




Epistemic norms on evidence-gathering

Carolina Flores¹ · Elise Woodard² 

Accepted: 25 April 2023 / Published online: 17 June 2023
© The Author(s) 2023

Abstract

In this paper, we argue that there are epistemic norms on evidence-gathering and consider consequences for how to understand epistemic normativity. Though the view that there are such norms seems intuitive, it has found surprisingly little defense. Rather, many philosophers have argued that norms on evidence-gathering can only be practical or moral. On a prominent evidentialist version of this position, epistemic norms only apply to responding to the evidence one already has. Here we challenge the orthodoxy. First, we argue that there is no significant normative difference between responding to evidence you have and gathering more evidence. Second, we argue that our practices of epistemically criticizing agents for their poor evidence-gathering indicate the existence of epistemic norms on evidence-gathering. Finally, we show that our thesis has important implications for recent debates about the relationship between epistemic norms and inquiry.

Keywords Evidence-gathering · Epistemic norms · Epistemic blame · Epistemic accountability · Epistemic bubbles

In this paper, we argue that there are epistemic norms on evidence-gathering. Though the view that there are such norms seems intuitive, it has found surprisingly little defense. Rather, many philosophers have argued that norms on evidence-gathering can only be practical or moral (Sect. 1). On a prominent evidentialist version of this position, epistemic norms only tell us about how our beliefs ought to relate to the evidence we have, remaining entirely silent on when and how we ought to gather evidence.

Carolina Flores and Elise Woodard contributed equally.

✉ Elise Woodard
ewoodard@mit.edu

Carolina Flores
floreshc@uci.edu

¹ University of California, Irvine, Irvine, CA, USA

² Massachusetts Institute of Technology, Cambridge, MA, USA

Here we challenge the orthodoxy. First, we consider the central sources of resistance to epistemic norms on evidence-gathering and argue that these do not preclude such norms (Sect. 2). In particular, once we are open to epistemic norms on responding to evidence, we should be open to epistemic norms on gathering evidence. Second, we argue that our practices of epistemically criticizing agents for poor evidence-gathering indicate that there are epistemic norms on evidence-gathering (Sect. 3). Our argument proceeds in two steps. First, we argue that if there is a legitimate practice of epistemically criticizing agents for whether and how they ϕ , then we have reason to think that there are epistemic norms on ϕ -ing. We then argue that there *is* a legitimate practice of epistemically criticizing agents for their evidence-gathering, given our reactions to a wide variety of cases where agents poorly gather evidence. More generally, incorporating norms on evidence-gathering helps us elegantly diagnose and explain a number of common and socially significant epistemic pathologies: epistemic bubbles, confirmation bias, laziness, gullibility, and closed-mindedness. We consider and rebut several objections, on which these agents are failing by evidentialist lights, are only morally or practically criticizable, and our view over-generates epistemic norms (Sect. 4).

Finally, we show that our thesis has important implications for debates about epistemic normativity generally and its relationship to inquiry (Sect. 5). The view we develop provides the basis for a moderate expansion of the epistemic realm, offering a novel argument for the conclusion that many norms governing inquiry are epistemic in nature. Plausibly, we epistemically criticize agents not only for their poor evidence-gathering practices, but for other aspects of how they inquire more generally. Thus, our paper offers an important diagnostic tool for identifying epistemic norms in general. By identifying epistemic norms of inquiry, we can defend epistemology against charges on which it looks myopically focused only the evidence the agent has. In doing so, we can make room for the idea that epistemology has much to say about Western democracies' much-discussed "epistemological crisis," which is largely about aspects of our epistemic life that do not reduce to internal responses to evidence.

1 Resistance to epistemic norms on evidence-gathering

Though the view that we can be epistemically required to gather evidence seems intuitive, the orthodox position is that there are no such epistemic requirements. One source of resistance to epistemic norms on evidence-gathering comes from a putative contrast between epistemic and practical normativity: epistemic norms apply independently of an agent's particular goals or interests, whereas practical norms are conditional upon them. On this view, reasons for gathering evidence are contingent upon an agent's goals, interests, or curiosity, and hence only practical or instrumental.¹ Thomas Kelly elaborates on this contrast:

¹ There are two ways of understanding instrumentality. A norm can be instrumental in virtue of being contingent upon *some* goal—independent of whether the agent has it—or relative to the agent's idiosyncratic goal. When we talk about instrumentality, we mean it in the latter sense, i.e. a norm is instrumental just in case it only applies to agents who have some relevant goal.

The reasons which one has to engage in practices of evidence-gathering and experimentation are instrumental reasons; once the experiments have been performed, however, what it is rational to believe is no longer a matter of instrumental (but rather epistemic) rationality. (Kelly 2003, 635)

The main point is this: our opponents tend to agree that *epistemic* normativity is *not* instrumental, and that norms on evidence-gathering are instrumental. Therefore, they cannot be epistemic.

Others have come to the same conclusion about evidence-gathering on evidentialist grounds. Evidentialism is the view that whether an agent's belief is epistemically justified depends entirely on the evidence that the agent has.² On a dominant version of the evidentialist picture, one's sole epistemic obligation is a synchronic requirement to have one's beliefs match one's actual evidence, while "diachronic questions [about whether and how to gather (more) evidence] are moral or prudential questions rather than epistemic questions" (Feldman 2000, 689). On this picture, evidence-gathering—which includes behaviors ranging from paying minimal attention to your environment to actively researching or seeking out new sources—is of little interest to epistemologists.

More broadly, some theorists take epistemic normativity to include only norms on *states*, not on *actions* (Feldman 2000; Foley 1987; Hedden 2015). Gathering evidence is an action and therefore not covered by epistemic normativity on this picture.³ On this state-centered view of epistemic normativity, epistemology tells us exclusively what states agents ought to be in; it does not tell us what actions they ought to perform to arrive at those states.

Evidentialism is a state-centered view, focused solely on what doxastic states agents ought to have in light of their evidence. But state-oriented pictures more generally are unfriendly to epistemic norms on evidence-gathering. For instance, Brian Hedden, who defends a fully synchronic picture of rationality, claims that "[i]t is no requirement of rationality that you seek evidence unrelated to the things you care about. After all, you are not rationally required to spend the rest of your days reading articles on wikipedia" (Hedden 2015, 188). And Stewart Cohen (a coherentist) claims, "There is no further (non-practical) rational requirement that one seek out further evidence. To argue for such a requirement is to conflate rationality with curiosity" (Cohen 2016, 852). While these philosophers' projects differ, all of them agree that there are no epistemic obligations to gather evidence—only practical or moral ones.

In sum, the orthodox rejection of epistemic norms on evidence-gathering is based on three considerations: the categorical—as opposed to instrumental—nature of epistemic normativity, popular evidentialist commitments, and a commitment to state-centered epistemology more broadly.

² See Feldman and Conee 1985 for a classic defense of evidentialism.

³ Those who prefer reasons-talk to norms-talk would say instead that all epistemic *reasons* are reasons for a state, while reasons to gather evidence are reasons for an *action*.

While others have argued for epistemic assessments of evidence-gathering before, they have failed to engage with the above motivations for thinking that norms on evidence-gathering cannot be epistemic.

Defenses of the epistemic place of evidence-gathering have come from three sources. First, virtue epistemologists have rejected evidentialism on the grounds that it licenses negligence by, for example, allowing agents to avoid gathering evidence (Miracchi 2019; Sosa 2021). However, this will not obviously generate *norms* on evidence-gathering. Second, some have argued that duties to gather evidence derive from a primary epistemic duty to believe truths (Hall and Johnson 1998). We find such claims implausibly strong and dialectically unpersuasive, given the prominence of evidentialism. Finally, the Bayesian literature on I. J. Good's 'value of learning' theorem (Good 1966) and the philosophy of science literature on good scientific practice often argue for the importance of good evidence-gathering, but they tend to focus on the instrumental rationality of gathering evidence given the agent's goals.⁴

In contrast to these approaches, we will defend the existence of epistemic norms on evidence-gathering while granting many of our opponents' starting points. In the next section, we directly address the three main sources of resistance to genuinely epistemic norms on evidence-gathering that we have outlined here. First, we will grant that epistemic norms are not instrumental but argue that non-instrumentality does not preclude epistemic norms on evidence-gathering. Second, we will argue that state-centered epistemology is excessively narrow. Yet once one expands epistemology to count in norms on responding to evidence, there is no obvious reason to stop there; instead we should also posit epistemic norms on evidence-gathering. Third, we will aim to show that epistemic norms on evidence-gathering are in the spirit of evidentialism.

2 Room for epistemic norms on evidence-gathering

2.1 Instrumentality

First, consider the instrumentality worry: that norms on evidence-gathering depend on agents' interests, and therefore cannot be epistemic. The idea seems to be that we could only possibly have duties to gather evidence on topics that we want to learn more about.

But why can't we have duties to gather evidence that are independent of our interests? Some topics might be important independently of our interests and thereby generate duties to gather evidence. More strongly, taking a stance on a question plausibly leads one to acquire duties to regulate one's conduct in responsible ways. Further, it is plausible that these include doing some amount of evidence-gathering. For these

⁴ Maher (1990) casts his argument for why scientists ought to gather evidence in terms of its instrumental rationality given the goals of scientific inquiry. Our opponents will deny that this suffices for an epistemic norm. Myrvold (2012) argues that there is epistemic value to one's belief state in light of learning new information. However, this doesn't show that there are norms on evidence-gathering, since it focuses on the value of a particular state.

reasons, perhaps we have duties to gather evidence merely in virtue of taking a stance on a question, or having a belief about a topic. Or perhaps we have a general, imperfect duty to do enough evidence-gathering on topics on which we have views.

One might worry that such purported norms are too demanding and therefore not real norms. These purported duties seem to be too easily over-ridden. For example, perhaps the beliefs at hand are trivial, or there are more pressing things to do than correctly answer the question at issue.

This worry ignores the possibility of *pro tanto* obligations. Indeed, echoing Chisholm, it is plausible to think that any purely intellectual requirement (i.e. epistemic norm) “is only a *prima facie* duty; it may be, and usually is, overridden by other, nonintellectual requirements, and it may be fulfilled more or less adequately” (Chisholm 1966, 14). Similarly, it ignores that epistemic norms on evidence-gathering might be imperfect duties (Stapleford 2013). Such duties can be fulfilled in many ways, leaving room for agential discretion as to how to discharge it.

Further, it is plausible to think that responding to evidence is also often over-ridden by other demands. As Friedman (2020) has compellingly argued, whether one all things considered ought to allocate effort to responding to evidence on completely trivial beliefs also depends on what else one could be doing and other norms at play. Indeed, as Grimm (2009) notes, these sorts of concerns about triviality motivate worries about the normative force of traditional epistemic norms. This suggests that this worry is of a piece with concerns about epistemic norms generally.

In sum, there is room for norms on evidence-gathering that are not instrumental in the relevant sense. Such norms might appear demanding, but their demandingness appears to be shared by standard epistemic norms, such as norms on responding to evidence.

2.2 State-centeredness

Let’s turn now to the worry that epistemology only includes norms on states, not on actions, and therefore norms on evidence-gathering cannot be epistemic.

We grant that gathering evidence is an action. State-centeredness is indeed incompatible with epistemic norms on evidence-gathering. But we reject state-centeredness. In our view, it leads to an overly narrow epistemology. It precludes even norms on responding to evidence. For instance, if state-centeredness is true, the claim that you should drop a belief that *p* when you acquire decisive evidence that not *p* does not state an epistemic norm, even though we epistemically ought not believe that *p* while having decisive evidence that not-*p*. On the state-centered view, all that epistemology offers are standards on which to say whether a mental state is good or bad (along epistemic dimensions), remaining entirely silent on how we get to those states. Indeed, on this view, epistemology is neutral on the quality of all acts of belief-formation. Such a defanged epistemology is unattractive.⁵

⁵ For additional arguments against a state-centered epistemology, see, for example, Fantl and McGrath (2009), Friedman (2020), Hawthorne and Stanley (2008), and Podgorski (2016). See Gerken and Petersen (2020) for an overview of this debate. Singer and Aronowitz (2022) and Booth (2006) defend the related idea that there can be epistemic *reasons* for action.

Though we reject state-centeredness, it is worth noting that a state-centered epistemology is compatible with evidence-gathering making a difference to epistemology. For instance, perhaps the epistemic standing of beliefs depends on whether the agent's evidence-gathering was sufficiently virtuous (Baehr 2009), not solely on whether the belief matches the evidence that the agent has. This does not yield norms on evidence-gathering, but it yields the state-centered analogue of such norms.

Be that as it may, suppose that you are willing to go along with us and accept that not all epistemic norms are state-centered. This is not a revolutionary move. Most epistemologists are friendly to norms on responding to evidence. And once you do, it becomes hard to motivate accepting only norms on responding to evidence one has. Indeed, even hard-line evidentialists acknowledge that there is no sharp line between evidence out there and evidence one has. As Feldman and Conee (2018) note, "what can be reasonably regarded as the mental evidence that is 'available' to someone, or 'had' by the person, varies considerably. No one specific accessibility relation is uniquely correct" (Feldman and Conee 2018, 81). On any plausible view, the line between evidence one has and evidence one lacks is vague and indeterminate. Correspondingly, it will be hard to distinguish the normative significance of evidence one has from that of evidence one could gather.

More generally, this raises the question: What is the big difference between responding to evidence and gathering evidence? There must be something that justifies the claim that the former is covered by epistemic normativity and the latter is not.

The burden is on our opponent to state such a difference-maker. We will here consider two candidate factors and argue that neither yields the deep normative difference our opponent needs.

The first difference-maker that one might suggest is epistemic access. Our opponent might claim that evidence we have is *transparent*: we are always in a position to know that we have it. By contrast, we are not always in a position to know what evidence is out there. Our opponent claims that this makes a normative difference. Agents are in a position to know how to satisfy norms on responding to evidence but not on gathering evidence. Genuine norms are such that agents are in a position to know how to satisfy them. Hence, there are no norms on evidence-gathering.⁶

Against this line of argument, note how strong it is: if it succeeds, it doesn't preclude only epistemic norms on evidence-gathering, but *any* norms on evidence-gathering. This should give us immediate pause. Beyond this, we will argue that there is in fact no deep difference in epistemic access here.

As Srinivasan (2015) argues, we are often not in a position to know what evidence we have. First, if Williamson (2000)'s anti-luminosity argument succeeds, then we are not always in a position to know what mental states we are in. Given

⁶ See Pryor (2001), Jackson (1991), Gibbard (1990), and Hudson (1989) for uses of this kind of consideration to motivate epistemic internalism.

that having evidence involves being in some mental state, this implies that we are not always in a position to know whether we have some item of evidence. Second, results from psychology suggest that we are often mistaken about our mental states, even in favorable conditions.⁷ This implies that we often are wrong about what evidence we have. Third, introspective experience itself suggests that introspection is difficult. It is often difficult to know what one is feeling, or what credence one has in a given proposition, even after careful consideration (Srinivasan 2015).⁸

Moreover, the view that there are deep differences in actual access to evidence we have and evidence we could gather overstates our ignorance about evidence out there. We are often in a position to know that there is evidence bearing on our beliefs that we could easily gather. For example, we often know that we could check a second news source, ask a different person, or read an academic article. We might not be in a position to know what those additional sources will say. But we are often in a position to know that it would be a good idea to check.

Indeed, as a matter of empirical fact, it looks like our actual access to our own minds isn't better—and might even be worse—than our access to the external world (Schwitzgebel 2011). Hence, we often are in a position to know how to satisfy norms on evidence-gathering, as long as such norms take into account one's evidence about what's out there in the world. Epistemic access does not drive a wedge between evidence one has and evidence one could gather.

At this point, our opponent might protest that there's an easy fix: restrict responding to evidence to cases where one has transparent epistemic access to the evidence. There are two problems with this response. First, this still wouldn't secure the requisite difference between responding to and gathering evidence. As we saw, there are cases where we are in a position to know whether there is relevant evidence in our environment. Second, such a restriction would result in severely myopic epistemic norms, silent on how to organize beliefs where one lacks transparent access to those beliefs. We conclude, then, that facts about epistemic access do not motivate claiming that there can be epistemic norms on responding to evidence but not on gathering evidence.

The opponent of the idea that epistemic norms cover evidence-gathering might suggest a second difference-maker: the involvement of agency. The argument proceeds as follows. Gathering evidence, unlike responding to evidence, is agency-involving. Epistemic norms are not agency-involving. Therefore, norms on evidence-gathering cannot be epistemic. Responding to evidence, by contrast, is in the clear.

The idea that epistemic norms are not agency-involving can be seen as a weakening of the idea that epistemic norms are norms on states. This weakening may be motivated by a desire to accommodate norms on *responding* to evidence. On

⁷ See Gopnik (1993), Carruthers (2011), Schwitzgebel (2006), and Schwitzgebel (2011) for discussion of these results and their philosophical implications.

⁸ Though these last two points are compatible with the claim that we are always *in a position to* know the mental states we are in, they shift the burden of proof to the defender of transparency, who needs to offer compelling explanations of these failures of introspection.

such a view, responding to evidence is not an action. Rather it is something passive: something that just happens to agents, a quasi-inevitability in the face of the force of the evidence. This passive conception of responding to evidence is often in the background in epistemology. Most saliently, it shows up in debates on pragmatism. There, the fact that responding to pragmatic reasons for believing that p requires performing actions to get yourself to believe that p is contrasted with the purported immediacy and effortlessness of responding to evidence.

But this narrative is wrong. Responding to evidence one has is often agency-involving. It is typically a matter of drawing inferences from evidence one has to new or revised beliefs.⁹ Inferring, like gathering evidence, is an active process, something we do. And it is not always automatic or easy. In many cases, drawing appropriate inferences is difficult, requiring summarizing and weighing large bodies of information or creatively seeing the wide-ranging implications of our beliefs. In other cases, drawing correct inferences requires over-riding motivational factors that make it hard for one to accept dispiriting facts (Mandelbaum 2019).

Once we properly acknowledge the ways in which responding to evidence is agency-involving, it doesn't look substantively different from gathering evidence. Gathering evidence might require physical movement: picking up a book, typing something onto a Google search box, or asking a question. But this does not make for a deep difference from the intellectual effort involved in many cases of responding to evidence. Further, gathering evidence can simply require having one's eyes and ears open to what's going on in one's environment. Moreover, there are plenty of cases where responding to evidence requires more of the agent than gathering evidence. Hence, differences in agency are unsuitable for generating epistemic norms on responding to evidence but not gathering it.

Note that, when it comes to both agency-involvement and epistemic access, we do not claim to have shown that there is *no* difference between norms on responding to versus gathering evidence. Rather, insofar as differences exist along the metrics discussed, they will be differences in degree rather than kind. Hence, they will be too weak to justify positing norms on the former but not the latter. Moreover, we have shifted the burden to the opponent of epistemic norms on evidence-gathering to propose an alternative difference-maker. In the absence of such a difference-maker, we should be as open to investigating the possibility of epistemic norms on evidence-gathering as we are open to epistemic norms on responding to evidence.

2.3 Evidentialism

Evidentialists in particular should be open to norms on evidence-gathering. The high-altitude reason for this is that such norms are of a piece with giving pride of place to evidence in belief regulation. For this reason, despite evidentialist opposition (Sect. 1), the view that there are norms on evidence-gathering is a natural

⁹ Immediate perceptual beliefs may be an exception: some hold that they are formed by *endorsing* the content of one's perceptual experiences. But most of our beliefs are not immediate perceptual beliefs, so this won't help our opponent.

consequence of what makes evidentialism appealing in the first place: namely, the intuitive idea that evidence should take the driver's seat in belief regulation.

Indeed, evidentialism, narrowly conceived, can accommodate the existence of epistemic norms on evidence-gathering. *Narrow Evidentialism* is only committed to some version of the claim that you ought to believe what your evidence supports. But this is compatible with the existence of additional norms on evidence-gathering. For example, consider Feldman's evidentialist thesis:

For any person S , time t , and proposition p , if S has any doxastic attitude at all toward p at t and S 's evidence at t supports p , then S epistemically ought to have the attitude toward p supported by S 's evidence at t . (Feldman 2000, 679)

This claim is compatible with the view that there are additional epistemic norms on evidence-gathering, so that S ought to adjust their beliefs to the evidence they have *and also* gather relevant evidence. In other words, the existence of epistemic norms on evidence-gathering does not entail the falsity of Narrow Evidentialism, but only its incompleteness. The evidentialist goes wrong if they decide, as many do (Sect. 1) to add to Narrow Evidentialism a completeness clause: "...and there are no other epistemic norms."

Instead, evidentialists ought to add norms on evidence-gathering to norms on responding to evidence. This results in an attractive two-stage view: for proper epistemic conduct, you must satisfy both (1) norms on evidence-gathering and (2) norms on responding to evidence.¹⁰ While we are open to the claim that evidentialists are getting *one* dimension of epistemic assessment right, we disagree that evidentialist norms exhaust epistemic normativity.

3 Epistemic practices and epistemic norms

In this section, we argue that our practices of epistemically criticizing agents give us positive reason to think that there are epistemic norms on evidence-gathering. In brief, we offer the following argument:

1. **Practice Indicates Norms:** If there is a legitimate practice of epistemically criticizing agents for whether and how they ϕ , then we have reason to think that there are epistemic norms on ϕ -ing.
2. **Criticizability for Evidence-gathering:** There is a legitimate practice of epistemically criticizing agents for whether and how they gather evidence.
3. **Evidence-gathering:** We have reason to think that there are epistemic norms on gathering evidence.

¹⁰ See Hughes (2021b) for a version of this view.

3.1 In defense of premise 1: practice indicates norms

Here we will defend **Practice Indicates Norms** by drawing on the role of norms in our epistemic lives. We argue that we can gain traction on epistemic norms based on our responses when they are violated.¹¹

Practice Indicates Norms claims that our legitimate practices of epistemically criticizing others for ϕ -ing provide *evidence* for thinking that there are epistemic norms on ϕ -ing. We are not claiming that what *makes* it the case that there is an epistemic norm at play is that one can be epistemically criticized for violating it if one lacks an excuse.¹² Rather, we are offering a diagnostic. The idea is that we can get a grip on whether there is a norm at play, and whether it is an *epistemic* norm, by looking to our reactions when it is violated. This is unsurprising, given that norms must bear some connection to our actual or possible practices of interpersonal regulation.¹³

For a practice to provide evidence of distinctively epistemic norms, the kind of criticizability involved must be distinctively epistemic, as opposed to, for example, moral. We want the type of criticism to match its target. For example, it wouldn't be appropriate to *morally* criticize someone for failing to act in their self-interest, unless the latter also constituted a moral failing. Similarly, it wouldn't be appropriate to morally criticize someone for a purely epistemic mistake, such as failing to believe in accordance with one's evidence, unless the latter also constitutes a moral failing (Clifford 1877). The central question then becomes: what marks epistemic criticism as *distinctively* epistemic?¹⁴

In recent years, philosophers working on epistemic blame and accountability have made progress on that question. For example, Antti Kauppinen argues that a distinctive way in which we hold agents epistemically *accountable* is by reducing epistemic trust in them—either on a specific topic or more generally (Kauppinen, 2018). Having one's credibility score reduced functions as a sanction because the less trustworthy or credible you seem, the less others will rely on you or participate in joint

¹¹ The motivations behind "Practice Indicates Norms" may generalize beyond epistemic norms. A tight connection between norms and criticizability is also argued for by Kelp and Simion (2021), who effectively argue for the opposite direction. They claim:

For any performance type, ϕ that is governed by a rule to the effect that one must: ϕ only if one meets C , if an agent has legitimately been criticised for ϕ -ing without satisfying C , a certain kind of response on the part of the defendant is *prima facie* appropriate, including (i) explanations invoking overriding norms, excuses appealing to (blameless) (ii) lack of control or (iii) ignorance, and (iv) apologies. (Kelp and Simion 2021, 24)

We are open to both ways of strengthening our premise, though we need not rely on those here. Thanks to an anonymous referee for these suggestions.

¹² See Kauppinen (2018) for a defense of this stronger claim. His claim is compatible with our own and may provide an explanation for the evidential connection we are proposing, but there may be other explanations.

¹³ Cf. Darwall (2014), who emphasizes the relationship between norms and accountability.

¹⁴ We use the term 'epistemic criticism' as a placeholder for various notions of criticizability that have been discussed in the literature, including epistemic blame and epistemic accountability. At times, we use criticizability and accountability interchangeably, except where the difference between blame and accountability is important for exegesis.

inquiry with you.¹⁵ Reductions in trust can take forward- and backward-looking forms: we might decrease our trust in claims that someone previously made, or we may reduce our trust in (some of) their claims going forward. Kauppinen argues that this way of holding others accountable is distinctively epistemic: the reactions and attitudes involved differ from those characteristic of moral blame or accountability. Moral blame does not characteristically involve reductions in epistemic trust. If you lie to me, I may blame you, but I won't think of you as a worse believer (Kauppinen 2018, 11).

The idea that epistemic criticizability characteristically involves epistemic trust-reductions finds further support in Cameron Boulton's recent work on epistemic *blame* (Boulton 2021a, b). According to Boulton, epistemic blame consists of a distinctive type of relationship modification, characteristically, suspending the presumption of trust in someone when it comes to testimony—at least with respect to some domain. Boulton emphasizes the importance of *judging* that someone is blameworthy or that they have impaired the relationship. More generally, we suggest that what matters is not just *that* we reduce epistemic trust, but also the justification we would offer for that reduction. For example, we might cite the fact that these agents poorly exercise their epistemic agency, or are less likely to contribute knowledge, understanding, or other epistemic goods.¹⁶ Boulton understands these adjustments within a more general framework of intention- and expectation-modification: we might epistemically blame others *without* reducing trust. For example, we might instead *reaffirm* our intentions to withhold epistemic trust from someone within some domain.

One might worry that this account is circular. Specifically, if the reduction in trust that marks the violation of epistemic norms is a reduction in distinctively epistemic trust, we cannot answer the question as to how to individuate epistemic norms by appealing to reductions in trust. This objection would have force if we were engaged in the metaphysical project of attempting to elucidate the nature of epistemic normativity. However, we are here only looking for a diagnosis of when there are epistemic norms at play, not for an explanation of what makes such norms epistemic. For the purposes of securing such a diagnosis, all that matters is that we have some independent grip on what constitutes a modification of epistemic trust. And we think it is clear that we have such a grip: we reduce epistemic trust when we are disposed to rely less on others for our beliefs and other doxastic attitudes, and are disposed to justify that reduction by appealing to factors mentioned above. This claim is compatible with multiple accounts of what makes a norm into an epistemic norm, including accounts that center connections to epistemic values.¹⁷

The central claim is that epistemic trust-reductions are a *distinctive* way of *epistemically* criticizing others. It can be fitting to reduce epistemic trust in others in light of their epistemic mistakes but not, for instance, their purely moral or

¹⁵ Being seen as less credible—justly or not—can harm one's capacity as a *knower*. This is a distinctively epistemic harm (Fricker 2007).

¹⁶ Thanks to Hille Paakkunainen and Geoff Sayre-McCord for helpful discussion.

¹⁷ We thank an anonymous referee for raising circularity worries and for prompting us to clarify the nature of our project. See Kelp (2021a) for an account of epistemic norms that centers epistemic values.

prudential mistakes. Hence, in cases where it seems fitting to criticize someone by reducing epistemic trust in them, this suggests that the agent in question is *epistemically* criticizable.¹⁸

There may be other ways to isolate distinctively epistemic criticism. For example, Jessica Brown suggests that we do not demand apologies or compensation for purely epistemic mistakes (Brown 2020). It seems like strategies for epistemic repair can differ from those involved in moral repair, even if there is significant overlap (Woodard 2021). We would not demand recompense from agents who commit purely epistemic mistakes, such as failing to accord their beliefs with the evidence or having (obviously) inconsistent beliefs. In addition, various forms of punishment seem inappropriate. This suggests further practice-based diagnostics for whether an epistemic norm was violated.

Of course, not all epistemic norm-violations justify criticism. First, some norm violations are excused. These include cases where an agent tries their hardest to comply with a norm but fails for reasons outside of their control. (The fact that an excuse is needed is a sign that a norm is at play (Stapleford 2013).) Second, we may withhold criticism when the subject is explicit about not following a norm but is instead relying on heuristics, quick searches, and the like due to time constraints.¹⁹ Hence, it is not required for there to be epistemic norms on ϕ -ing that we *always* epistemically criticize agents who do a poor job of ϕ -ing. What is needed is a *practice* of criticism.

Summarizing so far, our practices of criticizing others for ϕ -ing are indicative of norms on ϕ -ing. We can get a grip on the *type* of norm violated in part based on the types of reactions we have to a norm violation. In particular, epistemic norms are marked by the *availability* of reducing our trust in them in response to these norm violations, at least when they lack an excuse.

To illustrate, consider epistemic norms on responding to evidence you have. Unless they have an excuse, agents who fail to respond to their evidence are epistemically criticizable. Moreover, this criticizability renders it appropriate for you to reduce your epistemic trust in them, at least on some matters. You would not trust or rely on someone's beliefs if they failed to be evidence-responsive in this way. As we'll argue below, these points extend to evidence-gathering.

Finally, we should say something about what makes a practice of epistemic criticism *legitimate*. Minimally, we think that a legitimate practice within the epistemic domain must be reliability-conducive or, more generally, appropriately connected to epistemic goods. Sanford Goldberg justifies a reliability-condition as follows: if the standards for the practice were *not* conducive to agents forming true beliefs, then

¹⁸ On both Kauppinen and Boulton's views, epistemic trust reductions are distinctive of epistemic criticism. This does not entail that reductions in epistemic trust are either necessary or sufficient for epistemic criticizability. As noted, there are other ways to epistemically criticize others beyond reducing trust. Second, reductions in trust only count as a form of epistemic criticism if this reduction of trust is accompanied by a *judgment* of blameworthiness or relationship-impairment. For this reason, reducing our epistemic trust in an agent who claims that p simply because we have more information than them does not amount to epistemically criticizing them. Thanks to Sophie Horowitz for this example.

¹⁹ Thanks to Jane Friedman and Eduardo Martinez for pressing us to consider this point.

they would not help agents achieve their epistemic aims (Goldberg 2018, 170). For our purposes, we are happy to be more neutral about what counts as the appropriate connection or the relevant epistemic goods.²⁰

We need this legitimacy qualification because our practices are an imperfect guide to norms (Simion 2021). This is true not just of epistemic norms, but of moral norms as well. Suppose there is a practice of epistemically criticizing people for relying on the testimony of women or people of color. This practice is clearly problematic and cannot ground a genuine epistemic norm on which we ought not trust the testimony of women or people of color.²¹ We want to rule out such cases by focusing on *legitimate* epistemic practices. Clearly, dismissing testimony from women or people of color is not going to be conducive to epistemic goods.

To sum up, we have argued that practices of epistemic criticism, when legitimate, are a good indicator of epistemic norms. The motivating idea behind this claim is that the type of criticism should match the target. This led to the question: what makes epistemic criticism *distinctive*? Following Boulton and Kauppinen, we argued that reducing our epistemic trust in agents is a distinctive way of epistemically criticizing others, and we illustrated this idea using norms on responding to one's evidence.

3.2 In defense of premise 2: criticizability for evidence-gathering

We'll now turn to premise two, according to which there is a legitimate practice of epistemically criticizing agents for whether and how they gather evidence.

There is a growing list of examples in the literature of cases where agents are intuitively criticizable for their evidence-gathering practices (Hughes 2021; Lackey 2020; Boulton 2021a; Baehr 2009; Miracchi 2019; Sosa 2021). We'll now offer several further cases to support our claim that we have a robust practice of epistemically criticizing agents for their evidence-gathering.

First, consider agents in epistemic bubbles, i.e. "social epistemic structure[s] which ha[ve] inadequate coverage through a process of exclusion by omission" (Nguyen 2020, 2). For example, people who get all their news from one source, such as their carefully curated Facebook News Feed, or a single news channel, are in epistemic bubbles. We also find epistemic bubbles outside of the political domain. Consider Cloistered Claire:

Cloistered Claire: Claire gets all of her nutrition news from Guup, which tends to endorse fad diets that are not always scientifically backed. For example, this month, it encourages its readers to add 1 tbsp of coconut oil to their coffee each day. As it turns out, this is actually a scientifically backed sug-

²⁰ Goldberg (2018, 170) offers a fuller treatment of what renders a practice legitimate. For example, the practice must be ongoing, governed by intuitive standards, and not widely questioned. We think that all of these conditions are met in the case of evidence-gathering, as will become clearer from the examples in the next section.

²¹ Compare Simion 2021 for a similar objection to views that emphasize social practices.

gestion, but Guup does not offer good evidence for it. Claire believes Guup's claim, and she feels no need to check additional sources.

It seems like Claire is criticizable for not gathering more evidence, and perhaps more generally for not gathering evidence from more diverse sources. If Claire were to tell us that we ought to use coconut oil in our coffee each day, and we learned that she gets all of her nutritional information from Guup, we would trust her less on nutrition than if she diligently gathered more evidence. We would also encourage her to seek out additional sources for her nutritional beliefs. This is true even if Guup were a reasonably reliable source. The practice of encouraging people to diversify their sources is best accounted for by thinking that there are epistemic norms on gathering evidence that we are encouraging agents to comply with.²²

Our responses to agents in epistemic bubbles are parallel to responses we have to agents who do a poor job of responding to evidence they have. For example, suppose Claire read additional sources on nutrition, but did a poor job of integrating such sources with her Guup-derived information. In such a case, it would be natural to reduce trust in Claire, at least on nutrition-related topics, to encourage Claire to be more reflective and take into account all her evidence, and to criticize her belief-regulation. In other words, our practices of criticizing agents for poor evidence-gathering mirror those surrounding evidence-responding.

We don't just criticize agents in epistemic bubbles for their poor evidence-gathering. Consider how we respond to gullible agents, who form beliefs too quickly based on a single source, such as:

Gullible Gabe: Gabe tells you that there are 10% fewer jobs in finance this year than there were last year. You defer to him. You later learn that he got this fact from a dated *Economist* magazine that he read at his therapist's office, assuming that it was up-to-date despite the prevalence of dated magazines in therapists' office—something he should know about.

Whether or not finance jobs are on the rise, once you learn about his single evidence source, you will dock credibility points. This doesn't hinge on Gabe being generally gullible. However, once you learn that he is generally gullible and negligent in his evidence-gathering, you will dock even more points and begin to view him as epistemically untrustworthy. This is so even if he is justified in believing as he does based on the evidence he has and is not morally or practically blameworthy.

If you already believe the claim is true, you might remain just as confident but dock credibility points from Gabe, or gently encourage him to improve his

²² This is not to suggest that we *always* ought to diversify our sources; indeed, sometimes relying on one source is fine, if it is sufficiently reliable (though perhaps the agent must also be in a position to recognize this). Rather, our suggestion is that *when* we do encourage to diversify their sources—and criticize them for failing to do so—this can be explained in part by epistemic norms on evidence-gathering. Thanks to an anonymous referee for prompting us to clarify our position. Of course, there is much more work to be done on what norms on diversifying sources might look like, and how they interact with moral and political norms on gathering evidence. We hope to investigate such questions in future work. See Worsnip 2019 for discussion of norms on diversifying sources.

evidence-gathering habits. If Gabe is generally good at gathering evidence and this is just a one-off mistake, you may either think that Gabe has an excuse (e.g. he's under cognitive overload at his therapist's office), or you may hold him epistemically accountable in some other way besides reducing trust, such as by encouraging him to be more epistemically careful or self-aware.²³

Finally, consider our responses to agents who are lazy in their evidence-gathering: when offered additional relevant evidence, they lazily pass on it. For example:

Lazy Larry: Larry is a chemistry major, who forms his beliefs about the structure of the atomic nucleus based on over-simplifying and idealizing diagrams, depicting electrons as marble-like entities that orbit the nucleus in precise tracks. However, this is misleading: electrons actually are spread out diffusely within a massive region. The textbook includes this information, but Larry limits his efforts to just looking at the pictures.

Intuitively, Larry epistemically errs by maintaining such beliefs. Indeed, he is problematically resistant to evidence.²⁴ As a chemistry major, he should know that textbook representations are idealizations, and it would be naive to think that they accurately represent reality. His behavior is negligent: if he had read the textbook, rather than merely relied on pictures, he would know that electrons are massively diffused. Like Gabe, he might be justified in this belief given the evidence he has. Intuitively, it seems that he is epistemically criticizable for his failure to adequately gather evidence bearing on his beliefs, not for failing to respond to evidence he had. Moreover, even if his initial belief was unjustified by evidentialist standards, he is criticizable not *just* for that, as we emphasize in Sect. 4 below. We miss out on an important dimension of epistemic assessment if we fail to recognize that he is epistemically criticizable for his inquiring practices.

These examples illustrate how common it is to epistemically criticize agents for their poor evidence-gathering, whether out of negligence, distraction, laziness, or just as a one-off mistake. Additional cases in the literature support premise two. For example, we epistemically criticize agents who fall prey to confirmation bias (Klayman 1995; Nickerson 1998), such as Hughes' case of 'Likeable Levi,' who unconsciously seeks out evidence that disconfirms his fear that his teammates don't like him and avoids evidence that could confirm it. Intuitively, Levi is criticizable for his biased evidence-gathering practices, which could easily result in acquiring misleading evidence (Hughes 2021).²⁵ We likewise criticize agents who believe that secondhand smoke is not hazardous due to intellectual laziness or hastiness (Baehr 2009). Given Sect. 3.1, our reactions to such cases indicate the existence of epistemic norms on evidence-gathering.

²³ Thanks to Sarah Moss for discussion of cases where we don't reduce trust. Note that it seems like Gabe needs an excuse in order to not be sanctioned. This supports our argument that there is a norm on evidence-gathering at play here.

²⁴ See Simion 2021 and Goldberg 2018 for detailed discussion of cases of evidence-resistance, like Larry's.

²⁵ Compare also Lackey (2020)'s case of the racist who cherry-picks evidence.

Appealing to epistemic norms on evidence-gathering helps us diagnose bad epistemic conduct—ranging from epistemic bubbles, gullibility, laziness, hastiness, and biased inquiry. It also helps us diagnose close-mindedness, given that a central way of being closed-minded—i.e. of failing to engage seriously with relevant intellectual options (Battaly 2018)—is by failing to gather evidence. In addition, it explains what is wrong with agents who rest on their laurels. For instance, agents who engaged in research decades ago, arrived at a settled view and, since then, have not seen the point of keeping up-to-date with more recent findings or arguments are epistemically criticizable.²⁶

Finally, these practices of criticism are legitimate in a similar way that our practices of criticizing agents who fail to respond to the evidence they have are legitimate. Except in particularly non-ideal environments, gathering evidence is generally reliability-conducive. Agents who gather evidence are more likely to have true beliefs, at least in healthy political environments unsullied by disinformation campaigns. But even in cases where this isn't true, gathering evidence will often help agents achieve epistemic goals beyond true beliefs, such as understanding of various aspects of an issue. In short, on plausible accounts of legitimacy, there is legitimate practice of criticizing others for evidence-gathering practices.

4 Objections and replies

In this section, we consider three objections to our argument for epistemic norms on evidence-gathering. The first objection grants that we criticize agents for their evidence-gathering but claims that we can understand such criticism by appealing to violations of norms on responding to evidence. The second argues that we can instead account for our practices of criticism and accountability surrounding evidence-gathering by appealing only to occupational, instrumental, or moral norms. The third objection worries that our account overgenerates epistemic norms. We argue that each objection fails.

Objection 1: Failing by Evidentialist Lights The evidentialist might try to respond to our argument by arguing that these agents are, in fact, doing poorly by evidentialist lights. According to this objection, we criticize these agents for failing to respond to some piece of evidence that they already have—not for their poor evidence-gathering. Perhaps the evidence they have is just *bad* evidence.²⁷ Alternatively, perhaps they are failing to respond to evidence that there is evidence out there. For example, perhaps Larry is failing to respond to evidence that there is additional relevant evidence: in this case, that the textbook has more information. Call this evidence of evidence *further evidence*.

²⁶ In all of these cases, the goal is to explain the badness of epistemic conduct *when it is in fact bad*. This leaves open that there are cases where *not* gathering evidence is epistemically virtuous. When exactly agents ought to gather evidence depends on the exact content of epistemic norms on evidence-gathering.

²⁷ Thanks to Kevin McCain for this suggestion.

Response Let's take each suggestion in turn. First, consider the suggestion that these agents have *bad* evidence. We agree, but the question is whether the evidentialist has resources to say this without appealing to these agents' poor evidence-gathering practices.

Evidentialism claims that agents ought to rationally respond to the evidence they *have*, not that they ought to respond to the *good evidence* they have. Now the question is what counts as evidence. Generally speaking, evidentialism comes in two varieties depending on how one answers this question: internalist and externalist. (This is a spectrum: views can be more or less internalist or externalist.)

The more internalist one goes, the more difficult it will be to claim that these agents are failing to rationally respond to their evidence. For example, if evidence is reducible to sense data of which the agent is consciously aware (Audi 2001; Ayer 1936; Russell 1912; Wedgwood 2002), or more generally to one's conscious mental states (Conee and Feldman 2004), then it's unclear why Claire, Gabe, and Larry fail to count as rationally responding to the evidence they have.²⁸

Moving toward externalist evidentialism may help us deal with this worry. For example, on Simion's recent account of evidence, a fact e is evidence of p for S iff S is in a position to know e , and $P(p|e) > P(p)$ (Simion 2021). The position to know is indexed, in part, to an *average* cognizer. Hence, we might be able to claim that Gabe and Larry *are* in a position to know certain facts that render their beliefs unjustified.

In response, we want to make two points. First, it's not clear that this will work for all cases where we claim that agents violate epistemic norms on evidence-gathering. For instance, it's unclear whether Claire is in a position to know certain facts that render her belief about coconut oil unjustified. It could be that other sources actually *do* support this belief. Second, even if her belief *is* unjustified by externalist evidentialist lights, it's unclear why that exhausts our normative assessment. We should *also* claim that such agents are criticizable for how they gather evidence, not just for how they form their beliefs in light of that evidence. As Friedman (2020, 527) puts it, "Why should epistemology care about what to do with the information we happen to get but not about our getting the information we actually want and need?" Indeed, for agents to correctly respond to facts that they are in a position to know, they first need to become aware of those facts—by gathering evidence.

Similar problems undermine the suggestion that these agents fail to respond to evidence that there is evidence out there. We can imagine versions of each of these cases where the agent does not *have* the requisite further evidence. As we've seen, for many evidentialists, what it takes to *have* evidence is very restrictive. For example, if the evidence you have is what you're currently thinking about (Feldman 1988) or is indexed to your current mental states (Hedden 2015), then the agents we discussed will not count as *having* the relevant further evidence. Despite not having further evidence that there is something problematic or incomplete about their evidence in their particular cases, they are still, intuitively, criticizable.

²⁸ McCain expands this conception of evidence to include beliefs that an agent is disposed to bring to mind when reflecting on the question of whether that belief is true (McCain, 2014, ch. 3). But, we think that the agent can still be criticizable for failing to gather evidence even if they are not so disposed. Indeed, we would criticize our protagonists even if they lack this disposition.

Indeed, our criticism of agents for their poor evidence-gathering extends to cases where it is implausible that the agent has evidence that there is evidence out there. Consider the case of an agent who has been in an epistemic bubble for a long time and gets all their news from a small range of ideologically-aligned sources. They don't know concrete sources outside of the ones they usually consult, and adequately assessing the reliability of such sources from their limited epistemic perspective would be an uphill struggle. Perhaps if they paused and reflected, they would conclude that there likely are other reliable sources that they are not aware of. However, as we've noted, this will not always be the case. Moreover, this does not mean—by most evidentialists' lights—that they are doing poorly *now* by omitting to reflect long enough to have this further evidence. By contrast, our view accurately diagnoses what goes wrong in such cases while also prescribing how to fix it, namely by recommending that the agent acquire such evidence.

Finally, if one accepts Goldberg (2017)'s claim that there are cases where agents should have known that p , then there may be cases where agents are criticizable for not gathering evidence though they were in no position to know that there was such evidence. In particular, where agents should have known that there was such evidence out there (but have no evidence that there is), they will be criticizable for not gathering it. Importantly, beliefs or evidence that an agent *should* have are not identical to evidence that there is evidence out there (Lackey 2020; Goldberg 2018). But once we claim that an agent should *have* this evidence, we're only a short step away from claiming that they should *gather* it. Claims to the contrary drive an implausible wedge between the normative status of states and the normative status of actions necessary for achieving those states.

Objection 2: The Confound Charge Our opponent might deny that these agents are *epistemically* criticizable for their poor evidence-gathering. Rather, they are criticizable only because they violate occupational, practical, or moral norms. First, agents may occupy a role which requires good evidence-gathering (e.g. student, detective, or doctor). In such cases, they can be criticized for their poor evidence-gathering because they fail to meet the demands of that role. Second, sometimes achieving one's goals requires agents to gather evidence on a topic. In other words, agents can be practically irrational for failing to gather evidence by, for example, failing to gather evidence on topics they want to learn about. Third, failing to gather evidence can constitute a moral failing. If we can explain our reactions to agents discussed in Sect. 3.2 via these considerations, then our argument fails.

Response We agree that many cases of failures to gather evidence *also* constitute practical, moral, or occupational failures. But we maintain that agents can also violate epistemic norms when they fail to gather evidence.

As discussed in Sect. 3.1, criticizing others by reducing our trust in them for our beliefs is indicative of epistemic norms. Further, epistemic criticism for evidence-gathering can come apart from occupational, practical, and moral criticism.

First, we criticize agents even when they don't occupy a specific role that demands good evidence-gathering. This is true of Gabe and Claire: we haven't specified any professional or social role they occupy. Perhaps they are accountable for their poor

evidence-gathering because they occupy the role of potential testifiers. But that is an epistemic role; correspondingly, the criticism involved would be epistemic.

Second, we criticize agents for their poor evidence-gathering regardless of their interest in the topic. For example, we epistemically criticize people in epistemic bubbles regardless of whether they have an interest in the topics they rely on the bubble for information about. We don't only hold these agents accountable for failing to act in ways that help them meet *their* goals; we also hold them accountable for acting in ways that make them fail to meet paradigmatic epistemic goals, regardless of whether they care about them. Therefore, some of the norms regulating this practice of accountability aren't norms about how to get what you want, i.e., they aren't norms of practical instrumental rationality. The view that the only reasons for gathering evidence come from one's curiosity about a topic (Sect. 1) misrepresents our practices of accountability.

Third, one can be morally criticizable for gathering evidence but epistemically praiseworthy for doing so. This is the case for excellent gossipers who, according to many, are morally criticizable. However, their excellence at evidence-gathering (into others' lives) makes them epistemically praiseworthy, as reflected in your trust in them when it comes to beliefs about others' lives. Conversely, an agent may be morally praiseworthy for not nosily seeking out evidence about others' personal lives, but we would rely on such people less as sources of information on others' lives.²⁹

In sum, we epistemically criticize others for their evidence-gathering regardless of their interest in the topic, the moral significance of whether they gather evidence, and their social or professional roles. Given our argument in Sect. 3, we ought to think the norms involved here are epistemic.

Objection 3: Over-generation Have we proven too much? Does our view imply, implausibly, that there are epistemic norms on eating sandwiches, getting a good night's sleep, and so on? To see the worry, consider the following case:

Hungry Hank: You and Hank are working together on a project at the architecture firm. Hank has skipped breakfast and is now working through lunch in order to complete his part of the project. He comes back to you with his work, which includes a lot of detailed measurements and calculations. In fact, they are all correct, but Hank is visibly tired and irritable.³⁰

Wouldn't we epistemically criticize Hank for coming to believe that these measurements are correct on an empty stomach? Such a practice of criticism would be

²⁹ There will be cases where failing to gather evidence will involve moral or occupational failures as well. But it doesn't follow that such failures cannot also be epistemic (Goldberg 2018; Lackey 2020). Moreover, insofar as some failures to gather evidence do exhibit moral failings, they will *also* have to exemplify bad epistemic underpinnings—assuming a plausible and widely held epistemic condition on moral blameworthiness (Simion 2021).

³⁰ Thanks to Sophie Horowitz for this case.

reliability-conducive, assuming that people who form beliefs on an empty stomach are generally less reliable.³¹ Perhaps it would promote or instantiate other epistemic goods as well.

While we think that this objection may be a problem for other defenses of epistemic norms of inquiry, we think it leaves ours unscathed.³² The key is to clarify the locus of criticism. We are *not* criticizing Hank for whether and how he eats breakfast or lunch, whether he gets a good night sleep, and so on. Suppose Hank had not engaged in inquiry on complicated topics under these non-ideal conditions, or if he had carefully double-checked or gathered more evidence for his beliefs. (For example, suppose that he had asked a co-worker to double-check his math, or if he re-ran the calculations himself using a calculator.) In such a case, we would not have epistemically criticized him. Insofar as we *do* criticize people for failing to eat a good breakfast, it's for their belief-forming and inquiring practices, including whether they gather more evidence or double-check their beliefs.

Suppose you think that Hank *does* violate an epistemic norm by working on an empty stomach. On our view, the norm that is violated is not one telling you to eat a good breakfast before inquiring, but rather one telling you to gather evidence well. Hank fails to do this in part *because* he doesn't eat a good breakfast, but that is not what we criticize him for. (Compare: suppose an agent is more likely to lie if they didn't eat breakfast. If they lie, the norm they violate is one against lying, not one against failing to eat breakfast.) More generally, there is no practice of epistemically criticizing agents for whether and how they eat breakfast, though there is one of doling out such criticism for how they gather evidence. Hence, our view does not overgenerate epistemic norms on activities such as eating a well-balanced breakfast, getting a good night's sleep, and so on, even though such activities can often causally facilitate good inquiry. You can still be a good inquirer on an empty stomach.

5 The zetetic & the epistemic

Integrating evidence-gathering into epistemology allows for an attractive picture of the epistemically good life. This is a two-stage picture: epistemic normativity requires agents to both gather and respond to evidence in good ways (Hughes 2021). It is part of your epistemic obligations to maneuver in your environment in ways that ensure you make good use of its informational richness—not only to make good use of whatever information is in your head. This point can help us diagnose intuitively bad epistemic conduct in a simple and unified way, helping us account for a wide range of practices of epistemic criticism (Sect. 3.2).

Evidence-gathering, however, is only one component of inquiry more generally. In recent years, philosophers have become increasingly interested in norms governing inquiry and their relationship to epistemic norms (Friedman 2020; Hookway 2006; Hughes 2021; Kelp 2021b; Steglich-Petersen 2021; Thorstad 2021). Are there

³¹ They are, at least, harsher: <https://www.scientificamerican.com/article/lunchtime-leniency/>.

³² For example, it seems to affect Friedman (2020)'s defense, which emphasizes that inquiry promotes epistemic goods. See Thorstad 2022.

epistemic norms governing other aspects of inquiry? The argument we have offered in this paper illuminates this question, motivating expanding the epistemic to cover other norms on inquiry.

Existing arguments for the view that norms of inquiry can be epistemic have been indirect. For example, Kelp (2021b) takes the view that epistemology is the theory of inquiry as a methodological starting point, while Friedman (2020) argues that views to the contrary generate problematic tensions between zetetic and epistemic norms.³³ We offer a novel and direct argument for the existence of epistemic norms of inquiry. The argument is that we epistemically criticize agents not only for their poor evidence-gathering practices, but also for how they inquire more generally. For example, we epistemically assess how agents deliberate, double-check, and confirm their answers. Moreover, these practices seem legitimate. As a consequence, there are epistemic norms on inquiry more generally. Simply by attending carefully to our practices of epistemic accountability, we can see that epistemic norms extend beyond what we do with evidence we have. This argument has the advantage of being metaphysically neutral: detecting that epistemic normativity extends further than previously noticed does not require us to understand what the grounds of epistemic normativity are.

Our view provides a diagnosis for when norms governing inquiry are epistemic (as opposed to practical, for instance). This secures the intuitively appealing view that, though some norms on inquiry are epistemic, not all are. Some norms governing inquiry—such as norms on taking frequent breaks, snacking, or ingesting nootropics—appear paradigmatically practical, even if these activities are conducive to epistemic goods. Our paper not only gives us the resources to argue that some norms of inquiry are epistemic, but also to identify which ones are, based on our practices of epistemic criticism.

On our view, the epistemic and the zetetic are partially overlapping domains. There are norms governing inquiry that are non-epistemic, and epistemic norms that do not govern inquiry. The latter include evidentialist and coherence norms and, more generally, state-centered epistemic norms. By contrast, Friedman (2020) argues for the far more radical view that the epistemic and zetetic are identical normative domains. On our view, it would be too quick to jettison state-centered norms from the purview of epistemology and implausible to re-describe all such norms as zetetic norms. Instead, we should expand our conception of epistemic norms to include both norms on states and on action.

Admittedly, this expansion will lead to conflicts between various epistemic norms. However—and while we cannot argue for it here—we think that looking at many norms of inquiry as akin to *imperfect* duties will minimize direct conflicts between traditional epistemic norms and norms governing inquiry.³⁴

Finally, our view explains the significance of calling a norm ‘epistemic.’³⁵ The debate about whether zetetic norms are epistemic norms can look like a merely

³³ Kelp (2021a) offers a more detailed argument for this view, but one which is nonetheless indirect (cf. p. 139). Moreover, he relies on a teleological picture of epistemic normativity, which we worry many of our opponents will reject.

³⁴ For discussion of imperfect epistemic duties, see Stapleford (2013) and Lackey (2018).

³⁵ See also Thorstad (2022) for helpful discussion of the significance of calling a norm ‘epistemic.’

verbal dispute about how to apply the term-of-art ‘epistemic.’ Against this, we argue that distinctively epistemic norms are tied up with a range of real practices of criticism and accountability that make a difference in how we navigate the world. Epistemic norms license distinctively epistemic reactions and criticism. When we learn that an agent violates an epistemic norm, as opposed to a moral or practical one, this licenses distinctive forms of criticism, such as trust-reduction. Relatedly, our expectations for repair may look importantly different than if the failure were paradigmatically moral or practical. These facts give bite to the question of whether (some) norms on inquiry are epistemic. Specifically, the answer matters for the kinds of practices of criticism and inter- and intra-personal regulation to which inquiry is subject.

Like other philosophers advocating the zetetic turn, we encourage an expansion of our understanding of the epistemic. Including norms of inquiry within the purview of epistemology allows us to appreciate why everyday inquirers should care about epistemology (Friedman 2020). At the same time, there are limits to the epistemic domain. We have not argued that purely instrumental (in the desire-dependent sense) norms are epistemic, nor that a norm is epistemic merely because it promotes epistemic goods, such as true beliefs. Hence, our defense of epistemic norms governing evidence-gathering—and inquiring more broadly—is amenable to preserving the insights of traditional epistemology. This conception of the epistemic does not balloon epistemology into the unmanageable study of how anything affects knowledge or our access to the truth, nor does it give us epistemic obligations to eat sandwiches or read Wikipedia all day. At the same time, it avoids the charge that epistemology addresses only a myopic concern with the sanitized management of one’s internal doxastic life.

6 Conclusion

In this paper, we’ve defended the existence of epistemic norms on evidence-gathering, thereby contributing to the broader zetetic turn within epistemology. Our argument relied on articulating and defending our practices of epistemically criticizing agents for patterns of evidence-gathering. Without epistemic norms on evidence-gathering, epistemic normativity can look myopically focused on what we do with whatever evidence ends up in our ken. We hope to have made some progress in defending epistemology against this charge.

Acknowledgements For helpful discussion and feedback on earlier versions, we are very grateful to: Nathan Biebel, Kevin Blackwell, Sarah Buss, Laura Callahan, Richard Chappell, Tez Clark, Anna Edmonds, Jane Friedman, Alex Guerrero, Michael Hannon, Scott Hershovitz, Nick Hughes, Josh Hunt, Rishi Joshi, Adam Lovett, Annette Martín, Eduardo Martínez, Genae Matthews, Calum McNamara, Alex Meehan, Ram Neta, Eleonore Neufeld, and Ketan Ramakrishnan. We are especially thankful to Jim Joyce, Maria LasonenAarnio, Matthew McGrath, Sarah Moss, Susanna Schellenberg, Lani Watson, Brian Weatherston, and Chris Willard-Kyle for giving us detailed feedback on multiple drafts. For helpful follow-up discussions, thanks to Chris Blake-Turner, Sophie Horowitz, Jonathan Ichikawa, Jon Matheson, and Jonathan Vogel. We also want to thank an anonymous referee for *Philosophical Studies* for helpful comments and suggestions. Finally, we are thankful to audiences and participants at the Inquiry and

its Norms Conference, the 2022 Chapel Hill Normativity Workshop, the COGITO Epistemology Workshop, the 2021 Pacific APA, the Early Career Inquiry Network, the Rutgers Epistemology Workshop, the Michigan Epistemology Workshop, the USC Speculative Society, the Virtues Graduate Seminar at the University of Notre Dame, and Hrishikesh Joshi's Social Epistemology graduate seminar at Bowling Green State University. We have been fortunate to have gathered a lot of evidence on previous drafts, and we hope that the paper has improved in light of it.

Funding Open Access funding provided by the MIT Libraries'.

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

References

- Audi, R. (2001). An internalist theory of normative grounds. *Philosophical Topics*, 29(1/2), 19–46.
- Ayer, A. J. (1936). *Language, Truth and Logic*. London: V. Gollancz.
- Baehr, J. (2009). Evidentialism, vice, and virtue. *Philosophy and Phenomenological Research*, 78(3), 545–567.
- Battaly, H. (2018). Closed-mindedness and dogmatism. *Episteme*, 15(3), 261–282.
- Booth, A. R. (2006). Can there be epistemic reasons for action? *Grazer Philosophische Studien*, 73(1), 133–144.
- Boult, C. (2021a). There is a distinctively epistemic kind of blame. *Philosophy and Phenomenological Research*, 103(3), 518–534.
- Boult, C. (2021b). The significance of epistemic blame. *Erkenntnis*.
- Brown, J. (2020). What is epistemic blame? *Noûs*, 54(2), 389–407.
- Carruthers, P. (2011). *The opacity of mind: An integrative theory of self-knowledge*. Oxford University Press.
- Chisholm, R. M. (1966). *Theory of knowledge*. Englewood Cliffs.
- Clifford, W. K. (1877). The ethics of belief. In *The Ethics of Belief and Other Essays*, Prometheus Books, pp 70–97.
- Cohen, S. (2016). Theorizing about the epistemic. *Inquiry: An Interdisciplinary Journal of Philosophy*, 59(7–8), 839–857.
- Conee, E., & Feldman, R. (2004). *Evidentialism: Essays in epistemology*. Oxford University Press.
- Darwall, S. (2014). Agreement matters: Critical notice of derek parfit, on what matters. *Philosophical Review*, 123(1), 79–105.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an Uncertain World*. Oxford University Press.
- Feldman, R. (1988). Having evidence. In D. F. Austin (Ed.), *Philosophical Analysis* (pp. 83–104). Kluwer Academic Publishers.
- Feldman, R. (2000). The ethics of belief. *Philosophy and Phenomenological Research*, 60(3), 667–695.
- Feldman, R., & Conee, E. (1985). Evidentialism. *Philosophical Studies*, 48(1), 15–34.
- Feldman, R., & Conee, E. (2018). Between belief and disbelief. In K. McCain (Ed.), *Believing in accordance with the evidence* (pp. 71–89). Springer.
- Foley, R. (1987). *The Theory of Epistemic Rationality*. Harvard University Press.
- Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.
- Friedman, J. (2017). Why suspend judging? *Noûs*, 51(2), 302–326.
- Friedman, J. (2020). The epistemic and the zetetic. *Philosophical Review*, 129(4), 501–536.
- Gerken, M., & Petersen, E. N. (2020). Epistemic norms of assertion and action. In S. Goldberg (Ed.), *The oxford handbook of assertion* (pp. 683–706). Oxford University Press.

- Gibbard, A. (1990). *Wise Choices: Apt Feelings, a theory of normative judgment*. Harvard University Press.
- Goldberg, S. (2017). Should have known. *Synthese*, 194(8), 2863–2894.
- Goldberg, S. C. (2018). *To the best of our knowledge: Social expectations and epistemic normativity*. Oxford University Press.
- Good, I. J. (1966). On the principle of total evidence. *British Journal for the Philosophy of Science*, 17(4), 319–321.
- Gopnik, A. (1993). How we know our minds: The illusion of first-person knowledge of intentionality. *Behavioral and Brain Sciences*, 16(1), 1–14.
- Grimm, S. R. (2009). Epistemic normativity. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Epistemic Value* (pp. 243–264). Oxford University Press.
- Hall, R. J., & Johnson, C. R. (1998). The epistemic duty to seek more evidence. *American Philosophical Quarterly*, 35(2), 129–139.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *The Journal of Philosophy*, 105(10), 571–590.
- Hedden, B. (2015). *Reasons without persons: Rationality, identity, and time*. Oxford University Press.
- Hookway, C. (2006). Epistemology and inquiry: The primacy of practice. In S. Hetherington (Ed.), *Epistemology Futures* (pp. 95–110). Oxford University Press.
- Hudson, J. L. (1989). Subjectivization in ethics. *American Philosophical Quarterly*, 26(3), 221–229.
- Hughes, N. (2021). Epistemic feedback loops (or: How not to get evidence). *Philosophy and Phenomenological Research*.
- Jackson, F. (1991). Decision-theoretic consequentialism and the nearest and dearest objection. *Ethics*, 101(3), 461–482.
- Kauppinen, A. (2018). Epistemic norms and epistemic accountability. *Philosophers' Imprint*, 18.
- Kelly, T. (2003). Epistemic rationality as instrumental rationality: A critique. *Philosophy and Phenomenological Research*, 66(3), 612–640.
- Kelp, C. (2021a). *Inquiry, Knowledge, and Understanding*. Oxford University Press.
- Kelp, C. (2021b). Theory of inquiry. *Philosophy and Phenomenological Research*, 103(2), 359–384.
- Kelp, C., & Simion, M. (2021). *Sharing Knowledge: A functionalist account of assertion*. Cambridge University Press.
- Klayman, J. (1995). Varieties of confirmation bias. *Psychology of learning and motivation*, 32, 385–418.
- Lackey, J. (2018). The duty to object. *Philosophy and phenomenological research*, 101(1), 35–60.
- Lackey, J. (2020). Epistemic duties regarding others. In *Epistemic Duties*, pp. 281–295. Routledge.
- Maher, P. (1990). Why scientists gather evidence. *British Journal for the Philosophy of Science*, 41(1), 103–119.
- Mandelbaum, E. (2019). Troubles with bayesianism: An introduction to the psychological immune system. *Mind and Language*, 34(2), 141–157.
- McCain, K. (2014). *Evidentialism and epistemic justification*. Routledge.
- Miracchi, L. (2019). When evidence isn't enough: Suspension, evidentialism, and knowledge-first virtue epistemology. *Episteme*, 16(4), 413–437.
- Myrvold, W. C. (2012). Epistemic values and the value of learning. *Synthese*, 187(2), 547–568.
- Nguyen, C. T. (2020). Echo chambers and epistemic bubbles. *Episteme*, 17(2), 141–161.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2), 175–220.
- Podgorski, A. (2016). A reply to the synchronist. *Mind*, 125(499), 859–871.
- Pryor, J. (2001). Highlights of recent epistemology. *British Journal for the Philosophy of Science*, 52(1), 95–124.
- Russell, B. (1912). *The problems of philosophy*. Home University Library.
- Schwitzgebel, E. (2006). The unreliability of naive introspection. *Philosophical Review*, 117(2), 245–273.
- Schwitzgebel, E. (2011). *Perplexities of Consciousness*. Bradford.
- Simion, M. (2021). Resistance to evidence and the duty to believe. *Philosophy and Phenomenological Research*.
- Singer, D. J., & Aronowitz, S. (2022). What epistemic reasons are for: Against the belief-sandwich distinction. In B. Dunaway & D. Plunkett (Eds.), *Meaning, decision, and norms: Themes from the work of allan gibbard*. Maize Books.
- Sosa, E. (2021). *Epistemic explanations: A theory of telic normativity, and what it explains*. Oxford University Press.
- Srinivasan, A. (2015). Normativity without cartesian privilege. *Philosophical Issues*, 25(1), 273–299.

- Stapleford, S. (2013). Imperfect epistemic duties and the justificational fecundity of evidence. *Synthese*, 190(18), 4065–4075.
- Steglich-Petersen, A. (2021). An instrumentalist unification of zetetic and epistemic reasons. *Inquiry: An Interdisciplinary Journal of Philosophy*.
- Thorstad, D. (2021). Inquiry and the epistemic. *Philosophical Studies*, 178(9), 2913–2928.
- Thorstad, D. (2022). There are no epistemic norms of inquiry. *Synthese*, 200(5), 1–24.
- Wedgwood, R. (2002). Internalism explained. *Philosophy and Phenomenological Research*, 65(2), 349–369.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.
- Woodard, E. (2021). Epistemic atonement. In Shafer-Landau, R., editor, *Oxford Studies in Metaethics*, vol. 18. Oxford University Press.
- Worsnip, A. (2019). The obligation to diversify one's sources: Against epistemic partisanship in the consumption of news media. In C. Fox & J. Saunders (Eds.), *Media ethics: Free speech and the requirements of democracy* (pp. 240–264). Routledge.

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