Journal of the American Philosophical Association (2026) 235–254 © The Author(s), 2026. Published by Cambridge University Press on behalf of the American Philosophical Association. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (http://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.
DOI: 10.1017/apa.2025.10023

Epistemic magnetism‡

ABSTRACT: *An agent's epistemic prospects depend on a combination of that agent's individual characteristics and features of that agent's epistemic environment. Such factors cannot always be cleanly separated. Often, individual characteristics impact agents' epistemic prospects by shaping the epistemic environments in which individuals find themselves. In particular, features of individuals often repel or attract certain sorts of information, a phenomenon I label epistemic magnetism. I argue that epistemic magnetism is a ubiquitous and underrecognized phenomenon that sometimes promotes and sometimes frustrates the achievement of positive epistemic outcomes. Then, I consider a series of simple proposals concerning what distinguishes between beneficial and harmful forms of epistemic magnetism. I argue that these proposals cannot capture the impacts of epistemic magnetism. Instead, I offer a series of principles that serve to roughly characterize the consequences of this phenomenon. I conclude with some remarks on why epistemologists have thus far tended to overlook epistemic magnetism.*

KEYWORDS: coverage reliability, epistemic bubbles, epistemic environments, social epistemology, testimony

1. Introduction

Are epistemic successes and failures due to traits of individual believers or to features of the environments in which these individuals find themselves? As with certain structurally similar questions—whether individual traits or situational factors cause behavior, for example—the answer to this question is both obvious and, considered in itself, unsatisfying. Both individual factors and features of the environment contribute to epistemic successes and failures. But which of these factors is more important in particular cases, and how do they interact? One possible answer to the latter question is that individual and environmental factors simply combine, in a summative fashion, to yield epistemic outcomes (Bland 2022a: 22–23). Because epistemologists often investigate these factors independently (Bland 2022b: 67), much of the literature in epistemology might be taken to implicitly endorse this view. According to an alternative view, individuals and environments exhibit mutual influence on one another, and epistemic outcomes depend on both factors and on the ways in which they interact (Bland 2022a: 23–24).

This article contributes to the development of this alternative, *interactionist* view by focusing on a significant but neglected way in which individuals shape features of their local epistemic environments and thereby contribute to both positive and

‡This article has been updated since original publication. A notice detailing the change has been published.

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

negative epistemic outcomes. In particular, features of individuals regularly function to attract some sorts of information and to repel other sorts. I refer to this phenomenon as *epistemic magnetism*. My chief aim in this article is to show that epistemic magnetism is a common phenomenon that complicates attempts to explain epistemic successes and failures in terms of environmental factors, individual traits, or any simple combination of these. Instead, features of individuals regularly shape the epistemic environment as it is experienced by the individuals in question. What is more, epistemic magnetism is often rooted in factors that, on the face of things, have little to do with reliable or unreliable belief formation. Partially as a consequence of this, epistemic magnetism is often exercised without the awareness of the individuals in question.

In what follows, I highlight several ways in which features of the environment significantly impact epistemic outcomes. Then, I offer several examples of epistemic magnetism. Next, I argue that certain *prima facie* plausible ways of capturing the distinction between beneficial and harmful epistemic magnetism are unsuccessful. Having argued that such simple principles are unsuccessful, I then offer a series of generalizations meant to roughly capture some conditions under which various forms of epistemic magnetism are, generally speaking, beneficial or harmful. I conclude with some speculation as to why epistemic magnetism has thus far received little direct attention from philosophers.

2. Environment, Individual, and Epistemic Interactionism

To begin, it will be useful to clarify some terminology that will be extensively used in the remainder of the article. First, I will regularly use the terms *epistemic outcomes* and *epistemic prospects*. Epistemic outcomes are to be understood as the epistemically significant results of actions and other processes. I take this to be a broad category. For some examples, gaining knowledge, coming to believe, forming a rational or safe belief, gaining understanding, obtaining justification, obtaining defeaters, as well as the loss of these, are epistemically significant outcomes. Some outcomes, including gaining true or rational beliefs, knowledge, and understanding, are positive. Others, like coming to believe a falsehood or losing justification, are negative. Thus, I will often speak of good and bad epistemic outcomes. I describe increases in the chances of obtaining good epistemic outcomes as improvements in epistemic prospects, and increases in the chances of obtaining bad epistemic outcomes in terms of the worsening of epistemic prospects.

The concept of *epistemic environments* will also be important in what follows. For present purposes, the most important point to note here is that epistemic environments are characterized in part by social epistemic features. For example, it may be an important feature of a given epistemic environment that there is a strong norm of honesty operative among agents in that environment. Relatedly, the quality of the testimony within a given environment is an epistemically significant feature of that environment. More generally, the health of epistemic environments is determined in large part by the availability of good and bad information and, because much of our information is derived from others, social factors are central to the health, or lack thereof, of epistemic environments. Given the centrality of information derived from others, a crucial feature of any given epistemic environment is evidence concerning the reliability of various sources. Consider, as an extreme case, the echo chamber

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

operative in a particular cult (Nguyen 2020). A highly significant feature of this epistemic environment will be testimony, and perhaps other forms of evidence, that call into doubt the competence and/or benevolence of those outside of the cult.

As a final note on the concept of epistemic environments, it is worth acknowledging that ways of individuating epistemic environments will inevitably be somewhat arbitrary. To illustrate, consider first the well-known barn façade case (Goldman 1976). In that case, we imagine an agent who, based on visual perception of a real barn, comes to believe that there is a real barn before him. However, because he is in a region that includes many barn façades that would be mistaken for real barns (call these fake barns), the agent's true belief is unsafe and thus plausibly fails to constitute knowledge. Beyond this description of the barn façade case, we may specify that fake barn country is part of a broader environment within which, as in the real world, there are very few fake barns and many real barns. The relatively small fake barn country, and the broader world that includes many real barns and few fake barns, may be described as epistemic environments on different scales, and epistemic environments of many intermediate sizes could also be specified.

Turning to digital environments, we may say that an individual person's social media feeds, as well as the evolving stream of social media content more generally, constitute social epistemic environments of different "sizes." What is crucial for present purposes is that the conditions of local epistemic environments are sometimes what matters to agents' epistemic prospects. Thus, returning to the fake barn case, the condition of the agent's local epistemic environment might be inconducive to epistemic goods, even if the broader environment is largely conducive to positive epistemic outcomes. Similarly, the erection of many fake barns in a narrow geographical region might compromise epistemic outcomes in that region, even if they have no ill effects more generally.

Consider a salient example of how environmental factors might influence the epistemic prospects of those within the environment. Misinformation, a broad category of misleading information that plausibly includes subcategories like fake news and deceptive deepfakes, is widely thought to pose various threats to the acquisition of knowledge by way of jeopardizing various conditions thought to be necessary for knowledge. Insofar as it is misleading, misinformation promotes falsehoods over truths. Individuals may overcorrect for this threat, leading them to avoid forming beliefs even based on legitimate information. Finally, even where misinformation does not prevent the acquisition of true beliefs, it may keep such beliefs from being justified (Fallis 2021), safe (Harris 2021: n. 3; Matthews 2023; De Ridder 2024), or sensitive, or from having some further property arguably necessary for knowledge (Blake-Turner 2020; Harris 2022, 2024).

The focus on misinformation here is only illustrative and the basic point illustrated —that features of the environment can impinge in various ways on epistemic success and failure—holds quite generally. It is, in short, commonly recognized among epistemologists that features of an agent's epistemic environment can make a significant difference to that agent's epistemic prospects. Whereas misinformation can plausibly frustrate the acquisition of knowledge, other features of the environment may have more positive effects. Indeed, this can be seen in some proposed responses to the challenge of misinformation. Some have proposed, for example, that measures to

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

label online misinformation, to remove or otherwise suppress it, or to add additional accurate information might substantially mitigate its epistemic ill-effects. Although the point is not typically put in these terms, the thought guiding such actions is that these changes in the environment function to improve the epistemic prospects of social media users.

In light of the role of environmental factors in determining epistemic outcomes, some have proposed that individuals, states, and other agents ought to intervene in the environment in ways that promote positive epistemic outcomes. Shane Ryan, for example, has advocated for *epistemic environmentalism*, and has specifically proposed a "state register of epistemic polluters" (Ryan 2018, 2022). Relatedly, Neil Levy (2022a) argues that the inculcation of epistemic virtue is of limited use as a corrective for widespread epistemic failure, and instead recommends interventions on the epistemic environment. The epistemic environmentalist approach acknowledges the deep significance of epistemic environmental conditions to our epistemic prospects. The approach then recommends large-scale interventions on the epistemic environment as a means of bettering individuals' epistemic prospects.

Despite the relevance of epistemic environmental features to epistemic outcomes, features of the epistemic environment need not exhibit similar impacts on all individuals. For example, individual abilities to discriminate between genuine and inauthentic information can promote knowledge acquisition even in otherwise *hostile epistemic environments* (Battaly 2018a)—environments that include a substantial amount of lies and other sorts of misinformation. Suppose, for example, that somewhat realistic synthetic images and videos become common on social media. Such a state of affairs would at least arguably[1] present a significant epistemic threat to those individuals who are unable to distinguish between synthetic and genuine content. Indeed, the very distinction between synthetic and genuine content is increasingly blurred by the incorporation of AI-driven editing tools with smartphone cameras and other devices. Still, how realistic an image or video must be to be *convincing* is likely to vary dramatically between individuals, and it is arguably the convincingness of fakes that poses a threat. Suppose, for example, that some people—especially those aware of the existence of synthetic images and videos and their typical defects—can more reliably distinguish between authentic content and synthetic content. Plausibly, individuals belonging to this group will be less epistemically vulnerable than those with comparatively little awareness of, and ability to identify, synthetic media. The relevance of such discriminatory abilities is reinforced by more traditional epistemological work that has emphasized the importance of the ability to discriminate between certain relevant possibilities to the acquisition of perceptual knowledge (Goldman, 1976; Stroud, 1984, pp. 24–25). Attention to the real-world challenge of synthetic media serves to illustrate that, although features of the epistemic environment contribute to epistemic outcomes, such features do not obviate the relevance of individual properties.

[1] Two caveats are in order. First, even realistic synthetic media may be unconvincing if they are not shared through trusted channels (Harris 2021). Second, even synthetic media that appears highly realistic may in principle be identifiable, and indeed identified, as such by detection algorithms.

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

This section has highlighted some respects in which both features of epistemic environments and features of individuals factor into epistemic outcomes. One might be tempted to conclude, then, that epistemic outcomes depend on a simple combination of individual features and environmental features. But such a conclusion would obscure an important point, to be developed and defended in the remainder of this article. Sometimes, properties of individuals contribute to epistemic outcomes by shaping epistemic environments. This point may be understood by analogy with a parallel point that has long been recognized in the study of human behavior. Consider, first, two simple generalizations concerning the production of human behavior. According to the first, the way a given person will behave in any given instance is predictable based on traits specific to that individual. According to the second, the way a given person will behave in any given situation is predictable based on the situational variables operative in that instance. Following John F. Kihlstrom (2013), we may refer to these generalizations as the *doctrine of traits* and the *doctrine of situationism*, respectively. Psychologists have long realized that neither doctrine adequately accounts for human behavior, but nor can human behavior be adequately explained by a simple combination of independent individual and situational variables (Mischel 1973; Buss 1981; Kihlstrom 2013). This is because such variables are not independent of one another. Rather, as *interactionists* maintain, "situations are as much a function of the person as the person's behavior is a function of the situation" (Bowers 1973: 327). Specifically, a given person may shape the situational variables to which that person is then subject by selecting which situations to enter, eliciting behavior from others (Wachtel 1973: 330–331), or intentionally altering their social environments by, for example, using coercion to constrain others' behavior (Buss 1987).

Comparably to how interactionists in the psychology of behavior propose to explain behavior in terms of mutually interacting individual and situational variables, those that we might call *epistemic interactionists* propose to explain epistemic outcomes in terms of mutually interacting individual and environmental factors. The connection between interactionism in the psychology of behavior and in epistemology has been recognized most clearly by Steven Bland (2022a, 2022b, 2022c). The descriptive component of epistemic environmentalism may thus be understood as a form of epistemic interactionism. Similarly, recent work in virtue epistemology that aims to highlight how certain epistemic virtues improve the agent's epistemic prospects by modifying the agent's local epistemic environment (Bland 2022a, 2022c; Levy 2022b; Jarczewski 2024; Harris 2025) may be understood as embodying a virtue-oriented form of epistemic interactionism.

One might imagine rather dramatic illustrations of epistemic interactionism—for example, an eccentric billionaire reshaping a social media platform in epistemically deleterious ways to such an extent that he himself (among others) is epistemically harmed—but I focus in this article on simpler and more commonplace examples. Although the extent of individuals' capacities to shape their local epistemic environments plausibly depends in part on their social statuses, I aim to highlight in what follows that even those without atypical degrees of wealth or power regularly shape their local epistemic environments in significant ways. Additionally, whereas a virtue-oriented epistemic interactionism would tend to focus on ways in which

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

individuals intentionally shape their local epistemic environments for better or worse, I will not take such intentions to be central to the shaping of individuals' local epistemic environments in what follows. Even without the power to significantly alter the broader epistemic landscape, and even without intending to shape their local epistemic environments, individuals can, and routinely do, behave in ways that shape such environments in ways that significantly alter their own epistemic prospects.

3. Epistemic Magnetism

In section 2, I highlighted some ways in which features of the epistemic environment can substantially impact an individual's epistemic prospects. Most features noted in that section were common, in that they were not specific to particular individuals. For example, although token items of misinformation might be encountered only by particular individuals, the prevalence of misinformation is a property of an epistemic environment that is shared. Other things being equal, common features of the epistemic environment can plausibly impact one's epistemic prospects, even if one does not directly encounter them. Suppose, for example, that convincing fake news is extremely common in one's environment but, through sheer luck, one never happens to encounter it. Instead, one luckily happens to only encounter accurate news stories from reliable sources, and only forms beliefs about current events based on these. In this case, the commonality of fake news plausibly still compromises the warrant for one's beliefs and, in this way, prevents one from knowing various things about the world.

Importantly, however, features of the broader epistemic environment that would otherwise compromise knowledge acquisition can be mitigated. For one thing, as we have seen above, *individual* traits—especially discriminatory abilities—can plausibly prevent otherwise epistemically harmful features of the environment from doing damage. In this case, although the environment itself is not especially conducive to epistemic success, features of the individual serve to compensate for the hostile environment.

In other cases, positive epistemic outcomes can be secured, despite an otherwise hostile epistemic environment, by third party efforts to improve a target's local epistemic environment. Consider, as an example, Sanford Goldberg's (2007: ch. 8, 2008) case of a mother who, being aware of the deceptive tendencies of her son's uncle, works to prevent the uncle's testimony from reaching the child. As a second example, consider how, in the US context, judges withhold some information about defendants' character and past crimes from jurors on the grounds that jurors are likely to put too much weight on such evidence (Goldman 1991; Ahlstrom-Vij 2013: 29–30). More generally, those in positions of power and relative epistemic superiority may work to shield others from epistemically harmful phenomena. Here it is worth noting that efforts to shield individuals from features of the broader epistemic environment are not always helpful with respect to the epistemic conditions of those shielded. For one thing, such paternalistic efforts may compromise the development of skills for competently navigating hostile epistemic environments. In this way, some acts of epistemic paternalism may

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

improve targets' epistemic prospects in the short term but worsen them in the long term. Additionally, the consequences of such paternalistic actions may be skewed by misconceptions on the part of the paternalists. Consider, as an example, how science-skeptical parents may opt to homeschool their children due to the (false) belief that public schooling will indoctrinate their children.

The preceding examples illustrate that, whatever the condition of the broader epistemic environment, the regulation of local epistemic environments can dramatically alter individuals' epistemic prospects. Unlike many of the examples considered thus far, which focus on third party efforts to regulate others' local epistemic environments, I will in what follows focus mainly on ways in which properties of individuals regulate their own local epistemic environments by repelling or attracting certain sorts of information. Such properties are of special interest because, in contrast to third-party efforts to regulate epistemic environments, properties rooted in individuals can be expected to play relatively *consistent* roles in individuals' lives. I will refer to cases in which such regulation occurs as cases in which *epistemic magnetism* is manifested. It is worth clarifying why I use the term "epistemic magnetism" to refer to both the attraction and repulsion of information. Magnetism sometimes refers narrowly to an attractive force, as in the case of personalities described as "magnetic." But magnetism, as a physical phenomenon, refers to the capacity to both attract and repel objects with certain properties. "Magnetism," as I use the term in this article, better aligns with the definition from physics than with the term as it ordinarily applies to persons.

To get a feeling for the concept of epistemic magnetism, and especially its ubiquity, consider a series of cases:

> *Dictator:* Diane is a dictator with a reputation for punishing failures on the part of her subjects. Consequently, Diane's subjects will go to great lengths to avoid reporting negative outcomes to Diane.
>
> *Fanatic:* Fatima is an affable and well-liked member of her community, with a reputation for being caring and kind toward those around her. However, to the bafflement of those in her circle, she is a diehard supporter of a political candidate whose public life is plagued by scandal. Otherwise pleasant interactions with Fatima will turn sour if she is presented with evidence of her favored candidate's misdeeds. Thus, to maintain positive interactions with Fatima, those around her have entirely given up on presenting her with such evidence.
>
> *Gossip:* Gary is known to love gossip about celebrities, as well as people in his community. Consequently, people in Gary's circles are often eager to share the latest gossip with him, regardless of its credentials.
>
> *Open Mind:* Olga has a reputation for being highly thoughtful and open-minded, even when it comes to information that is known to conflict with her beliefs. Thus, many people in her community, including those with very different worldviews, look forward to discussing news events with her.

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

Replication: Knowledge of the replication crisis has shaken Ricardo's faith in the social sciences. He regularly criticizes his academic colleagues for talking up new and surprising results from the social sciences. "Talk to me when it's been replicated," he often says. Consequently, Ricardo's colleagues only share with him results that have been replicated.

Shock News: Sae regards *Shock News* as a hyperpartisan and unreliable news source. Consequently, she has in the past heavily criticized anyone who discussed reporting or commentary from *Shock News* in her presence. Thus, people have generally stopped discussing material from *Shock News* around Sae.

Sickness: Sal has four adult children, all of whom know that Sal tends to panic and to fear for the worst when he hears that one of his children is sick. To avoid causing him stress, none of Sal's children ever tell Sal when they are feeling sick.

Voice: A podcaster, prone to sharing heterodox and mistaken takes about politics, history, and health, and for platforming others with similar views, is extremely popular among Vineet's friends and colleagues, who used to regularly recommend the podcast and to send episodes and clips to Vineet. However, because Vineet finds the podcaster's voice to be whiny and annoying, and has told his friends and colleagues as much, they have ceased to send him such material.

War: A war has broken out between two foreign countries. Wanda finds news and footage from the conflict to be highly distressing. For this reason, she adjusts the settings on her social media accounts to filter out any posts containing words, photos, or videos related to the war. As it happens, misinformation concerning the conflict is rampant on social media but, because of Wanda's filters, she does not encounter it.

In virtue of their past actions or their broader traits, individuals in each of these cases tend to either repel some sorts of information, attract some sorts of information, or both. An important point illustrated by such cases is that a wide range of different factors can underlie epistemic magnetism of the sort highlighted here.

I focus here on tendencies to attract and repel information that are rooted in features specific to individuals. I adopt this focus because such cases show especially clearly how individuals' properties shape the local epistemic environments encountered by those individuals. It is worth mentioning, though, that individual but shared traits might likewise have predictable effects on individuals' local epistemic environments. For example, being a human rights attorney or having a certain socioeconomic background is likely to impact the sorts of information others tend to present to a person. Some work in standpoint epistemology may be understood as exploring some of these impacts (Rolin 2009: 220; Dror 2023: 622). However, because I think the effects of epistemic magnetism are clearest

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

when individuals' magnetism is rooted in relatively idiosyncratic properties, I focus on cases involving such properties here. In *Dictator*, for example, a sort of magnetism is due to an implied threat that is made evident by Diane's past behavior. In *Sickness*, by contrast, Sal's tendency to worry is a salient cause of the repulsion of certain information about the health of his children. Although some of the examples presented above are extreme, these examples are intended to illustrate—albeit sometimes very dramatically—features of individuals that repel and attract information in a broad range of real-world cases. For example, *Dictator* is a rather unusual case, divorced from most of our everyday experiences, but similar dynamics plausibly shape the flow of information in, for example, many employer-employee relationships.

These cases also help to illustrate that, although the attraction and repulsion of information is rooted in individual properties, whether these properties actually lead to the attraction and repulsion of information is often dependent on the magnetic agent's social context. For example, it is because members of her community want to retain positive relations with Fatima that they avoid sharing evidence of her favored candidate's misdeeds with her. If those around her were less interested in positive interactions with Fatima, they might share unwelcome evidence with her. Despite the contingent effects of her individual properties, Fatima is magnetic in the sense that she exhibits a tendency to repel certain sorts of information in hospitable environments—namely, those that include agents who want to maintain positive relations with her.

Relatedly, with the exception of *War*, the examples of epistemic magnetism presented above involve agents whose magnetic properties impact what information others share, or decline to share, with them. In *War*, Wanda's repulsion of misinformation is instead enabled by her deliberate filtering of information about the war out of her social media feeds. Similar cases, involving the effects of personalization algorithms on the local epistemic environments of internet users, can easily be constructed. Because epistemologists and other researchers have discussed the epistemic impacts of personalization algorithms at length (Pariser 2011; Miller & Record 2013; Nguyen 2020), I do not focus on them here. However, it is worth noting that, in effect, epistemic magnetism is about the personalization of information. Thus, such personalization is not a novel phenomenon, made possible by social media and search engine algorithms. Rather, such algorithms compound existing forms of epistemic magnetism.

One final feature of epistemic magnetism, as illustrated by these examples, is worth highlighting. In some cases, epistemic magnetism is exercised more-or-less intentionally. Ricardo, for example, deliberately discourages others from sharing information about unreplicated social scientific findings with him. In other cases, however, epistemic magnetism is exercised without the agent making any attempt to attract or repel any kind of information. For example, it is Sal's known tendency to worry that discourages his children from sharing information about their health with him. Because epistemic magnetism can be but need not be exercised intentionally, agents may vary in the degree to which they recognize themselves to be repelling or attracting information. One consequence of this is that agents may vary in their awareness as to how their own magnetic properties skew the

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

information that reaches them. The epistemological implications of this latter point will be discussed in section 5.

4. Substantive and Procedural Epistemic Magnetism

In section 5, I offer a partial account of when (and why) epistemic magnetism either benefits or harms its possessor. First, it is worth considering an intuitive but ultimately mistaken explanation of the difference between positive and negative forms of epistemic magnetism. The shortcomings of this explanation will inform the discussion in the following section, ultimately leading to a less neat, but more accurate, understanding of beneficial and harmful epistemic magnetism.

In general, an item of information might be repelled or attracted in virtue of either its *contents* or its *credentials*, where the latter may include facts about its source, others' attitudes toward it, the process by which it was generated, and so on. Where information is repelled or attracted based on content, let us say that *substantive epistemic magnetism* has occurred. Where information is repelled or attracted based on credentials, let us say that *procedural epistemic magnetism* has occurred. The terminology here is inspired by the distinction between substantive and procedural justice and the distinction between substantive and procedural accounts of autonomy and independence (Stoljar, 2024). Comparably to how substantive justice concerns the quality of outcomes, substantive epistemic magnetism is oriented toward the actual content of information. Comparably to how procedural justice concerns the quality of procedures by which outcomes are decided, substantive epistemic magnetism is oriented toward the quality of the processes that produce information. For an example of substantive epistemic magnetism, notice that testimony along the lines of "x is sick," where x is one of Sal's children, is unlikely to reach Sal's ears. Consider next an example of procedural epistemic magnetism. In *Replication*, Ricardo is interested only in information with certain credentials, specifically, in social scientific results that have been replicated. One might be tempted to generalize from these cases to conclude that substantive epistemic magnetism leads to bad epistemic outcomes, while procedural epistemic magnetism leads to good epistemic outcomes. But this conclusion is premature.

First, it is far from clear that substantive epistemic magnetism always leads to negative outcomes. Plausibly enough, Wanda's filtering out of information about the war is beneficial, in light of the rampant misinformation concerning the war on social media. If Wanda were embedded in a different epistemic environment, within which misinformation about the war was scarcer, and accurate information more common, then her policy might well be epistemically harmful. Still, at least in some contexts, substantive epistemic magnetism is very plausibly epistemically beneficial.

Second, procedural epistemic magnetism will only be conducive to positive epistemic outcomes if the credentials in question are reliable ones. For example, one might easily imagine an agent whose hostility to science is well-known, and thus who is rarely exposed by those in his circle to scientific evidence. Alternatively, one might imagine an agent who is known to be especially interested in anecdotal reports on certain topics, and thus tends to be overexposed to weak and likely misleading evidence of this sort. In short, procedural epistemic magnetism can lead one astray.

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

One final concern for the present proposal is that the distinction between substantive and procedural epistemic magnetism is not always sharp. Recall from above that social epistemic factors, especially testimony, make up important parts of epistemic environments. This is significant, as magnetism with respect to testimony may be understood in terms of both substantive and procedural epistemic magnetism. To see this, consider the *Shock News* case, in which Sae repels discussions of material from *Shock News*. This case is perhaps most naturally described in terms of procedural epistemic magnetism, as Sae effectively repels evidence that is sourced from *Shock News*. Here, the sourcing of the information is regarded as problematic. But this case can also be described in terms of substantive epistemic magnetism. After all, Sae systematically repels information with certain contents—specifically, information of roughly the form "it was reported on *Shock News* that…" or "so-and-so, from *Shock News*, says that…" Given that the same instance of epistemic magnetism can be described as either substantive or procedural, the appeal to this distinction, alone, cannot account for the difference between beneficial and harmful forms of epistemic magnetism.

In this section, I have critiqued what I take to be an initially attractive, but ultimately misguided, account of the difference between beneficial and harmful epistemic magnetism. But, although there is good reason to reject this approach, the distinction between substantive and procedural epistemic magnetism will figure into the discussion below concerning beneficial and harmful epistemic magnetism.

5. When (and Why) is Epistemic Magnetism Beneficial?

So far, I have sought to show that epistemic magnetism plays an important role in the determination of epistemic outcomes. But by what mechanisms does it do so? This section addresses this question by taking up the issues of when epistemic magnetism improves epistemic outcomes and when it worsens them, as well as the explanation of this distinction. These issues are complicated, for several reasons. First, as we have seen, epistemic magnetism comes in different forms. Second, epistemic outcomes are themselves diverse. In a particular case, epistemic magnetism could in principle contribute to some positive epistemic outcomes and some negative ones. Third, as we have seen above in the *Open Mind* case, the consequences of epistemic magnetism may depend on features of the magnetic individual independent of their magnetic properties. For example, differences in relevant background knowledge might impact whether the attraction of information is helpful or harmful. Particularly uninformed agents are likely to be especially susceptible to incoming misinformation, while especially well-informed agents are likely to be better able to interpret incoming information—even where much of the incoming information is inaccurate. Finally, properties of the epistemic environment—especially the relative prevalence of good and bad information—can affect the impact of epistemic magnetism. For example, if an epistemic environment includes a great deal of misleading information then, other things being equal, tendencies to repel information are likely to be more beneficial than tendencies to attract information. In light of these considerations, it is not to be expected that there are simple principles that capture the distinction between helpful and harmful epistemic magnetism. Still, this does not preclude the formulation of

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

some generalizations—albeit piecemeal and incomplete ones. The ultimate aim of this section is to offer such generalizations.

Let us begin with a relatively straightforward observation about how epistemic magnetism impacts epistemic prospects. As a general matter, epistemic magnetism affects the ability to know based on *coverage reliability*. Sanford Goldberg (2010: ch. 6) first introduced *coverage reliability* as a property of communities. Goldberg's basic insight is that a subject's attainment or retention of knowledge can sometimes be due to that subject not having received certain information. For example, one's knowledge that the leader of one's country has not been assassinated is not typically traceable to any particular news report or item of information more generally. Rather, one's knowledge of this is typically based on not having encountered any information to the contrary. One knows the leader has not been assassinated because, if the leader had been assassinated, one would have heard about it. Crucially, knowledge of this kind requires that information of the sort in question is produced and disseminated within the subject's community.

Goldberg argues that an agent's coverage-reliability-based knowledge also depends on the receptivity of that agent to contrary information (2010: 164). In short, if an agent is to know some proposition *p* based on not having received any indications to the contrary, it is not enough that such indications would be produced and shared within the agent's community. Instead, it is necessary that such indications would make their way to the agent in question. Some forms of epistemic magnetism compromise this *receptivity condition*, as can be illustrated by a return to the *Sickness* case. In that case, if Sal knows on a given occasion that his children are healthy, this is owed at least in part to the truth of the following conditional: if one of Sal's children was *not* healthy, he would have heard about it. It is precisely this conditional whose truth is jeopardized by Sal's tendency to repel information about the ill-health of his children. More generally, substantive epistemic magnetism that consists in the repulsion of information jeopardizes coverage reliability and, with it, certain positive epistemic outcomes. On the other hand, a recognized openness to information, as displayed by Olga in *Open Mind* for example, will promote coverage reliability.

How, in turn, does coverage reliability promote good epistemic outcomes? In part, it does so by making available an important pathway to knowledge. Only those whose epistemic networks meet the conditions of coverage reliability can enjoy knowledge that *p* via inferences of the form, *if* p *were false, I would have heard about it by now*. While I have discussed the effects of coverage reliability mostly in terms of knowledge, it is worth noting that such reliability also promotes related positive epistemic outcomes, including justification, the formation of true beliefs, and the abandonment of false beliefs. For example, given adequate coverage reliability pertaining to a certain subject matter, a given person is likely to give up false beliefs related to that subject matter (De Ridder 2024: 2970). On the other hand, insofar as epistemic magnetism reduces coverage reliability by repelling information, negative epistemic outcomes are to be expected.

One further connection between repulsive epistemic magnetism and coverage reliability is worth making here. The attainment of certain positive epistemic outcomes depends not only on *actual* coverage reliability, but also on *perceived*

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

coverage reliability. Believing and hence knowing that *p* based on not hearing that ~*p* requires trusting that the sources in one's environment would report that ~*p* if that were the case (Ryan 2021). To illustrate, suppose that Sal finds out that his children have been hiding information from him and confronts them about this. Suppose that his children then promise, sincerely, to no longer hide information about their illnesses from Sal. In this case, the relevant sort of coverage reliability is restored just so long as the children are committed to following through on their promise. Still, if Sal does not believe their promise, he will not believe a proposition along the lines of *if one of my children became sick, I would hear about it*. Hence, he will likely decline to infer that his children are healthy from not hearing anything to the contrary, and may thereby miss out on both the true belief and thus the knowledge that his children are healthy on occasions when they are. Put differently, even the false perception that one repels certain forms of information may compromise one's attainment of certain epistemic goods.

The effects of epistemic magnetism on coverage reliability are closely related to its effects on the *safety* of the magnetic agent's beliefs.[2] In an influential paper, Ernest Sosa defines safety as follows:

> *Safety*: S's belief that *p* is safe iff S would believe that *p* only if it were so that *p* (1999: 142).

In other words, a belief that *p* is safe iff there are no nearby possible worlds in which the subject believes that *p* and *p* is false. In part, certain forms of epistemic magnetism compromise safety for much the same reason that they compromise coverage reliability. Consider *Sickness* again. In that case, it can easily occur that, although Sal believes his children to be healthy, one or more are in fact sick. Thus, even if Sal's belief is true on a particular occasion, the belief may well be unsafe on that occasion. More generally, cases like *Sickness* illustrate that epistemic magnetism that functions to repel information can render an agent's beliefs unsafe.

Here it is worth noting that *Sickness* is modeled after a case first introduced by Nozick (1981, p. 179). In that case, a grandmother knows that her grandson is well when she sees him, despite the fact that, if the grandson were not healthy, the parents would hide this information from her and she would continue to believe that he was well. That case serves as a counterexample to a straightforward sensitivity condition on knowledge. Nozick addresses this concern by including a reference to the method by which a given belief is formed in the analysis of sensitivity. Similarly, a plausible analysis of safety must, in effect, center on safe methods (Sosa, 1999, p. 149). Given such an analysis, Sal's belief that his children are healthy can be safe when it is based on seeing his visibly healthy children before him, even though it is not safe when it is based only on the absence of reports to the contrary. The relevant upshot, for present

[2] Much of what is said here about the connections between epistemic magnetism and safety could also be applied to the connections between epistemic magnetism and *sensitivity*. However, because safety and sensitivity are typically put forward as rival necessary conditions on knowledge, and because it is often thought that sensitivity is too strong to play such a role (Sosa 1999), I focus here on safety.

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

purposes, is that even if epistemic magnetism forecloses certain routes to positive epistemic outcomes, it need not foreclose those outcomes *tout court*.

Thus far, the discussion of coverage reliability and safety might lead one to think the following: epistemic magnetism that repels information worsens epistemic outcomes, while epistemic magnetism that attracts information improves epistemic outcomes. In cases like *Sickness*, it is the repulsion of information that compromises both coverage reliability and the safety of the agent's beliefs and, in cases like *Open Mind*, the agent's attraction of information promotes both coverage reliability and, given plausible assumptions, the safety of her beliefs. But it would be a mistake to generalize too far from these cases. This is because sometimes, when it comes to information, less is more. *Sickness, War, and Dictator* are cases in which certain kinds of information are repelled. Such repulsion seems to be epistemically harmful in *Sickness* and *Dictator*, but is arguably epistemically beneficial in *War*. Whereas Sal and Diane are deprived of accurate—albeit unwelcome—information, Wanda is deprived of various forms of potentially seductive but misleading information. On the other side of things, *Gossip* and *Open Mind* involve agents who tend to attract certain kinds of information. But, whereas Gary's magnetism is likely to attract thinly-sourced information, Olga's recognized policy of open-mindedness arguably epistemically benefits her, by inviting others to expose her to a wide range of perspectives and bodies of evidence.

To say that Olga's recognized open-mindedness may benefit her is not to say that openness to information is invariably beneficial, or beneficial when taken to an extreme. For example, if many people in Olga's community tend to offer low-quality information, and especially if they do so deliberately in an attempt to deceive Olga, then recognized openness to information may be harmful in that context. Additionally, excessive openness to information—for example a willingness to treat outlandish conspiracy theories as explanations on par with any other—might predictably lead to negative epistemic outcomes (Battaly, 2018a, 2018b; Levy, 2022b). Especially for an agent that is particularly gullible and/or lacking in relevant background knowledge, a high degree of openness to information can easily lead to bad outcomes. However, one can acknowledge these points while still allowing that, among a community of mostly good-faith, reasonably competent interlocutors, a reputation for open-mindedness is likely to attract information in epistemically beneficial ways.

The preceding points may be put in terms of *epistemic bubbles*. In Nguyen's (2020) influential taxonomy, epistemic bubbles are characterized by the exclusion of information, typically according to its source. In effect, repulsive forms of epistemic magnetism function to sustain localized epistemic bubbles. On the face of things, this might suggest that epistemic magnetism of this sort worsens the insulated agent's epistemic prospects. However, because information can be wrong, misleading, or otherwise confusing, epistemic bubbles can function to shield epistemic agents from negative outcomes (Lackey 2021; Campbell 2023). Indeed, this is a rationale for programs that attempt to remove misinformation and disinformation from the epistemic environment.

Here, too, it is important not to oversimplify. If Gary's attraction of gossip is epistemically harmful to him, it is plausibly because such gossip tends not to be

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

reliable and because he lacks the background knowledge to distinguish between truths and falsehoods in this category.[3] More generally, and as noted above, the effects of epistemic magnetism will depend on further properties of both the magnetic agent and the broader environment in which that agent is embedded.

As a further complication, it is worth reiterating that a given case of epistemic magnetism might be beneficial in some respects and harmful in others. The attraction of information may promote the formation of both true and false beliefs, depending on the quality of the information attracted. This is why, for example, open-mindedness like that displayed by Olga may be a double-edged sword. Similarly, the repulsion of information may prevent the formation of both true and false beliefs. Sae's repulsion of information from *Shock News* might prevent her from being exposed to, and coming to believe, certain misconceptions that are common in her community. At the same time, however, unless *Shock News* is so unreliable as to be *anti-reliable*—reporting only falsehoods—Sae's repulsion of information is likely to repel some truths. However, as an anonymous referee pointed out, this cost can be mitigated or offset relatively easily, so long as Sae has access to alternative channels of information that share these truths.

Let us turn to a final complication. As noted above, individuals may vary in the degree to which they recognize their own epistemic magnetism. Sal, for instance, may be entirely oblivious to his tendencies to repel information about his children's health. In contrast, Sae likely recognizes that she actively discourages others from sharing material from *Shock News* with her. There is some reason to expect recognized epistemic magnetism to be preferable to unrecognized magnetism. An agent who knowingly repels or attracts information is likely to be at least somewhat better able to estimate the degree to which the information she receives is representative of the full range of available information than an agent who is unknowingly magnetic. This is significant for multiple reasons. First, the former sort of agent is better able to proportion her beliefs to the strength of her evidence, and will likely thus be better positioned to form true beliefs and to avoid false ones. Suppose, for example, that Sae recognizes that Shock News spends an inordinate amount of time critiquing a particular political figure and that she is rarely exposed to such critiques. In this case, Sae can recognize that she is likely missing some evidence that would point to flaws on the part of that political figure. Second, and relatedly, an agent aware of her own magnetism will be in a better epistemic position from which to evaluate her own epistemic condition. In short, epistemic magnetism of which the agent is aware is relatively conducive to higher-order knowledge.

The discussion thus far in this section indicates that the prospects for uncovering exceptionless principles concerning the consequences of epistemic magnetism are dim. Because epistemic outcomes depend on both features of the environment and features of the individual, it is not to be expected that any particular form of epistemic magnetism reliably produces any particular epistemic outcome. Still, the

[3] Plausibly enough, Gary's attraction to gossip is also epistemically harmful insofar as it diverts his attention from more worthy matters. However, a defense of this claim would require some means of ranking the epistemic values of beliefs concerning various subject matters, and this proves to be a contentious matter (Pritchard 2021: 30–33).

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

observations made here are consistent with certain *ceteris paribus* generalizations concerning epistemic magnetism. Here are some plausible candidates:

(1) Given a sufficiently knowledgeable agent and a sufficiently friendly epistemic environment, attractive epistemic magnetism, relative to repulsive epistemic magnetism or the mere absence of epistemic magnetism, tends to promote good epistemic outcomes by, among other things, improving coverage reliability.
(2) Given a sufficiently unknowledgeable agent and/or a sufficiently unfriendly epistemic environment, repulsive epistemic magnetism, relative to attractive epistemic magnetism or the mere absence of epistemic magnetism, tends to mitigate bad epistemic outcomes.

It is worth highlighting here that (1–2) dovetail neatly with recent work in virtue epistemology concerning open-mindedness and closed-mindedness. Heather Battaly (2018a), for instance, has argued that although closed-mindedness is an epistemic vice in some epistemic environments, it is a virtue in others. Relatedly, Levy (2022b) has argued that open-mindedness is likely to lead to bad epistemic consequences in some environments. One upshot of the preceding cases is that the effects of epistemic virtues and vices, as employed in various sorts of epistemic environments, can be approximated by characteristics of a person that are independent of epistemic motivations. For example, Wanda's tendency to be distressed by content concerning a particular conflict may have nothing to do with epistemic motivations, but may, in her environment, have similar epistemically protective effects to those that would be provided by closed-mindedness.

To these generalizations, we may add that:

(3) Procedural epistemic magnetism oriented toward truth-conducive credentials tends to promote positive epistemic outcomes.

Principle (3) accords well with a wide range of recent work emphasizing our epistemic dependence on others, and especially on epistemic authorities (Hardwig 1985; Goldman 2001). Notably, epistemological debates in this area have tended to focus on whether and how individuals can identify reliable authorities. Viewing this issue through the lens of epistemic magnetism helps to illustrate that, in principle, features of a person other than epistemic motivations, strategies, or competencies might make one relatively responsive to information from reliable sources and well-credentialed information more generally, and comparatively unresponsive to information lacking such credentials. Thus, for example, Vineet's aversion to a particular podcaster's voice incidentally functions to repel certain low-quality information. By a similar token, physical attraction to a particular newscaster might function to increase one's exposure to well-credentialed information.

Consider, now, a final *ceteris paribus* generalization based on the remarks above:

(4) Epistemic magnetism of which the magnetic agent is aware is more conducive to positive epistemic outcomes than epistemic magnetism of which the agent is unaware.

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

Because agents aware of their own magnetism will be better able to assess the representativeness of their evidence, such agents will be better positioned to achieve positive epistemic outcomes. To reiterate, (1–4) are offered as *ceteris paribus* generalizations. It is to be expected that rare exceptions can be found but, nonetheless, these principles are plausibly true across a wide range of cases. As one final caveat, it is worth noting that I do not regard these principles as exhaustive of the generalizations that connect epistemic magnetism to epistemic outcomes.

In formulating these ceteris paribus claims, and in the preceding discussion more generally, I have focused primarily on the contribution of epistemic magnetism to the production and prevention of true and false beliefs. But what of further epistemic outcomes, like rational belief? I will assume here that the rationality of a given belief is a matter of the conformity of that belief to one's evidence and to one's other beliefs. I will not comment here on the effect of epistemic magnetism on the achievement of rational belief. Instead, what I wish to highlight is how epistemic magnetism may narrow or widen the gap between rational belief and true belief. On the one hand, as illustrated by cases like *Open Mind*, epistemic magnetism can in principle function to attract high-quality information. In such cases, it is to be expected that, if the magnetic agent forms her beliefs rationally, she is likely to form true beliefs. On the other hand, as in cases like *Sickness*, epistemic magnetism may leave the magnetic agent with an impoverished body of evidence, such that even a rational processing of this evidence will not reliably yield true beliefs. The more general lesson can be summarized in the following principle:

> (5) Epistemic magnetism can narrow the gap between rational and true belief by attracting accurate information or repelling misleading information, and can widen the gap between rational and true belief by repelling accurate information or attracting misleading information.

Unlike (1–4), (5) does not require a *ceteris paribus* caveat. This is due to both the relatively weak, modal nature of the claim and to the metaphorical language of "narrowing" and "widening" gaps. Although I will not attempt to fully precisify (5) here, one further remark is in order. To the extent that an agent has reason to think that the information available to him is skewed by the operations of epistemic magnetism, as in the modified *Sickness* case above, he may have reason to doubt that the information he receives is representative of the relevant information that makes up the broader epistemic environment. In this case, he may have higher-order evidence that recommends against taking his first-order evidence at face value. Consequently, even if his first-order evidence supports a given belief, his total evidence may not. Thus, evidence of the operations of epistemic magnetism may further complicate the relationship between rational and true belief.

6. Concluding Remarks

This article has argued that properties of individuals regularly exert influence on those individuals' local epistemic environments in ways that are highly significant to their own epistemic prospects. Rather than being passive recipients of information,

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

individuals can shape the flow of information into their local epistemic environments —often without recognizing this. Although some instances of epistemic magnetism depend in part on magnetic agents have an atypical degree of power, as in *Dictator*, I have sought to show that epistemic magnetism does not in general depend on outsized social status.

Why is it that epistemic magnetism has received little direct attention from philosophers? Part of the answer, I suspect, is that cases in which epistemic magnetism plays an important role in epistemic outcomes can be easily confused with cases in which another factor, like distrust or motivated reasoning, is dominant. One might think of a case like *Fanatic*, for example, as simply a case in which a person distrusts and intentionally avoids information that would challenge her beliefs. However, while such factors no doubt contribute to epistemic outcomes in this and similar cases, it is important to recognize that these factors are compounded by, and intermingled with, ways in which the agent's traits and behavior shape her local epistemic environment. For example, not only does Fatima distrust and attempt to avoid discordant information, her behavior discourages others from even attempting to present her with such information. One benefit of the epistemic magnetism concept is that it draws attention to how individuals' behaviors and traits that—unlike patterns of trust and motivated reasoning—have no obvious connection to belief formation may nonetheless exert considerable influence on our epistemic prospects.

KEITH RAYMOND HARRIS
PHILOSOPHY, UNIVERSITY OF VIENNA, AUSTRIA
keithraymondharris@gmail.com

Acknowledgment. I would like to thank M. Dentith, Stephen Gadsby, and Jonas Wittwer for their helpful feedback on an early version of this paper. I would also like to thank two anonymous referees for their comments on an earlier draft of this paper. This research was funded by the Austrian Science Fund (FWF) [10.55776/COE3]. For open access purposes, the author has applied a CC BY public copyright license to any author-accepted manuscript version arising from this submission. The sole responsibility for the content of this publication lies with the author.

References

Ahlstrom-Vij, Kristoffer. (2013) *Epistemic Paternalism*. London: Palgrave Macmillan UK.
Battaly, Heather. (2018a) 'Can Closed-mindedness be an Intellectual Virtue?' *Royal Institute of Philosophy Supplements*, 84, 23–45.
Battaly, Heather. (2018b) 'Closed-mindedness and Dogmatism'. *Episteme*, 15, 261–282.
Blake-Turner, Christopher. (2020) 'Fake news, relevant alternatives, and the degradation of our epistemic environment'. *Inquiry*, 1–21.
Bland, Steven. (2022a) '*Interactionism, Debiasing, and the Division of Epistemic Labour*'. *Social Virtue Epistemology* (New York: Routledge), 15–38.
Bland, Steven. (2022b) 'An Interactionist Approach to Cognitive Debiasing'. *Episteme*, 19, 66–88.
Bland, Steven. (2022c) '*Commentary from Steven Bland*'. Social Virtue Epistemology (Routledge), 131–133.
Bowers, Kenneth S. (1973) 'Situationism in psychology: An analysis and a critique.' *Psychological Review*, 80, 307–336.
Buss, Allan R. (1981) '*The Trait-Situation Controversy and the Concept of Interaction*'. The Psychology of Social Situations (Elsevier), 227–234.

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

Buss, David M. (1987) 'Selection, evocation, and manipulation.' *Journal of Personality and Social Psychology*, 53, 1214–1221.
Campbell, Douglas R. (2023) 'In defense of (some) online echo chambers'. *Ethics and Information Technology*, 25, 45.
De Ridder, Jeroen. (2024) 'What's so bad about misinformation?' *Inquiry*, 67, 2956–2978.
Dror, Lidal. (2023) 'Is there an epistemic advantage to being oppressed?' *Noûs*, 57, 618–640.
Fallis, Don. (2021) 'The Epistemic Threat of Deepfakes'. *Philosophy & Technology*, 34, 623–643.
Goldberg, Sanford C. (2007) *Anti-Individualism: Mind and Language, Knowledge and Justification*. 1st edition. Cambridge University Press.
Goldberg, Sanford C. (2008) 'Testimonial Knowledge in Early Childhood, Revisited'. *Philosophy and Phenomenological Research*, 76, 1–36.
Goldberg, Sanford C. (2010) *Relying on others: an essay in epistemology*. Oxford: Oxford university press.
Goldman, Alvin I. (1976) 'Discrimination and Perceptual Knowledge'. *The Journal of Philosophy*, 73, 771.
Goldman, Alvin I. (1991) 'Epistemic Paternalism: Communication Control in Law and Society'. *The Journal of Philosophy*, 88, 113.
Goldman, Alvin I. (2001) 'Experts: Which Ones Should You Trust?' *Philosophy and Phenomenological Research*, 63, 85–110.
Hardwig, John. (1985) 'Epistemic Dependence'. *The Journal of Philosophy*, 82, 335.
Harris, Keith Raymond. (2021) 'Video on demand: what deepfakes do and how they harm'. *Synthese*, 199, 13373–13391.
Harris, Keith Raymond. (2022) 'Real Fakes: The Epistemology of Online Misinformation'. *Philosophy & Technology*, 35, 83.
Harris, Keith Raymond. (2024) *Misinformation, Content Moderation, and Epistemology: Protecting Knowledge*. 1st edition. New York: Routledge.
Harris, Keith Raymond. (2025) 'Social virtue epistemology and epistemic exactingness'. *Episteme*, 1–16.
Jarczewski, Dominik. (2024) '*For the Sake of Knowledge: The Epistemic Value of Other-Regarding Epistemic Virtues*'. Acta Analytica.
Kihlstrom, John F. (2013) 'The Person–Situation Interaction'. In Donal Carlston (ed.). *The Oxford Handbook of Social Cognition* (Oxford University Press), 786–805.
Lackey, Jennifer. (2021) 'Echo Chambers, Fake News, and Social Epistemology'. In Sven Bernecker, Amy K. Flowerree, & Thomas Grundmann (eds). *The Epistemology of Fake News* (Oxford University Press), 208–227.
Levy, Neil. (2022a) *Bad beliefs: why they happen to good people*. First edition. New York: Oxford University Press.
Levy, Neil. (2022b) '*Narrowing the Scope of Virtue Epistemology*'. *Social Virtue Epistemology* (New York: Routledge), 113–130.
Matthews, Taylor. (2023) 'Deepfakes, Fake Barns, and Knowledge from Videos'. *Synthese*, 201, 41.
Miller, Boaz & Record, Isaac. (2013) 'Justified belief in a digital age: On the epistemic implications of secret internet technologies'. *Episteme*, 10, 117–134.
Mischel, Walter. (1973) 'Toward a cognitive social learning reconceptualization of personality.' *Psychological Review*, 80, 252–283.
Nguyen, C. Thi. (2020) 'Echo chambers and epistemic bubbles'. *Episteme*, 17, 141–161.
Nozick, Robert. (1981) *Philosophical Explanations*. Cambridge, MA: Belknap Press of Harvard University Press.
Pariser, Eli. (2011) *The filter bubble: what the Internet is hiding from you*. New York: Penguin Press.
Pritchard, Duncan. (2021) 'In Defense of Veritism'. *Epistemology & Philosophy of Science*, 58, 22–37.
Rolin, Kristina. (2009) 'Standpoint Theory as a Methodology for the Study of Power Relations'. *Hypatia*, 24, 218–226.
Ryan, Shane. (2018) 'Epistemic Environmentalism'. *Journal of Philosophical Research*, 43, 97–112.
Ryan, Shane. (2021) 'Fake News, Epistemic Coverage and Trust'. *The Political Quarterly*, 92, 606–612.

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press

Ryan, Shane. (2022) 'State Epistemic Environmentalism'. In Karyn L. Lai (ed.). *Knowers and Knowledge in East-West Philosophy*. Palgrave Studies in Comparative East-West Philosophy (Cham: Springer International Publishing), 275–292.

Sosa, Ernest. (1999) 'How to Defeat Opposition to Moore'. *Noûs*, 33, 141–153.

Stoljar, Natalie. (2024) 'Feminist Perspectives on Autonomy'. In Edward N. Zalta & Uri Nodelman (eds.) *The Stanford Encyclopedia of Philosophy*. Stanford: Stanford University. https://plato.stanford.edu/archives/sum2024/entries/feminism-autonomy/.

Stroud, Barry. (1984) *The Significance of Philosophical Scepticism*. New York: Oxford University Press.

Wachtel, Paul L. (1973) 'Psychodynamics, behavior therapy, and the implacable experimenter: An inquiry into the consistency of personality.' *Journal of Abnormal Psychology*, 82, 324–334.

https://doi.org/10.1017/apa.2025.10023 Published online by Cambridge University Press