Synthese (2024) 203:81
https://doi.org/10.1007/s11229-024-04504-3

ORIGINAL RESEARCH

Environmental epistemology

Dallas Amico-Korby[1] · **Maralee Harrell**[1] · **David Danks**[1]

Received: 30 March 2023 / Accepted: 18 January 2024 / Published online: 26 February 2024

Abstract
We argue that there is a large class of questions—specifically questions about how to epistemically evaluate *environments*—that currently available epistemic theories are not well-suited for answering, precisely because these questions are not about the epistemic state of particular agents or groups. For example, if we critique Facebook for being conducive to the spread of misinformation, then we are not thereby critiquing Facebook for being irrational, or lacking knowledge, or failing to testify truthfully. Instead, we are saying something about the social media *environment*. In this paper, we first propose that a new branch of epistemology–*Environmental Epistemology*–is needed to address these questions. We argue that environments can be proper objects of epistemic evaluation, and that there are genuine epistemic norms that govern environments. We then provide a positive account of these norms and conclude by considering how recognition of these norms may require us to rethink longstanding epistemic debates.

1 Introduction

It has become fashionable to critique social media platforms like Facebook, Twitter, and TikTok for our social epistemic woes. They've been accused of permitting and amplifying misinformation (Massachi, 2022), giving us a distorted view of ourselves and the world around us (Bail, 2021), fomenting political polarization (Barrett et al.,

✉ Dallas Amico-Korby
damico@ucsd.edu

Maralee Harrell
mharrell@ucsd.edu

David Danks
ddanks@ucsd.edu

[1] UC San Diego, 230 South Weitzel Street Apt 1 Oceanside California, San Diego, California 92054, USA

2021), creating filter bubbles (Sunstein, 2017), and shortening our attention spans (Hari, 2022). These are diverse accusations, but they all imply that platforms are failing to meet some epistemic standard. However, what epistemic standards apply to social media platforms?

In what follows, we start by arguing that there is a large class of questions—specifically questions about how to epistemically evaluate *environments*—that currently available epistemic theories are not well-suited for answering as those theories focus on epistemic agents (in many different forms), rather than the environments within which they operate. However, a critique of Facebook for being conducive to the spread of misinformation is not thereby critiquing *Facebook* for being irrational, or lacking knowledge, or failing to testify truthfully, or falling short on any number of other normative epistemic accounts. Instead, we are saying something about the social media *environment* itself. Specifically, that the environment is designed in such a way that it makes it easy for misinformation to spread, and that this property of the environment is epistemically problematic. This critique implies that there are epistemic standards for evaluating environments.

This class of questions about the epistemic evaluation of environments is exceedingly broad. Social media platforms are obviously epistemic environments, but so are nations, cities, schools, classrooms, businesses, forests, parks, and homes. Moreover, many of these questions are particularly pressing. For example, the US has seen a marked decline in trust in institutions and other people, including social groups and communities (Barrett et al., 2021). Many individuals are ill-equipped to navigate the political and media landscape as it is currently constituted (Wineburg & McGrew, 2019; Pavlounis et al., 2021; Vo, 2023). And facts about political and media environments clearly help explain these changes.

Many philosophers have begun to notice the epistemic importance of environments. For example, Nguyen (2021) argues that environments can be epistemically hostile to beings like us. O'Connor and Weatherall (2020) and Levy (2019) maintain that people's beliefs are not just the result of their individual actions, but also their social networks. Even Kant, in *What is Enlightenment?*, noticed the importance of environments. On his view, enlightenment is an individual achievement, and the failure to achieve it is often self-caused (due to laziness). But Kant also argued that environmental features play a role in determining whether individuals will achieve enlightenment. Specifically, while it is difficult to achieve enlightenment in isolation, "it is more nearly possible for a public to enlighten itself…if only the public is given its freedom" (Kant, 1995, p. 56). 'Freedom' here is a feature of the social structures—that is, the environments—within which the public thinks and reasons, and so enlightenment requires quite radical changes in the environment. For example, it seems to require significant changes in social institutions so that military commanders, the clergy, and even governments do not command but instead reason with those they wish to act in various ways (fight, pray, pay taxes, etc.).

This prior work has often focused on the causal role that environments play in determining what epistemic states individuals can and do form. However, we propose something more radical: environments themselves—and not just agents or groups within them—should also be a locus of (epistemically) normative evaluation. Importantly, our claim is not that we can epistemically evaluate environments indepen-

dently of the agents that occupy them; what makes an environment an epistemically good environment will depend fundamentally on how it affects epistemic agents (as a type, rather than as tokens). However, we are proposing that evaluations of epistemic environments are not *reducible* to evaluations of individual or group cognitive states, and that the tools of traditional epistemology are thus insufficient for such evaluations.

In this paper, we develop and defend a new subfield of epistemology—"Environmental Epistemology"–which has two branches.

Environmental epistemology The study of

A. The effects of environments on epistemic phenomena, and.
B. The epistemic norms that govern environments.

Many philosophers have thought carefully (though not always systematically) about the first branch (A), and so we will not focus (much) on that aspect. Instead, we focus here on a positive account of (B), the epistemic norms that govern environments, and a framework to articulate and evaluate them. We argue that there are two distinct types of epistemic norms for the evaluation of environments: specific norms that arise from the function of the environment in question, and general norms that apply to and govern all environments. We further argue for a health-based understanding of environments: environments are unhealthy to the extent that they fail to satisfy the norms that apply to them.

Our arguments also have implications for traditional epistemic theories. Whatever it is to achieve cognitive success or to possess epistemic virtue surely depends to some degree on the environments that we're in. Epistemic tools that are useful in one environment may not be useful in others, and standards that are satisfactory in one environment may be inadequate in others. Thus, facts about the environment should inform our accounts of cognitive success and epistemic virtue. Relatedly, what counts as an epistemically healthy environment will also depend on the epistemic capabilities, tools, and resources of the agents that occupy them. Thus, we suggest the following: when things go best, there will be an epistemic harmony between individual (or group) and environment—each well suited for the other. Thus, our epistemic theories should provide an account of norms for the evaluation of cognitive agents and the environment, and they should explain the conditions necessary for this harmony.[1] And, when this harmony is absent, we should seek to identify the explanation and means of restoring it.

Section 2 motivates the need for a new subfield of epistemology. We explain our initial social media case in more detail and argue that the tools of traditional epistemology are not well suited for answering the questions this case poses. In Sect. 3, we begin our positive project. We provide a working account of what an environment is

[1] We note an interesting parallel with information theory, where people often focus on the signal, but it is also critical to ensure that we have a channel that is appropriately structured to carry the type of intended signals. Information theorists have developed notions such as channel capacity to describe and evaluate the quality of this type of "environment." Our proposal here calls for an analogous effort for epistemic environments more generally.

and introduce and defend branch (A) of Environmental Epistemology. In Sects. 4–6 we introduce and defend branch (B), including a positive account of the norms that govern environments. Finally, in Sect. 7, we discuss potential payoffs of this new research program, including implications for debates in traditional epistemology (e.g., between internalists and externalists), as well as avenues for future research.

2 The insufficiency argument

This paper provides an extended argument for our contention that we need a new type of epistemology to address questions about the epistemic norms that govern environments. But we begin by motivating the need for such an account with a pressing example, and by arguing that the epistemic tools and concepts that have been developed to this point are not well suited for epistemically evaluating environments. Our motivating example: how do we evaluate social media platforms epistemically? (We consider other types of environments below.)

Social media platforms are online *environments* where people interact and communicate.[2] While these platforms have creators, designers, builders, and maintainers, these people are not the platform. And in the same way that we can evaluate the quality of a home or an office building without evaluating the architects and builders, so we can evaluate a social media platform without evaluating its architects and builders. For example, a house may be a bad house because it lacks a working bathroom. But imagine that it lacks a bathroom because of a fire caused by a lightning strike. The house is bad, but this does not reflect poorly on the architect or builders.[3] So too for social media platforms: we can evaluate them as epistemic environments without evaluating their designers, builders, maintainers, or residents.[4]

Critically, we can evaluate environments without evaluating the individuals that occupy them, though we do need to consider the *type* of individuals who will (or might) occupy them. Whether a house is a good house depends on things like: the affordances it provides (or fails to provide) to "ordinary" people, the effects it has on such people, and so on. These evaluations can (in many cases) be performed with no one living in the house, and without knowing about the specific needs, preferences, or interests of the people who might someday live there. For example, any house with no bathroom is a bad house for individuals who need bathrooms (i.e., all of us), regardless of exactly who lives there. Similarly, online environments can be evaluated epistemically without evaluating any individual's epistemic states; we only need to know about the type of individual who is likely to occupy or use the environment.

Our aim in this section is to show that extant epistemic tools are insufficient for normatively evaluating online *environments* (social media platforms, for example).

[2] We will soon say more about what environments are, but for now we rely on an intuitive sense of the term.

[3] Assuming, of course, that all appropriate safety standards were met.

[4] Of course, the quality of the house often reflects on the architects and builders, and they can bear some responsibility for its quality. The behaviors of the people living in the house may also impact its quality. We discuss these points further in Sects. 2 and 5.

Epistemology has provided us with excellent tools for evaluating epistemic agents and groups—important loci of epistemic evaluation—but we need new tools for epistemically evaluating environments.[5]

We begin by considering what we will call "traditional epistemology." Traditional epistemology has (at least) two foci. First, it focuses on the nature of, and norms for, epistemic states. For example, epistemologists have long attempted to give accounts of beliefs, credences, knowledge, rationality, understanding, etc. for individual people and groups.[6] Second, traditional epistemology has long been concerned with testimony or the social transmission of knowledge, as even Socrates was concerned with determining whether virtue can be taught.[7]

These approaches clearly cannot provide a way to epistemically evaluate social media platforms. These epistemic theories may give us the tools to evaluate Facebook *users*, but when we criticize Facebook for being conducive to the spread of conspiracy theories, we are not criticizing any *individual* (or group of individuals) for their epistemic states. Instead, we're criticizing an *environment* that makes it easy for a problematic epistemic phenomenon (conspiracy theories) to gain traction.[8] Of course, when conspiracy theories do gain traction, particular individuals often form false beliefs. Our traditional theories will tell us why this is epistemically problematic for them, and the fact that the environment is conducive to such beliefs will—as we will show—help explain why Facebook might count as an epistemically unhealthy environment (if the charge is true). But tools to assess individuals' (false) beliefs do not tell us how to assess the environment in which those beliefs are formed.

Similarly, tools to evaluate testimony seem to target the wrong thing. When one complains about Facebook being conducive to the formation of filter bubbles,[9] that criticism does not target Facebook as a testifier, nor as one who receives testimony.[10] Instead, the criticism targets the environment for having features and functionality that are conducive to the formation of a problematic epistemic phenomenon.

Of course, traditional epistemology does not exhaust the field of epistemology. Of late, there has been what we might call an action-guiding turn and a social turn, but these too prove insufficient for answering the questions our case presents.

[5] In the remainder of this section, we will paint with very broad brushstrokes. There are *many* details we will omit, and subtleties will go unremarked. Nonetheless, we hope to convey the lack of resources in most epistemological work to account for epistemic evaluation of environments in themselves. We also, for reasons of space, omit discussion of work in some other areas of philosophy that mention environments (e.g., on measurement and instrumentation in philosophy of science), as those similarly focus on the agent (and so parallel our proposed branch (A)).

[6] Views about what these amount to for individuals abound. For two examples of views about what it might mean for groups see Lackey (2021) and Miller (2015).

[7] See, for example, Protagoras (313d-314a; 320a) and Meno (89e, 91d).

[8] We take no stand on whether social media platforms can have beliefs or credences.

[9] Following Nguyen (2018), a filter bubble is a social epistemic structure in which a person primarily sees and interacts with perspectives similar to their own. So, for example, if one is in a filter bubble of Taylor Swift fans, one may only read or hear perspectives of people who think her music is excellent. See also Sunstein (2017).

[10] We are agnostic about whether social media platforms might play the role of testifier (e.g., when a new policy is announced). Even if so, evaluating Facebook in that capacity is not the same as evaluating Facebook's environment epistemically.

First, consider group epistemology and research on the ethics of belief. Group epistemology attempts to apply the concepts that traditional epistemology applied to individuals (e.g. beliefs and knowledge) to groups (e.g., List, 2005; List & Pettit, 2011; Lackey, 2021). Similarly, the primary focus of the Ethics of Belief literature has been on norms that govern individuals' beliefs, such as whether moral or pragmatic reasons can provide reasons for individuals to believe similarly to epistemic reasons such as evidence (e.g., Pascal, 1941; James, 1896). Both of these traditions, however, are insufficient for the same reasons as traditional epistemology: (1) we are not trying to evaluate either an individual or a group, but an environment, and (2) we are not trying to evaluate beliefs or credences.

Turn next to zetetic epistemology (e.g., Friedman, 2020; Thorstad, 2021). Zetetic epistemologists study the norms that govern inquiry and attempt to answer questions like: "if you want to figure out whether *p*, what should you do?". Notice that answers to this question will necessarily be sensitive to facts about the environment in which one is inquiring. For example, if *p* refers to the proposition "sea water becomes drinkable after boiling," then appropriate actions to determine whether *p* is true will differ if one is on a desert island with no cell reception versus a university campus with a smartphone. However, when we critique social media platforms in the ways that concern us here, we're not interested in the norms of inquiry that govern them, and so the tools of zetetic epistemology will not help.[11]

Finally, social epistemology has largely focused on epistemic norms governing individuals' doxastic states, albeit within distinctively social environments. For example, literatures regarding peer disagreement (e.g., Kelly, 2005; Christensen, 2007; Elga, 2007), testimony, and expertise[12] have largely been concerned with how individuals should alter their doxastic states in response to these phenomena. To the extent this is so, this *portion* of social epistemology too is insufficient for answering the questions at hand. Having said that, some social epistemologists have begun to bump up against the need for the type of evaluative norms we defend here, and we consider that work in the next section.

Instead of broadening epistemology, one might instead consider ways to broaden the notion of an 'epistemic agent'. Embodied accounts of the mind/cognition (e.g., Shapiro, 2007) imply that epistemic agents like us necessarily have physical embodiment, as some aspects of cognition occur through and in our bodies. Or one could turn to extended theories of mind (e.g., Clark & Chalmers, 1998) in which cognition can occur even outside of one's body, as long as the relevant cognitive activities satisfy various criteria, such as being reliably accessible. One common thread in many of these accounts of cognition is the importance of thinking about the environment within which the agent is acting, including the "fit" between the agent and environment. However, while these approaches do consider environments, the locus of epis-

[11] Importantly, this does not mean that norms of inquiry have no bearing on questions related to the spread of misinformation on social media platforms. For example, when considering what individuals ought to do when trying to learn using social media, norms of inquiry will be relevant. Further, we may be concerned with organizational norms of inquiry, such as how Facebook ought to conduct research or organize their research teams.

[12] For a paradigmatic example, consider the literature on moral deference including: McGrath, 2009, 2011; Hills, 2009; Sliwa, 2012; Howell, 2014.

temic evaluation remains the agent (albeit, an agent who might be extended beyond the brain or body), rather than the environment itself (albeit, relative to an entire class or type of agent). These approaches can help us to make sense of claims such as "Facebook has some of my memory," but not "Facebook is epistemically unhealthy" (assuming the Facebook servers do not, for example, randomly delete information). However, the details of the specific account of an embodied or extended mind will matter for the importance of, and criteria for, epistemic roles of environments, and so we leave that issue to future research.

While this has been merely a cursory discussion of an exceedingly rich field, we hope that the general point is clear: the epistemic tools and concepts so far considered are inadequate for epistemically evaluating social media environments like Facebook.[13,14] Traditional epistemic tools are designed for the evaluation of individual and group epistemic states, which are not what concern us when we say that Facebook is epistemically problematic. But Facebook is just an instance of a very broad class. Environments differ in many ways, but what they have in common is that their epistemic evaluation does not turn on whether they have beliefs or credences, nor what those are if they do. In order to evaluate any environment epistemically, we need a new set of tools. To this end, we turn to our positive project.

3 How environments affect epistemic phenomena

Environmental Epistemology as we define it has two branches. The first branch is the study of the effects of environments on epistemic phenomena.[15] In this section, we begin motivating this as a field of study that should be interesting to epistemologists (and we note that it is already a field that interests many non-epistemologists). We start with a broad brushstroke working definition of what we will call *epistemic environments*, and a broad brushstroke account of the constitutive features that determine

[13] Again, we note that many epistemologists *have* done excellent and important work on the social epistemology within the internet, but they have not typically focused on how to epistemically evaluate online and other *environments* themselves. For example, Frost-Arnold (2014) has provided compelling analyses of anonymity and deception online, but she also claims that "any problem of deception can be approached in two ways: (i) by focusing on the speaker and attempting to make her more honest, or (ii) by focusing on the hearer and attempting to shield her from dishonesty […] or increase her abilities to detect and reject falsehoods" (p. 65). We are arguing here that one can *also* evaluate and intervene on the environment. Similarly, Miller and Record (2013) point out that agents often form beliefs on the basis of internet technologies that they do not understand, and they consider the implications of this for the justificatory status of those beliefs. These are valuable contributions, but do not answer the questions we pose. Miller and Record (2017) consider what search engines (construed as subjects) must do to count as epistemically responsible subjects, rather than treating them as environments that contain subjects or agents.

[14] To emphasize how cursory our discussion has been, there are also important traditions that have focused on evaluating agents (see, for example, Zagzebski, 1996; Sosa, 2007) and methods which we have not discussed here (see, for example, Popper, 2002). We believe similar arguments as those considered above apply to these traditions as well.

[15] This is closely related to what Nguyen (2021) calls "hostile epistemology," but his focus is on the ways that environments can be hostile to epistemic agents like us, and not on ways that environments may have positive effects or be conducive to good epistemic phenomena. We wish to expand the scope of this exciting project. See footnote 39 for further discussion of his view.

the nature and character of epistemic environments. We do not intend to provide a full-blown conceptual analysis of epistemic environments nor of their constitutive features.[16] But we do aim to provide the reader with an intuitive grasp of the sort of environmental features that will matter for doing environmental epistemology. As will be clear, our working account makes it such that many things count as environments. We intend this. As our examples will show, environments come in many shapes and sizes.

Epistemic environments (for simplicity, we henceforth refer to *environments* with 'epistemic' implied) are the settings or conditions in which agents engage in epistemic activities.[17] More specifically, we might say that epistemic environments are regions of space (where this includes regions of cyberspace)[18] whose nature and character is determined by (at least) five constitutive features: (1) the type and quantity of present objects (whether physical or virtual/informational), (2) the structure or arrangement of those objects, (3) the norms (moral, pragmatic, epistemic, etc.) that govern the space, (4) the social conventions that hold in the space (if any), and (5) the agents in the space (if any).[19] Each of these features can affect the ways people interact with each other, the content and tone of their interactions, the choices available to them, their conceptualizations of the problems and opportunities that they face, their understanding of the space of possibilities, and ultimately the beliefs and credences they form.

Consider feature (1). Part of what makes swimming pools different from skateparks is the *quantity* of water present. Similarly, a significant environmental difference between a redwood forest and most other forests is the *type* of tree that grows there. The type and quantity of objects present in a space play a role in determining what kind of environment that space is. And these environmental differences have epistemic effects on (or provide epistemic affordances to) agents occupying these spaces. So, for example, the scope of possibilities in a swimming pool and skatepark

[16] While it ultimately will be important to have a more fully worked out account of what an environment is, answering these difficult metaphysical questions is not our purpose here.

[17] We intend *agent* in a minimal sense that will include many non-human animals. While our focus here will primarily be on the health of epistemic environments for people, we take it that non-human animals also engage in some kinds of epistemic activities in environments (for example, a fox may investigate whether a cave is suitable for a den), and so we could identify a notion of epistemic environmental health for them as well. We're grateful to an anonymous reviewer for emphasizing this point.

[18] We do not wish to put our hat in for the substantivalist by referring to *regions of space*, and readers should feel free to replace this with their preferred nomenclature (see Dasgupta, 2015 for discussion of this view and a competitor). We do wish to rely on an intuitive and flexible understanding of *regions of space*. For example, we wish to be able to say things like: a redwood forest is a region of space made up of redwoods, a skate park is a region of space designated by the city for skating, and Twitter is a region of cyberspace made up of (at least) Tweets. What sets the boundaries for such regions depends on the environment in question; in some cases it is social decision or convention (the city decides where the skatepark ends), in other cases not (the redwood forest ends where there are no more redwoods). We are grateful to an anonymous reviewer for encouraging us to think about this issue, and to Tim Juvshik for immensely helpful discussion.

[19] A more complete definition would also address the persistence conditions of environments. While we don't take on this task here, it is worth noting that the duration of an environment will sometimes have epistemic effects. For example, prolonged oppressive environments are likely to have more significant epistemic effects on agents than oppressive environments that are quickly corrected or dissolved.

are different, which can affect what and how people innovate. At the pool, one might discover a new swan dive or cannonball variation, but in a skatepark, one would be unlikely to spend much time even thinking about cannonballs (let alone innovating new variations of them).

Consider (2) next. A major difference between coffee shops and lecture halls is the *arrangement* of objects in the space. Coffee shops typically have clusters of seats that face each other. Lecture halls typically have all the chairs facing the same direction. Even if the two spaces share the exact same type and quantity of seats, their arrangement makes a difference to the environment. Moreover, these environmental differences can have epistemic effects. Coffee shops are often conducive to small group conversations: the seats face each other so it is easy to converse, pick up on non-verbal cues, etc. Lecture halls are not conducive to small group conversation, but rather are more conducive to effective unidirectional communication. These observations do not imply that coffee shops are epistemically better, only different.

Feature (3) concerns the norms that govern the space. A psychiatrist's office and a tabloid writer's office may often have a similar type, quantity, and arrangement of objects, but different *norms* govern these offices. Specifically, one space is governed by stringent confidentiality norms while the other is not (or at least, not to the same degree). And again, this environmental difference will have epistemic effects: many people will be willing to share private information in a psychiatrist's office that they may not be willing to share in the tabloid writer's office precisely because the former (but not the latter) has a norm of confidentiality.

Finally, (4) and (5) capture how the people who occupy a space make a difference to the environment. As a simple example, consider that people often refer to a work environment as "unhealthy" because of difficult or morally problematic colleagues. In other cases, they may have in mind workplace conventions regarding who has power, who reports to whom, etc. These clearly can all have epistemic effects. For example, difficult bosses or unwieldy reporting structures are likely to dampen communication and information sharing.

To this point, we've offered examples of the ways that each of the features of environments can affect epistemic phenomena (i.e., branch (A) of environmental epistemology).[20] Most of our examples have been physically localized environments, such as a park, coffee shop, business, etc. Many environments are more extended in space than these, which can sometimes make it possible for them to have much more significant effects on epistemic phenomena. Appreciating this is important for recognizing the full potential of environmental epistemology.

Consider, for example, the formation of the first agrarian cities. Whatever the causes of their formation, the resulting increased population density and reduction in the need to hunt and gather food had significant epistemic effects. The most obvious epistemic effect is the increased ability to share information. Quite simply, an envi-

[20] While our examples have mostly been about created environments, found or natural environments (and hybrid environments) can also affect epistemic phenomena. For example, it is easier to share information verbally next to a peaceful lake than in the midst of a hurricane, and a forest at dusk is less conducive to the formation of correct visual beliefs than the same forest at noon (at least for epistemic agents with stereotypically human perceptual capabilities).

ronment that brings people into more frequent contact with each other, also provides increased opportunity for information sharing (Guptill et al., 2016).

Another important epistemic effect of the agrarian city was the opportunity for specialization in a wide variety of domains. Once individuals developed the ability to grow more food than was necessary for themselves and their families, it was no longer the case that everyone had to focus the majority of their time on food production. This created the possibility for people to develop expertise in new domains and permitted the creation of more sophisticated divisions of labor (Guptill et al., 2016).

Further examples of large-scale environmental changes that had effects on epistemic phenomena include: the invention of the printing press, the creation of systems of rapid post, and the development of the internet. Each of these had significant effects not only on peoples' ability to share reasons and ideas, but also on the democratization of knowledge. Although we do not discuss these examples further, they highlight some ways in which environments extend beyond local physical spaces.

There is clearly a need for further research in branch (A). But studying the effects of environments on epistemic phenomena may seem like work for non-philosophers. Historians offer analyses of how environmental changes like invention of paper and the printing press changed information access (Gunaratne, 2001); environmental and educational psychologists study how features of the environment affect the way humans learn and behave (Weston, 2018); sociologists and race and gender scholars offer analyses of the way that social structures affect educational opportunity (Ferguson, 2003); economists explore whether changes in choice architecture will improve human rationality (Thaler & Sunstein (2008); communications scholars study the way that internet pollution (like misinformation) affects online climates (Phillips & Milner, 2021); and so on. What is the role of the epistemologist here?

First, a general point: philosophers often contribute to empirical fields, and this case need not be different. More specifically, epistemologists can offer insight into the relevant states that environments are affecting. Epistemologists have studied and provided detailed accounts of belief, credence, justification, rationality, understanding, etc. Any investigation into the effect of the environment on these states will be improved by clarity about the nature of such states.

Further, philosophers can help classify the ways that environments can affect epistemic phenomena. For example, it appears that environments can have flaws that roughly correspond to traditional accounts of knowledge. Specifically, environments can fail to be conducive to agents forming justified beliefs, forming true beliefs, or believing at all, as well as forming justified true beliefs that are not accidental. If an environment is filled with misleading evidence, it will be conducive to the formation of false beliefs.[21] In some cases, too much misleading evidence will cause people to not take a position at all ("I don't know what to think!"). If an environment (perhaps a political one) tends to whip people into a frenzy, this may be conducive to the formation of unjustified beliefs (perhaps by increasing the likelihood that people make inferences in support of their favored position regardless of the evidence).[22] Finally,

[21] We take this sort of case to be the primary focus of O'Connor and Weatherall (2020) and Levy (2019).

[22] While their primary interest is moral responsibility, Doris and Murphy (2007) offer the potential example that individuals in combat environments are typically cognitively degraded, and that cognitive degradation

for an example where an environment is (on some accounts) conducive to Gettier violations, we need only consider the classic fake barn case. Imagine that a small town celebrates harvest each year by putting up barn decorations that, from the street, look just like real barns. But since this is a farming town, immediately behind most of the decorations is an actual barn. For those passing through during harvest, such an environment would be conducive to Gettier violations. After all, they may form justified true beliefs like "this town sure has a lot of barns" on the basis of observing the realistic barn decorations. Nonetheless, these justified true beliefs may seem objectionably accidental. Given this, we might find it useful to classify environments according to the ways they affect agents' ability to obtain justified, true, non-Gettierized beliefs.

Moreover, there are plenty of actual examples where philosophers have contributed to understanding the epistemic role of environments. First, consider Elijah Millgram's (2015) claim that the hyperspecialization of our world undermines our ability to be intellectually autonomous. On his view, the environmental fact of hyperspecialization forces people to frequently rely on specialists in fields they have no expertise in. But part of thinking for oneself—of being intellectually autonomous—is (on his view) to be able to come to conclusions on one's own. One should not need to rely on the word of experts, nor should one lack an understanding of the defeasibility conditions of claims that are central to their belief systems. Thus, for Millgram, environmental facts explain why it's difficult for individuals to obtain an epistemic good.

Next, consider Miranda Fricker's (2007) account of hermeneutical injustice.[23] A hermeneutical injustice occurs when a person "has a significant area of [their] social experience obscured from collective understanding owing to" a prejudicial lack in shared resources for social interpretation (p. 158). So, for example, if one is raised in an environment where the concept of sexual harassment is not known or commonly used, then one may have trouble describing or understanding experiences of sexual harassment, thereby producing hermeneutical injustice.

In our view, both Millgram and Fricker are (in these passages) engaging in the first branch of Environmental Epistemology—the study of the effects of environments on epistemic phenomena.[24] This branch, and these questions, are extraordinarily important. But we want to emphasize a limitation of this branch, and thereby note

can involve a person's powers of rationality being disrupted.

[23] Fricker joins a long line of scholars of gender and race who have recognized the epistemic effects that our social environments have on us. Examples include: hooks (1990), Haraway (1992), and Mills (1997).

[24] And of course, there are many other examples, such as: (1) Descartes seems to rely on the environmental facts that god exists and is not a deceiver to argue that the world must be understandable for beings like us (Descartes, 1984). (2) Spinoza claims that environmental conditions that the Israelites experienced in Egypt prevented them from being able to understand god's decrees (Spinoza, 2007, TTP ii.46). (3) Harding (1978 p. 205–206) discusses ways that the communities we occupy can shape our values and influence our inquiries. (4) The literature on higher order evidence and irrelevant influences is often concerned with the effects of environments on beliefs (see Street, 2006 and Schoenfield, 2014, respectively; thanks to an anonymous reviewer). (5) Goldman (1999) likens online environments to fake barn country that have features that impact epistemic goods for individuals. (6) Nguyen (2018) analyzes the sources of epistemic bubbles—social epistemic structures in which other relevant voices have been left out—and echo chambers—social epistemic structures from which other relevant voices have been actively excluded and discredited.

a difference between the two projects. Even if we somehow developed a complete descriptive theory of the interactions between environments and epistemic phenomena, we would still have two further unanswered normative questions: which environments are best for which purposes, and which environments should we build?

To make these questions more vivid, recall our original motivating example of wanting to determine how to evaluate social media platforms epistemically. Even if we had a complete account of the effects they have on epistemic phenomena, we would not thereby know what makes social media platforms epistemically *healthy*, or what platforms it is *permissible* to build. For this reason, we need a complementary branch of research; we need an account of the epistemic norms that govern environments.[25]

Fricker's work (and much more[26]) begins to bridge the gap between these two branches, as hermeneutical injustices make an environment epistemically bad (in one way). However, a more complete bridge is needed—we need a theory of the epistemic norms that govern environments. With that, we turn to introducing branch (B) of environmental epistemology.

4 Specific epistemic environment norms

In our view there are two types of "epistemic environment norms" (EENs): *general* epistemic norms that govern all environments, and *specific* epistemic norms that sometimes arise in specific types of environments. We define epistemic health for environments in terms of these norms: to the extent that environments satisfy the applicable EENs, they are epistemically healthy; environments that fail to satisfy applicable epistemic norms are, to varying degrees, epistemically unhealthy.[27] At a high level, the relevant epistemic norms are those that, when an environment satisfies them, enable the environment to be suitable *for* the epistemic flourishing of the agents that occupy it (even if the agents might fail to flourish for other reasons). These norms govern environments, and thereby have implications for environment designers, builders, maintainers, and residents. At a minimum, if an environment would violate an EEN, then that provides a defeasible reason to not build it. Of course, the strength of these defeasible reasons will depend on the violated norm and extent of violation (see more in Sect. 6).

[25] Of course, evidence about the effects of environments on epistemic phenomena will be essential for philosophers to consider when identifying these norms.

[26] Other examples of work that begins to bridge this gap can be found in Rini (2017) and Record and Miller (2022)'s discussion of norms of communication on platforms. Both suggest platform level changes (*environmental* changes) will help cultivate appropriate norms in users. One further example is Simon's (2010) work which is concerned with the ways that web environments effect epistemic phenomena and suggests that platform designers have obligations to increase transparency of how the environments work in certain cases.

[27] The metaphor of health has often been used for evaluating natural environments (cf. UN environment program, 2021) and living environments (Berg, 2022). We acknowledge that other language could be used instead. For example, some working on social media have relied on metaphors of pollution (Phillips & Milner, 2021) and trash (Frost-Arnold, 2023), though we note the connections between health and pollution in non-epistemic environments.

We introduce and defend specific EENs first. There is a wide range of environments that people may create—coffee shops, concert halls, online chatrooms, dating apps, parks, etc.—all with a variety of reasonable ways of setting them up. But sometimes the type of environment one aims to build provides constraints on constructing the environment and standards for evaluating it. For example, imagine that one sets out to build a baseball stadium. To succeed, one must both build a field for playing baseball and some way for spectators to observe it. If a purported stadium designer built something that failed to meet these conditions—for example a tennis court—and called it a day, they would have failed to achieve their end.

Even if our stadium designer actually succeeds in building a baseball stadium, not all baseball stadiums are equally good. Products can be better or worse at fulfilling their ends. A silly example makes this clear. Imagine a baseball stadium where a wall of water separated the spectators from the players, such that the spectators had to watch the game distorted through the water. While this may be interesting in a number of ways (perhaps artistically) it would clearly be worse—as a baseball stadium—than Fenway Park or Yankee Stadium, for example. Well-constructed and well-maintained baseball stadiums make it easy for spectators to observe the game.

In the same way, environments sometimes have epistemic ends that give rise to epistemic constraints on constructing the environment, and hence *specific* EENs for evaluating it. For example, part of what is required to create an auto-mechanic school is to create an environment where people are enabled to *learn* about auto mechanics. More specifically, creating an auto-mechanic school requires creating an environment that is conducive to the transmission of skills and information relevant to being an auto-mechanic. This may include a suitable location, suitable objects (perhaps cars to practice on), suitable teachers with appropriate knowledge, and so on. If one fails to do these things, then one has failed to build an auto-mechanic school, even if one perhaps built something else instead. As a silly example, if someone built a baseball stadium but claimed it was an auto-mechanic school, then something clearly went awry.

Of course, schools can do a better or worse job of satisfying the relevant epistemic standards. So, for example, a school that failed to properly vet its teachers (e.g., ensuring they have the requisite knowledge and skills) is likely to end up with ineffective teachers. And this will make the school less conducive to the transmission of the relevant information and skills. Similarly, a school that blasted disruptive music over lectures or lacked sufficient heat or air conditioning (McCracken 2022) will not be conducive to student learning. To the extent that the school fails to meet these specific EENs we can say that it is less epistemically healthy than schools with properly vetted teachers and no disruptive music.

More generally, it is essential to recognize that there are *three* loci of epistemic evaluation in cases like this—agents who transmit information, agents who consume it, and the environment in which these interactions occur. When agents fail epistemically, the explanation for the failure may be attributable to any of these three. If a student decided to doodle or daydream instead of listening to their teachers or if the student failed to respond appropriately to the evidence communicated by their teachers, then the failing would be the student's rather than being due to the school being an epistemically unhealthy environment or the teacher failing in their duties as

a teacher or testifier. Similarly, if an appropriately vetted teacher shirked their duties, then the fault lies with the teacher and not the school environment nor the student. But finally, and most importantly for our purposes, students may be foiled in their endeavor to learn by an epistemically unhealthy environment. In fact, they may fail while both they and the teacher are epistemically blameless.

Specific EENs give us a framework for evaluating environments epistemically. Environments have epistemic functions (as imparted by the agents that build, maintain, and occupy them), and they should be well suited to the fulfilment of those functions. To the extent that they are, then they will satisfy the epistemically evaluative norms (or standards) that apply to them, and thereby count as epistemically healthy. To the extent that they are not, they are epistemically unhealthy.

5 General epistemic environment norms

At this point we've argued that the epistemic function of an environment[28] sets epistemic standards for evaluating the environment—specific EENs. However, we must also normatively evaluate the environment's epistemic functions or aims themselves. To see this, note that one might aim to set up an environment that's conducive to converting people to their personal cult. Specific EENs give one the tools to determine whether one has done a good job of designing, building, and maintaining such an environment (i.e., making it conducive to cult conversion). We contend, though, that an effective cult conversion environment would *not* therefore be epistemically healthy, even if it conforms to the specific EENs.

In our view, there are general epistemic environment norms that govern all environments with epistemic functions. These are evaluative norms that determine the conditions that epistemic functions must meet to count as epistemically healthy, and limit what environments may be (epistemically) permissibly designed, built, and maintained.

Consider an analogy to ethics. Sometimes environments give rise to specific moral norms of evaluation. For example, therapists' offices should morally be designed to protect privacy (in parallel with specific EENs), because of the sort of environment therapists desire to create—one that makes patients feel comfortable sharing personal information. If you knew your therapist was live tweeting your session or that the mirror was a 2-way mirror, you'd be unlikely to share freely, thereby defeating the purpose of therapy. Therapists' offices may thus be evaluated by how well they fulfill this moral function of privacy preservation. And therapists arguably have obligations arising from this to create environments that are conducive to private sharing.

But there are also general moral norms that govern all environments, and that therefore govern the way that environments should be designed and operated. So, for example, an environment with the function of non-consensually humiliating and degrading the people who occupy them, would be morally bad. This moral norm applies to all environments, such that a therapist's office designed to non-consensu-

[28] At the end of this section, we discuss the interesting issue of what it means for natural environments to have epistemic functions.

ally humiliate and degrade patients would be a morally unhealthy environment, even if the environment was created to have this function and fulfilled it well. Again, this has implications for action. When building social environments, we must ensure that they conform to both general and (relevant) specific moral norms. In the same way, we maintain that there are general EENs that apply to all environments, and that failure to respect these general norms will lead to epistemically unhealthy environments.

But what are these norms in the epistemic case? We begin by considering one kind of environment these norms must rule out, as it will not only clarify how these norms function in our account, but also serve as proof that general EENs exist. Consider the classic brain in a vat ('BiV') skeptical scenario (Harman, 1973). In this scenario, an agent's brain is hooked up to a super-computer which causes the agent to have experiences that are qualitatively indistinguishable from the sorts of experiences that you and I ordinarily have. For example, they might have the experience of writing a paper, making friends, catching a wave, etc. But the agent is not in fact doing any of those things; the agent is in fact, in a vat hooked up to a supercomputer.

In BiV scenarios, an agent appears to respond to their evidence appropriately, but nonetheless still ends up with false beliefs. For example, they might believe they are surfing because they had the experiences of planning to go to the beach, walking to the beach, getting in the water with their board, and paddling to catch waves. In ordinary environments, forming this belief in response to this evidence would result in the agent successfully believing the truth. But in this deviant environment, the agent believes something false.

The most obvious locus of failure (i.e. deviance) in BiV scenarios is clearly not the agent.[29] We can spell out the scenario in such a way where there is no obvious mistake in the agent's reasoning processes or responses to evidence.[30] Something clearly goes wrong in this scenario—something is deviant—and the obvious candidate is the environment. After all, BiV environments are such that even the ideally rational would end up massively deceived.

A BiV environment is clearly not an epistemically healthy environment, but it does not obviously violate any specific EENs. In fact, it may well fulfill the epistemic function intended by the creator of the BiV environment. Thus, there must be some other norms that explain why this is so—general EENs. Plausibly, the same general EENs would rule out other environments like those designed to make one susceptible to cult conversion, or those designed to systematically produce false beliefs. In light of this, we suggest as a candidate general EEN the following: any environment with the function of massively deceiving its (human) occupants is epistemically unhealthy.[31]

[29] Of course, the skeptic disagrees. But we here ignore skeptical hypotheses.

[30] Certain externalists will disagree (see, Lyons, 2013 e.g.). In a BiV scenario, forming beliefs in response to the evidence from one's senses is not to form beliefs on the basis of a reliable process. For that ordinarily reliable process is not reliable in that scenario. For the moment, we will set this concern aside. But in Sect. 7 we will address it and argue that Environmental Epistemology sheds new light on the internalist/externalist debate over such cases.

[31] Plausibly much more stringent norms can be defended. For example, one might think that a sufficiently lower threshold of deception would count as sufficient for violating a general EEN. It is worth noting that

We take this to be good evidence that such norms must exist. But what more can be said about them? First, whatever the correct account of such norms, they must be permissive of a wide range of epistemic functions. After all, they should be such that coffee shops, schools, concert halls, online chatrooms, dating apps, chemistry labs, parks, etc., may all count as having permissible epistemic functions. Second, they should be influenced by the epistemic good for individuals. We're interested in giving an account of what makes environments epistemically healthy *for* agents like us.[32] Thus, this account should be closely connected to the epistemic good for individuals. Epistemic functions that general EENs permit should be ones that are—at least—consistent with the epistemic flourishing of individuals. Similarly, those that are not permitted will plausibly be those that are inconsistent with the epistemic flourishing of individuals.[33]

What we said in the previous paragraph should be instructive for determining the shape of substantive constraints on what the general EENs are. Plausibly, there are also coherence constraints on general EENs. For example, it is plausible that an environment with inconsistent epistemic functions is epistemically unhealthy.

When describing the moral analogue of such norms, we suggested that it was morally impermissible to create environments that have the function of non-consensually humiliating and degrading the people who occupy them. This suggests that agents—by consenting—can make such environments morally acceptable. We contend that the same idea does *not* apply to the epistemic case. An environment that is ruled out because it has a function that is in violation of the general EENs will still count as an epistemically unhealthy environment, even if people wish to build, maintain, and/or occupy it. There may be perfectly good non-epistemic reasons (in special cases) to build, maintain, and/or occupy such environments (e.g., comedic value), but that does not thereby make the environment epistemically healthy.

At the same time, one may wonder whether the norms we've been discussing are genuinely epistemic norms as opposed to pragmatic or instrumental ones. There are a few things to say on this count. First, general and specific EENs are norms for evaluation, and those evaluations are along what is clearly an epistemic dimension. Moreover, whether environments satisfy these norms is closely tied to whether agents and groups will satisfy the epistemic norms that govern their beliefs and actions. Of course, specific EENs do operate like instrumental norms. Environments have epistemic functions which are often imbued by agents' aims, and an environments epistemic health depends on fulfilling those functions well. However, which specific EENs have normative force depends on the general EENs, and general EENs are not instrumental. Finally, other (plausibly) epistemic norms, like norms of inquiry, also plausibly have instrumental dimensions (Hall & Johnson, 1998; Kelly, 2003; Friedman, 2020).

more stringent norms along these lines will plausibly have implications for the permissibility of building many novel technologies including Deepfake technologies and Large Language Models.

[32] To reiterate a point made previously, this does not mean that we couldn't develop such an account of what makes environments epistemically healthy for different types of agents (e.g. non-human animals).

[33] To the extent that found and natural environments can have epistemic functions, then these norms will apply to them as well. If they have functions, general EENs will apply to them and determine whether the fulfilling those functions well would make the environment epistemically healthy.

One might object that the previous discussions have presupposed that environments have epistemic functions, but one might reasonably wonder whether natural or found environments have such functions. One response is to note that, on our account, epistemic evaluations of environments are relative to a (possibly broad) class of agents, and that relativization can thereby provide specificity about the relevant epistemic functions (e.g., for humans, one is forming true or accurate beliefs). A different, not mutually exclusive response, is to note that it is deeply intuitive to think of natural environments as a proper locus of epistemic evaluation. For example, we can imagine a natural environment that is exactly like a BiV world but where no agent was behind the creation of the BiV environment. It still makes sense to say that this environment is epistemically unhealthy (i.e., it is not conducive to human epistemic flourishing), and plausibly this is because it serves to deceive humans. Natural environments seem like they can be conducive, hostile, or irrelevant to human epistemic flourishing in the same ways that created environments can be, primarily depending on the relevant evolutionary histories and selection processes.[34]

Finally, note that there is a natural relationship between EENs and agent-centric epistemic norms. Agent-centric epistemic norms—e.g. norms of rationality, norms of inquiry, etc.—govern the epistemic acts of agents. They tell us whether agents do well or poorly when they act. Importantly, these norms are (and should be) sensitive to environmental features, as when the demands of rationality depend on the environment (cf. Lyons, 2013, and Sect. 6 below).[35] EENs govern environments. They tell us when environments are epistemically healthy or unhealthy. On our view, EENs are thereby sensitive to the epistemic flourishing of the agents that occupy them. A complete epistemology must give an account of *both* sets of norms.

6 Action-guiding epistemology

Notably, EENs are action-guiding. The epistemic health of an environment is closely tied to people's ability to be epistemically successful in that environment, and epistemic success is often necessary for people to flourish and successfully live their lives as they please. Further, we have significant control over the epistemic health of environments. We design, build, maintain, and occupy environments, and in each capacity have some degree of control over the epistemic health of the environment.

[34] One could build out this response in more detail by appealing to Cummins's (1975) analysis of functions in which ascription of a function to something is to ascribe it a set of corresponding dispositions (p. 758). In our analysis, this would require that environments regularly have dispositions to bring about various epistemic phenomena, which often seems correct. Thus, natural environments would have epistemic functions. Of course, one might object to Cummings's analysis, and so we include these details only in a footnote. More generally, if one rejects the claim that natural environments have epistemic functions, then one may still accept our view as applying to a more restricted set of environments.

[35] Ideal epistemologists should arguably also take this into consideration. It may be that in the same way that we usefully idealize individual epistemic agents in order to answer various questions, we may find it fruitful to idealize the environments they're in. Given our view that norms of cognitive success should be sensitive to facts about the environment, it may be interesting to determine (a) what an ideal epistemic environment is like, and (b) what the correct norms for cognitive success are. We are grateful to an anonymous reviewer for encouraging us to consider this.

Thus, to the extent that we have obligations of beneficence, we will have defeasible obligations to act in such a way as to ensure that those environments we have control over become epistemically healthier. To the extent that we have obligations of non-maleficence, we will have defeasible obligations to not make environments epistemically less healthy.[36]

We currently face a pressing need for guidance about how to address our epistemic failings. Concerns over disinformation, propaganda, conspiracy theories, and echo chambers fill newspapers worldwide. Perhaps partially for this reason, epistemologists have become increasingly interested in studying what (indirect) control we have over our doxastic states. Zetetic epistemology and virtue epistemology both do this to some degree. But this focus has primarily been on the agent. As should now be clear, it is necessary to also focus on the environments in which agents think, reason, believe, testify, and know. From the perspective of asking questions about what we can do about our social epistemic woes, environmental epistemology is especially important because of the extensive control we have over our environments (in contrast to the degree of control we have over our own or others' doxastic states). And by designing and maintaining epistemically healthy environments we may be able to have positive effects on the rationality and knowledgeability of large groups—effects that individuals would be unable to achieve on their own.

We are now in a position to make sense of the critique levied against social media platforms that we initially considered. When we critique a platform for being too conducive to the spread of misinformation, we are critiquing it for being epistemically unhealthy: that is, we are asserting that it fails to satisfy one or more EENs that applies to it. To vindicate this critique, then, we must thus establish that there is such an EEN and that the platform does in fact fail to satisfy it. To the extent that those accusations are correct, then those with control over the design and maintenance of the platform have obligations to make changes to the platform.[37] More generally, the framework of EENs provides the conceptual resources to epistemically evaluate environments in productive ways, even if many details remain for future work.[38]

[36] We also believe that there are genuine epistemic reasons for action that are generated by these norms but lack the space to discuss these issues here (see, Hall & Johnson, 1998; Kelly, 2003; Booth, 2006; and Friedman, 2020 for discussion).

[37] There is, of course, the further question about whether those in charge of the platform will find this convincing. While we'd like to think the answer is yes, it's far from clear. One thing we will say, however, is that even if a platform company were to reject the deeper normative implications of our view, it is hard to deny the pull of specific EENs. After all, these arise as a result of the aims of the platform. So, for example, if someone decided to create an online school, there are clearly better and worse ways of satisfying their aims. Thus, even if they reject the more fundamental picture, we hope they would recognize that they have instrumental reason to do what they've set out to do well.

[38] We close with some additional observations about the relationship between our view and Nguyen (2021). Nguyen provides an account of one way that environments might be hostile to epistemic agents like us. Interpreted through the lens of our view, we would say he identified an important way that some environments are in violation of the epistemic norms that govern them—i.e. epistemically unhealthy. What we have done, is provide a more general normative framework for epistemically evaluating *all* environments. We, thus, think of our projects as friendly companions exploring similar terrain. We suggest that similar things can be said about Mills (1997) and Fricker (2007).

7 Payoffs and directions for future work

Environmental Epistemology has many theoretical and practical implications. First, it provides a framework to guide our efforts to make environments more well-suited to epistemic beings like us. We can make them more understandable, more conducive to information transmission, etc. We can design environments that work for individuals like us, if we recognize that environments can be a locus of epistemic evaluation.

Second, and relatedly, environmental epistemology can (and should) inform our theories of epistemic virtue and our theories of cognitive success.[39] Whatever it is to achieve cognitive success or to possess epistemic virtue surely depends to some degree on the environments that we're in (cf. Pritchard, 2010). After all, epistemic tools that are useful in one environment may not be useful in others, and standards that are satisfactory in one environment may be inadequate in others. To the extent that we care about getting things right, then, the environment should inform our accounts of cognitive success and epistemic virtue.

Of course, what counts as an epistemically healthy environment will also depend on our epistemic capabilities, tools, and resources. Thus, we suggest the following: when things go best, there is an epistemic harmony between individual and environment. Similarly, when our general epistemic theories get things right, they will explain and account for this harmony.

Third, without Environmental Epistemology, it has sometimes been the case that epistemic deviance that ought to have been attributed to the environment was attributed to individuals. Not only is this in error, but taking it into account will likely force us to reconsider many epistemic debates and many epistemic positions. After all, agent centric epistemologies neglect a key locus of evaluation. To see this, consider a long-standing debate between internalists and reliabilists over the new evil demon problem. Here is Lehrer and Cohen (1983) posing the problem as an objection to reliabilism:

> Imagine that, unknown to us, our cognitive processes, those involved in perception, memory and inference, are rendered unreliable by the actions of a powerful demon or malevolent scientist. It would follow on reliabilist views that under such conditions the beliefs generated by those processes would not be justified. This result is unacceptable. The truth of the demon hypothesis also entails that our experiences and our reasonings are just what they would be if our cognitive processes were reliable, and therefore, that we would be justified in believing what we do if the demon hypothesis were true as if it were false (p. 192).

The intuitive problem for the reliabilist is that our evidence seems to be the same in the actual world as the demon world, and thus it seems like our justification status should be the same as well (even though the reliability of the process differs).

[39] By "cognitive success" we mean epistemic success states like the state of being justified, the state of knowing, the state of understanding, etc.

One response that the reliabilist can make to this objection is to grant that we have the same evidence in both worlds, but that some further thing explains the difference in justification. Environmental Epistemology can provide the theoretical resources to ground this response. In some environments, believing *p* when one possesses evidence *x*, *y*, *z* may be highly likely to result in success. But in other environments (e.g. ones where there's an evil demon on the loose), this may not be so. Given that we take getting things right in the environment we're in to place a constraint on our theories of cognitive success, then changes in the environment should result in changes in our accounts of justification.[40]

In light of this, the objection may appear like it has been turned back on its originators. Given that there should be a harmony between environment and justification, doesn't the internalist fail to respect this by claiming that our justification status in both worlds is the same?[41] But the internalist can respond as follows. Our accounts of justification should be sensitive to facts about the environment, but human animals are only so malleable. In the demon world, there is no plausible altering of our justificatory practices that would restore the harmony between a theory of justification and the environment. And, most importantly, the problem in the demon world—the reason the harmony is broken—has nothing to do with the agent. The agent didn't cease to be able to see or hear or to respond to their perceptions appropriately. Instead, the environment has been turned against them.

Importantly, this response puts pressure on a wider range of reliabilist responses to the problem as well. For example, Williamson (2000) has claimed that agents in the demon world have different evidence than those in the actual world, because one's evidence is identical to what one knows. And, because those in the demon world know a lot less (because many of their beliefs fail to be true), they possess different evidence. Thus, when the demon world agent forms beliefs in response to their non-veridical perceptions, they believe what is unsupported by their evidence. But as we pointed out, these failures seem to be most appropriately attributed to the *environment*, while justification is a way of evaluating *agents* epistemically. Thus, accounts such as Williamson's mistakenly count deviant environmental factors against the agent.

Of course, these considerations are far from conclusive. But they do highlight the important implications that Environmental Epistemology may have on traditional epistemic debates. We've long been trying to develop general epistemic theories without reference to a central and important locus of epistemic evaluation. We should expect that the recognition of a new kind of epistemic norm will require us to reconsider our previous theories. And, at the very least, we will have to determine the relationship between EENs and recognized epistemic norms.

A third payoff focuses on social and political epistemology. There are a wide range of pressing epistemic issues that currently confront us: polarization, the spread of conspiracy theories and disinformation, etc. And many of the most important ques-

[40] Another way of putting this: there must be a connection between justification and truth in the relevant environment.

[41] We take Lyons (2013) to be giving voice to an objection like this.

tions and difference-making solutions are environmental.[42] We need to know what exactly social media platforms and governments owe us environmentally: what kinds of epistemic environments do they have obligations to design, build, and maintain?

8 Conclusion

We have argued that we need a new subfield of epistemology to address questions about the epistemic norms that govern environments. We have also defended the existence of epistemic norms that govern environments and have introduced and begun to defend an account of these norms. Specifically, we have argued for a health-based model where conformity to two types of epistemic norms determine the health of environments. General EENs are grounded in the epistemic good for epistemic agents and determine which epistemic functions for environments may count as healthy. Specific EENs arise from the function of the environment and set standards for determining whether an environment satisfies its epistemic functions well. The more an environment conforms to its general and specific epistemic norms, the healthier it is. We have also argued that these norms are action guiding in the following way: to the extent that people have control over the design, build, or maintenance of an environment, they have defeasible obligations to do what they can to ensure its epistemic health.

We recognize that we have presented a relatively programmatic account of the epistemic norms that govern environments. Many complexities have been glossed over in service of arguing for the core points that environments themselves can be objects of epistemic evaluation that are governed by epistemic norms. We do not suggest that the framework of environmental epistemology is somehow finished, nor that we have answered all questions about it. We believe that significant intellectual work remains to be done in order to fill out this framework. Nonetheless, the account provided here enables more focused questions about the epistemic nature, role, design, and evaluation of environments.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s11229-024-04504-3.

Acknowledgements This research was partially supported by a grant from the Office of Naval Research to MH and DD. We are also grateful to Jennifer Rose Carr, Craig Callender, David Brink, and two anonymous reviewers for their many helpful comments on previous drafts of this paper. Many thanks also to the audiences at the 2022 International Association for Computing and Philosophy and UC San Diego's Spring 2022 Reasons and Rationality Reading group.

Declarations

Conflict of interest All authors contributed to all aspects of this paper.

[42] While some individual-level solutions show promise (see, e.g., McGrew & Wineburg 2019; Pavlounis et al.,2021), even these solutions require very specific facts about the environment to hold.

References

Bail, C. (2021). *Breaking the social media prism: How to make our platforms less polarizing*. Princeton University Press.

Barrett, P., Hendrix, J., & Sims, J. (2021). Fueling the fire: How social media intensifies U.S. political polarization–and what can be done about it. *NYU Stern, Center for Business and Human Rights*.

Berg, N. (2022). 04 13). *Here's what it's like to live in a windowless dorm designed by a billionaire*. Retrieved from Fast Company: https://www.fastcompany.com/90740511/heres-what-its-like-living-in-a-windowless-dorm-built-by-a-billionaire.

Booth, A. R. (2006). Can there be epistemic reasons for action. *Grazer Philosophische Studien*, *73*, 133–144.

Christensen, D. (2007). Epistemology of disagreement: The good news. *The Philosophical Review*, *116*(2), 187–217.

Clark, A. (1998). D. Chalmers (Ed.), The extended mind. *Analysis*5817–19.

Cummins, R. (1975). Functional analysis. *The Journal of Philosophy*, *72*(20), 741–765.

Dasgupta, S. (2015). Substantivalism vs relationalism about space in classical physics. *Philosophy Compass*, 601–624.

Descartes, R. (1984). In J. Cottingham, R. Stoothoff, & D. Murdoch (Eds.), *The philosophical writings of descartes*. Cambridge University Press.

Doris, J., & Murphy, D. (2007). From my Lai to Abu Ghraib: The moral psychology of atrocity. *Midwest Studies in Philosophy*, *31*(1), 25–55.

Elga, A. (2007). Reflection and disagreement. *Noûs*, *41*(3), 478–502.

Ferguson, R. F. (2003). Teachers' perceptions and expectations and the Black–White test score gap. *Urban Education 38*(4).

Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press.

Friedman, J. (2020). The epistemic and the Zetetic. *The Philosophical Review*, *129*(4), 501–536.

Frost-Arnold, K. (2014). Trustworthiness and truth: the epistemic pitfalls of internet accountability. *Episteme 11*(1), 63–81.

Frost-Arnold, K. (2023). *Who should we be online? A social epistemology for the internet*. Oxford University Press.

Goldman, A. (1999). *Knowledge in a social world*. Oxford University Press.

Gunaratne, S. (2001). Paper, printing and the printing press: A horizontally integrative macrohistory analysis. *International Communication Gazette*, *63*(6), 459–479.

Guptill, A., Copelton, D., & Lucal, B. (2016). *Food and society: Principles and paradoxes*. Wiley.

Hall, R. J., & Johnson, C. R. (1998). The epistemic duty to seek more evidence. *American Philosophical Quarterly*, *35*(2), 129–139.

Haraway, D. (1992). Otherworldly conversations; terran topics; local terms. *Science as Culture*, *3*(1), 64–98.

Harding, S. G. (1978). Four contributions values can make to the objectivity of social science. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 199–209.

Hari, J. (2022). *Your attention didn't collapse it was stolen*. Retrieved from The Guardian.

Harman, G. (1973). *Thought*. Princeton University Press.

Hills, A. (2009). Moral testimony and moral epistemology. *Ethics*, *120*(1), 94–127.

hooks (1990). *Yearning: Race, gender and cultural politics*. South End Press.

Howell, R. J. (2014). Google morals, virtue, and the asymmetry of deference. *Noûs*, *48*(3), 389–415.

James, W. (1896). In S. M. Cahn (Ed.), *The Will to believe: And other essays in Popular Philosophy*. Longmans, Green and Co.

Kant, I. (1995). What is enlightenment? In M. Perry, *Sources of the Western Tradition* (Vol. II), third edition, 56–57.

Kelly, T. (2003). Epistemic rationality as instrumental rationality: A critique. *Philosophy and Phenomenological Research*, 612–640.

Kelly, T. (2005). The Epistemic significance of disagreement. In J. Hawthorne (Ed.), *Oxford Studies in Epistemooogy* (Vol. 1, pp. 167–196). Oxford University Press.

Lackey, J. (2021). *The Epistemology of groups*. Oxford University Press.

Lehrer, K., & Cohen, S. (1983). Justification, truth, and coherence. *Synthese*, *55*(2), 191–207.

Levy, N. (2019). Due deference to denialism: Explaining ordinary people's rejection of established scientific findings. *Synthese*, *196*, 313–327.

List, C. (2005). Group knowledge and group rationality: A judgment aggregation perspective. *Episteme*, *2*(1), 25–38.
List, C., & Pettit, P. (2011). *Group Agency: The possibility, design, and Status of Corporate agents*. Oxford University Press.
Lyons, J. C. (2013). Should reliabilists be worried about demon worlds. *Philosophy and Phenomenological Research*, *86*(1), 1–40.
Massachi, S. (2022, October 13). *How social media amplifies misinformation more than information* Retrieved from New York Times.
McCracken, J. (2022, September 01). *School's out: As temperatures rise, some students sent home because of lack of AC*. Retrieved from Grist: https://grist.org/buildings/american-schools-not-prepared-for-heat-days/.
McGrath, S. (2009). The puzzle of pure moral deference. *Philosophical Perspectives*, *23*(1), 321–344.
McGrath, S. (2011). Skepticism about moral expertise as a puzzle for moral realism. *The Journal of Philosophy*, *108*(3), 111–137.
McGrew, S. W. (2019). Information, lateral reading and the nature of expertise: Reading less and learning more when evaluating digital. *The Teachers College Record*, *121*(11).
Miller, B. (2015). Why (some) knowledge is the property of a community and possibly none of its members. *Philosophical Quarterly*, *65*(260), 417–441.
Miller, B., & Record, I. (2013). Justified belief in a digital age: On the epistemic implications of secret internet technologies. *Episteme*, *10*, 117–134.
Miller, B., & Record, I. (2017). Responsible epistemic technologies: A social epistemological analysis of autocompleted web search. *New Media & Society*, *19*(12), 1945–1963.
Millgram, E. (2015). *The great endarkenment: Philosophy for an age of specialization*. Oxford University Press.
Mills, C. (1997). *The racial contract*. Cornell University Press.
Nguyen, C. T. (2018). Echo chambers and epistemic bubbles. *Episteme*, *17*(2), 141–161.
Nguyen, C. T. (2021). The seductions of clarity. *Royal Institute of Philosophy Supplement*, *89*, 227–255.
O'Connor, C., & Wetherall, J. (2020). *The misinformation age: How false beliefs spread*. Yale University Press.
Pascal, B. (1941). In T. W. Trotter, & T. M'Crie (Eds.), *Pensées: The provincial letters*. Modern Library.
Pavlounis, D., Johnston, J., Brodsky, J., & Brooks, P. (2021). *The digital media literacy gap: How to build widespread resilience to false and misleading information using evidence-based classroom tools*. CIVIX Canada.
Phillips, W., & Milner, R. M. (2021). *You are here: A field guide for navigating polarized speech, conspiracy theories, and our polluted media landscape*. The MIT Press.
Popper, K. (2002). *The logic of scientific discovery*. Routledge.
Pritchard, D. (2010). Cognitive ability and the extended cognition thesis. *Synthese*, *175*(Knowledge, rationality & action. The nature and scope of information), 133–151.
Record, I., & Miller, B. (2022). People, posts, and platforms: Reducing the spread of online toxicity by contextualizing content and setting norms. *Asian Journal of Philosophy*, *1*(2), 1–19.
Rini, R. (2017). Fake news and partisan epistemology. *Kennedy Institute of Ethics*, *27*(2), 43–64.
Schoenfield, M. (2014). Permission to believe: Why permissivism is true and what it tells us about irrelevant influences on belief. *Noûs*, *48*(2), 193–218.
Shapiro, L. (2007). The embodied cognition research programme. *Philosophy Compass*, *2*(2), 338–346.
Simon, J. (2010). The entanglement of trust and knowledge on the web. *Ethics and Information Technology*, *12*, 343–355.
Sliwa, P. (2012). In defense of moral testimony. *Philosophical Studies*, *158*(2), 175–195.
Sosa, E. (2007). *Apt belief and reflective knowledge, volume 1: A virtue epistemology*. Oxford University Press.
Spinoza. (2007). In J. Israel (Ed.), *Theological-political treatise*. Cambridge University Press.
Street, S. (2006). A darwinian dilemma for realist theories of value. *Philosophical Studies*, *127*(1), 109–166.
Sunstein, C. R. (2017). *#republic: Divided democracy in the age of social media*. Princeton University Press.
Thaler, R., & Sunstein, C. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.
Thorstad, D. (2021). Inquiry and the epistemic. *Philosophical Studies 178*(9), 2913–2928.

UN environment programme, & April (2021 13). *Six reasons why a healthy environment should be a human right*. Retrieved from UN environment programme: https://www.unep.org/news-and-stories/story/six-reasons-why-healthy-environment-should-be-human-right.

Vo, L. T. (2023, June 9). *Vietnamese youtuber is filling information voids with Newsmax and Breitbart*. Retrieved from The Markup: https://themarkup.org/news/2023/06/09/vietnamese-youtuber-is-filling-information-voids-with-newsmax-and-breitbart.

Weston, A. (2018). *Teaching as the art of staging*. Stylus Publishing.

Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.

Zagzebski, L. T. (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge University Press.