1371/journal.pmed.0020124>
- Irzik, G., & Kurtulmus, F. (2019). What is epistemic public trust in science? *The British Journal for the Philosophy of Science*, 70(4), 1145–1166. <https://doi.org/10.1093/bjps/axy007>
- Kripke, S. A. (2011). *Philosophical troubles. Collected papers* (1st ed.). Oxford University Press.
- Kruglanski, A. W., & Boyatzi, L. M. (2012). The psychology of closed and open mindedness, rationality, and democracy. *Critical Review*, 24(2), 217–232. <https://doi.org/10.1080/08913811.2012.711023>
- Kwong, J. M. C. (2016). Open-mindedness as engagement. *The Southern Journal of Philosophy*, 54(1), 70–86. <https://doi.org/10.1111/sjp.12163>
- Kwong, J. M. C. (2017). Is open-mindedness conducive to truth? *Synthese*, 194(5), 1613–1626. <https://doi.org/10.1007/s11229-015-1008-6>
- Landmore, H. (2013). *Democratic reason: Politics, collective intelligence, and the rule of the many*. Princeton University Press.
- Lehrer, K., & Wagner, C. (1981). *Rational consensus in science and society*. D. Reidel Publishing Company.

- Levy, N. (2006). Open-mindedness and the duty to gather evidence. *Public Affairs Quarterly*, 20(1), 55–66.
- List, C. (2011). Group communication and the transformation of judgments: An impossibility result. *Journal of Political Philosophy*, 19(1), 1–27. <https://doi.org/10.1111/j.1467-9760.2010.00369.x>
- Madison, B. J. C. (2019). Is open-mindedness truth-conducive? *Synthese*, 196(5), 2075–2087. <https://doi.org/10.1007/s11229-017-1571-0>
- May, K. O. (1952). A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica*, 20(4), 680–684. <https://doi.org/10.2307/1907651>
- Mercier, H. (2016). The argumentative theory: Predictions and empirical evidence. *Trends in Cognitive Sciences*, 20(9), 689–700. <https://doi.org/10.1016/j.tics.2016.07.001>
- Mercier, H., & Sperber, D. (2017). *The enigma of reason*. Harvard University Press.
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63(2), 81–97. <https://doi.org/10.1037/h0043158>
- O'Connor, C., & Weatherall, J. O. (2019). *The misinformation age: how false beliefs spread*. Yale University Press.
- Ranalli, C., & Lagewaard, T. (2022). Deep disagreement (Part 1): Theories of deep disagreement. *Philosophy Compass*, 17(12), e12886.
- Riggs, W. (2010). Open-mindedness. *Metaphilosophy*, 41(1–2), 172–188. <https://doi.org/10.1111/j.1467-9973.2009.01625.x>
- Sharon, A., & Spectre, L. (2010). Dogmatism repuzzled. *Philosophical Studies*, 148(2), 307–321. <https://doi.org/10.1007/s11098-008-9330-1>
- Song, Y. (2018). The moral virtue of open-mindedness. *Canadian Journal of Philosophy*, 48(1), 65–84. <https://doi.org/10.1080/00455091.2017.1335566>
- Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origg, G., & Wilson, D. (2010). Epistemic vigilance. *Mind & Language*, 25(4), 359–393. <https://doi.org/10.1111/j.1468-0017.2010.01394.x>
- Spiekermann, K. (2020). Epistemic network injustice. *Politics, Philosophy & Economics*, 19(1), 83–101. <https://doi.org/10.1177/1470594X19872505>
- Stanovich, K. E. (2003). The fundamental computational biases of human cognition: Heuristics that (sometimes) impair decision making and problem solving. In J. E. Davidson & R. J. Sternberg (Eds.), *The psychology of problem solving* (pp. 291–342). Cambridge: Cambridge University Press.
- Tversky, A., & Kahneman, D. (1974). Judgment Under Uncertainty: heuristics and biases. *Science*, 185(4157), 1124–1131. <https://doi.org/10.1126/science.185.4157.1124>
- Wilholt, T. (2013). Epistemic trust in science. *The British Journal for the Philosophy of Science*, 64(2), 233–253. <https://doi.org/10.1093/bjps/axs007>
- Zagzebski, L. T. (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge University Press. <https://doi.org/10.1017/CBO9781139174763>

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