



Normative Defeaters and the Alleged Impossibility of Mere Animal Knowledge for Reflective Subjects

Giacomo Melis¹

Received: 2 May 2022 / Revised: 13 February 2023 / Accepted: 5 May 2023
© The Author(s) 2023

Abstract

One emerging issue in contemporary epistemology concerns the relation between animal knowledge, which can be had by agents unable to take a view on the epistemic status of their attitudes, and reflective knowledge, which is only available to agents capable of taking such a view. Philosophers who are open to animal knowledge often presume that while many of the beliefs of human adults are formed unreflectively and thus constitute mere animal knowledge, some of them—those which become subject of explicit scrutiny or are the result of a deliberative effort—may attain the status of reflective knowledge. According to Sanford Goldberg and Jonathan Matheson (2020), however, it is impossible for reflective subjects to have mere animal knowledge. If correct, their view would have a number of repercussions, perhaps most notably the vindication of a dualism about knowledge, which would frustrate attempts to provide a unified account of knowledge-attributions to human adults, very young children, and non-human animals. I discuss Goldberg and Matheson's proposal, outline some of the ways in which it is insightful, and argue that it is ultimately unsuccessful because it neglects the inherent temporal dimension of knowledge acquisition. While the article is pitched as a reply to Goldberg and Matheson, its primary aim is to highlight significant connections between the debates on the relation between animal and reflective knowledge, propositional and doxastic justification, and the theory of epistemic defeat.

Keywords Animal and reflective knowledge · Propositional and doxastic justification · Epistemic defeaters · Diachronic epistemology · Sanford Goldberg · Jonathan Matheson

✉ Giacomo Melis
giacomo.melis1@stir.ac.uk; giacmelis@gmail.com

¹ University of Stirling, Stirling, United Kingdom

1 Introduction

Contemporary epistemology is hospitable to the category of unreflective, or animal, knowledge. This is commonly taken to be knowledge that can be obtained by subjects who are unable to formulate thoughts about the epistemic status of their doxastic attitudes. Unreflective knowledge sits well with epistemic reliabilism (e.g. Goldman 1979), theories of knowledge that aim to account for the way in which empirical sciences attribute it to animals (e.g. Kornblith 2014), philosophical accounts of animal rationality (e.g. Glock 2019, Dretzke, 2006, Bermudez, 2003), defences of the view that the source of epistemic normativity lies in agent-independent facts (e.g. Williamson, Forthcoming; Lord 2018), and the mentalist variety of epistemic internalism (e.g. Conee & Feldman 2001, Wedgwood, 2017: ch. 6).

Philosophers who are open to animal knowledge often presume that it complements the reflective knowledge achievable by subjects capable of critically assessing the epistemic standing of their attitudes. Roughly, the idea is that while many of the beliefs of human adults are formed automatically or unreflectively and may constitute mere animal knowledge, some of them—those which become subject of explicit scrutiny or are the result of a deliberative effort—may attain the status of reflective knowledge. Ernest Sosa has been, to date, the most articulate champion of a view along these lines.

According to Sanford Goldberg and Jonathan Matheson (2020), however, it is impossible for reflective subjects to have mere unreflective knowledge—that is, animal knowledge which is not also reflective knowledge. Their claim is noteworthy for several reasons. First, it challenges directly the most developed theory of animal and reflective knowledge (Sosa's). Second, if correct, it may provide support to transformative accounts of rationality against additive accounts (as illustrated, e.g., in Boyle 2016, 2018). Third, and relatedly, it might vindicate a dualism about knowledge, which would frustrate attempts to provide a unified account of knowledge-attributions to human adults, very young children, and non-human animals.¹ Fourthly, given that all but die-hard access-internalists (according to whom, roughly, all sources of knowledge and justified belief must be accessible to the agent on reflection) and advocates of the view that no knowledge is ever the result of reflection or deliberation are in principle open to the distinction between reflective and unreflec-

¹ Roughly, in the framework developed by Boyle, transformative theories of rationality have it that the acquisition of the capacity for reflection altogether transforms one's cognitive abilities. By contrast, additive theories hold that the acquisition of reflective abilities merely supplement perception-induced belief formation, which remains essentially the same in reflective and unreflective subjects. Since the acquisition of knowledge is one of the main achievements of rational agents, we can envisage a parallel framework concerning the relation between knowledge attributed to human adults (reflective subjects) and that attributed to very young children and non-human animals (commonly taken to be unreflective subjects). In a slogan, we might say that if transformative theories are correct, reflective rationality and knowledge differ in kind from their unreflective counterparts; if additive theories are right, the difference may be one of degree. If Goldberg and Matheson are right that mere animal knowledge is impossible for reflective subjects, the transformative camp would gather some support, as one plausible candidate for a rational state common to reflective and unreflective subjects—mere animal knowledge—would have to be ruled out.

tive knowledge, Goldberg and Matheson's suggestion is relevant for most theories of knowledge and justification.

The impossibility of mere animal knowledge for reflective subjects is thus a claim that deserves to be taken seriously. I discuss Goldberg and Matheson's proposal, outline some of the ways in which it is insightful, and I argue that it fails because it neglects the inherent temporal dimension of the acquisition of knowledge and justified belief. In the process, I highlight some connections in the debates on the relation between animal and reflective knowledge, propositional and doxastic justification, and epistemic defeat.

2 The Argument Against Animal Knowledge in Reflective Subjects

Following Goldberg and Matheson (2020: 830), let us agree that "to have reflective knowledge that *p*, the subject's perspective must be such that (from that perspective) the truth of *p* looks non-accidental." The subject who acquires reflective knowledge must meet the Reflection Condition:

(RC) The Subject *S* herself has an epistemic perspective according to which it is not an accident that her belief-forming method produced a true belief on this occasion.

As it will have been noted, (RC) leaves open what might determine that the truth of the subject's belief looks (or, in the good cases, is) not accidental, and what it takes exactly for someone's belief to meet (RC). This is intended, as Goldberg and Matheson aim to provide an argument that holds for all ways of drawing the distinction between animal and reflective knowledge. What matters for present purposes is that reflective knowledge requires the subject to have an epistemic perspective on the relevant belief. Having an epistemic perspective, in the appropriate sense, entails that the subject is in the position to assess the normative standing of the target belief (if only in a pre-theoretical way) and that such assessment would be positive. An important part of the discussion in Sects. 3–6 below will concern what it takes for someone's belief to meet (RC).

(RC) is an internalist condition not required for unreflective knowledge which, by contrast, is conceived as an externalist notion.² A corollary of this is that (RC) can be satisfied regarding both true and false (first-order) beliefs: if (RC) could be satisfied only in relation to true first-order beliefs, its internalist ambitions would be frustrated.³ This reveals that, even if (RC) is advanced as a condition on reflective knowl-

² The relevant notions of internalism and externalism are to be understood in terms of acceptance or rejection of so-called "access-internalism" respectively. The mentalist variety of internalism does not pose a requirement along the lines of (RC); see Bonjour (2010: § 1) for a critique of mentalist-internalism and Wedgwood (2017: 166–7) for a critique of access-internalism.

³ This can be appreciated by revisiting the so-called "argument from illusion", the reconstruction of which I loosely borrow from Wedgwood (2017: 162). Suppose that Sam and his counterpart in another world have the same experiences, apparent memories, seemings, etc., and that they go through the same processes of reasoning, forming, maintaining, and revising exactly the same beliefs in exactly the same

edge, it is best understood as primarily a condition on reflective justification—in the traditional non-factive sense. (RC) becomes a condition on reflective knowledge in so far as reflective knowledge demands reflective justification.⁴ As we'll see shortly, Goldberg and Matheson explicitly endorse the traditional connection between justification and knowledge by arguing that failing to meet (RC) defeats one's knowledge *via* defeating one's justification. Accordingly, the primary focus of this paper will be on (RC) as a condition on reflective justification.

Goldberg and Matheson (2020: § 2) begin their argument by assuming that mere animal knowledge can only be had in cases where (RC) is not met, while the other conditions for knowledge—whatever they might be—are met. They invite readers to suppose the following:

(A) There is a time *t* at which reflective subject *S** has mere animal knowledge that *p*. (831)

They go on to build on the characterization of reflective subjects as subjects capable of taking a stand on the epistemic status of their beliefs:

Since *S** is a reflective subject (at *t*), *S** (at *t*) *does* have a reflective perspective on her own belief-forming methods, including those belief-forming methods responsible for her belief that *p*. (ibid.)

Finally, they suggest that, for such subjects, failure to satisfy (RC) constitutes a defeater for the relevant belief. After all, if *S** were to reflect on matters in a case where (A) holds, she would realize that her belief may be false, or just accidentally true. That, in turn, would defeat her justification for, and knowledge that, *p*—*any* justification and related knowledge, that is, including animal or unreflective ones.

This is, in summary, the argument against the possibility of animal knowledge for reflective subjects. Here's how Goldberg and Matheson (2020: 833) formalize it:

1. If (RC) is not met for *S** with respect to *p*, then *S** has a defeater for any justification *S** has for believing *p*.
2. If *S** has a defeater for any justification *S** has for believing *p*, then *S** does not have animal knowledge that *p*.
3. If (RC) is not met for *S** with respect to *p*, then *S** does not have animal knowledge that *p*. (1, 2).

ways. Yet, as it happens, Sam's counterpart is tricked by an evil demon who ensures that many of his experiences are misleading. The argument from illusion concludes that Sam and his counterpart do not differ with respect to what beliefs they are justified in holding. What matters for us is that saying that they differ with respect to the epistemic perspectives they have on their first-order beliefs (i.e. by saying that only Sam's beliefs meet (RC)) requires making the relevant epistemic perspective external to them just like the truth of their first-order beliefs. And that would frustrate the very point of having (RC) as an internalist requirement. Thanks to an anonymous reviewer for pushing me to clarify this point.

⁴ To put it differently, meeting (RC) plays its epistemic role by contributing to reflective justification, which may support both false and true propositions. It is no accident that Goldberg and Matheson use non-factive terms in introducing (RC) in the passage quoted at the beginning of this section.

4. If (3), then mere animal knowledge is impossible for reflective subjects.
5. Mere animal knowledge is impossible for reflective subjects. (3, 4).

I sympathise with much of this argument. I agree that, in general, unanswered doubts—or, rather, unanswered doubts that are generated by one's epistemic predicament—about the good-standing of one's belief-forming methods work as defeaters for one's (first-order) justification and knowledge (premise 1).⁵ I also find rather persuasive that the defeater in question should affect animal justification and thereby animal knowledge, as well as reflective justification and knowledge (premises 2 and 3).⁶

What I am suspicious about is the presupposed characterization of mere animal knowledge as incompatible with (RC). In other words, I wish to challenge premise (4), which is a consequence of such characterization. The goal of what follows is *not* to present a counterexample to Goldberg and Matheson's account, whereby a reflective subject fails to meet (RC) while having mere animal knowledge. Rather, I aim to show why animal knowledge is possible for reflective subjects, even if there are no counterexamples to the claim that failure to meet (RC) provides one with a defeater for any justification one has for the relevant belief. To anticipate, I will suggest that mere animal knowledge for reflective subjects is possible in cases where (RC) *is* met and the subject has not considered the question whether it is met, thereby challenging Goldberg and Matheson's assumption that mere animal knowledge can only be had in cases where (RC) is *not* met. I will make my proposal by outlining a picture of the normative demands of epistemic justification that strikes me as a plausible development of Goldberg and Matheson's appeal to normative defeaters but which, of course, I do not mean to ascribe to them. As my primary goal is that of exploring the possible implications of Goldberg and Matheson's suggestion, and not merely to refute them, I trust it that engaging with my line of argument will be worthwhile even if it involves an element of independent development of their proposal.

⁵ Some philosophers disagree, and Goldberg and Matheson do offer a helpful discussion of the issue at pp. 830–1 and § 4.2. The reason why I am on board with this claim lies in the link between normative defeaters and the dimension of propositional justification, which will be developed in Sects. 3–6 below. To illustrate at this stage, the relevant doubts typically emerge from ignored evidence that one has, or from misguided assessment of the evidence one has. For example, if due to confirmation bias towards *p*, I fail to notice (or to properly assess) evidence against *p* that I have acquired and thereby fail to question *p*, my belief that *p* is no longer justified. For a nice way of illustrating the relation between epistemic defeat and propositional justification, see Graham and Lyons (2021: 42). Thanks to an anonymous reviewer for pushing me on this.

⁶ If one characterizes animal knowledge in purely reliabilist terms, one might wonder, along lines developed in Beddor (2015), whether the notion of epistemic defeat has genuine application. Two comments in reply to this. First, some reliabilist accounts of epistemic defeat have been recently proposed; see, for example, Constantin (2020) and Graham and Lyons (2021: § 3.5). Second, it is possible to account for animal knowledge in terms of responses to reasons or evidence, provided that such responses do not involve grasp of the concepts of reasons or evidence. The seeds of such accounts may be found in, e.g. Bermudez (2003), Dretzke (2006), Glock (2019). In such accounts, the capacity for animal knowledge involves, at a minimum, the capacity to respond to *overriding defeaters*—that is, evidence that not-*p*. Whether unreflective subjects can also respond to undermining defeaters (which typically suggest that one's warrant for *p* is not in good standing) is an open question. See Sect. 8 below for some comments on the significance of the overriding/undermining distinction. Thanks to an anonymous reviewer for prompting me to clarify this point.

The plan is the following. I will first illustrate how accepting (RC) as a requirement on reflective knowledge while claiming that mere animal knowledge is impossible for reflective subjects raises questions concerning verdicts on what beliefs are justified (Sects. 3–5). I will then go on to argue (Sects. 6–7) that the discussion of the problematic cases illustrated suggests that justified belief, and knowledge, have an inherent diachronic dimension which, in turn, enables us to appreciate why there can be mere animal knowledge for reflective subjects. I conclude (Sect. 8) by briefly exploring how to move the debate forward. But let's proceed a step at a time: my concerns are best introduced by a general discussion of Goldberg and Matheson's proposal and their defence of premise 1.

3 Normative Defeaters and the Reflection Condition

One reaction to the argument proposed by Goldberg and Matheson may be inspired by the consideration that we need to pay the temporal dimension its due. Consider that the time at which S^* begins to reflect, call it t_1 , is an earlier time than the one at which reflection is concluded—call it t_2 . One may thus argue that there will be cases in which S^* has mere animal knowledge at t_1 , but loses it at t_2 when she realizes that (RC) is not met. According to this suggestion, reflecting on the way in which S^* 's belief that p was formed leads to the defeat of the (unreflective) justification for believing p , which S^* possessed at t_1 . This is in line with fallibilism about knowledge,⁷ and it is compatible with the possibility of mere animal knowledge for reflective subjects: S^* did have mere animal knowledge of p at t_1 , but loses it at t_2 after realizing, on reflection, the presence of a defeater.

Goldberg and Matheson take this suggestion to be incompatible with S^* 's merely *having* an epistemic perspective on her belief. This is revealed by what they say about the scenario where S^* does not reflect at all—a scenario relevantly similar to S^* 's situation before she reflects. They argue that the justification for S^* 's belief was already defeated *before* S^* managed to appreciate as much through reflection. They write:

[O]ur opponent will ask us to reflect on cases in which, while it is true that were our reflective subject to reflect on matters she would recognize that (RC) fails to hold in connection with her belief that p , nevertheless she never does reflect on this, and so never does take a perspective on her belief that p or the processes that produced it. Regarding such cases, our opponent invites us to endorse the verdict that the reflective subject has mere animal knowledge. (833)

They comment:

In response, we submit that even in the envisaged type of scenario, it will be true of the reflective subject S^* that the justification for her belief is defeated. In the terms introduced by Jennifer Lackey (1999), our subject will possess a

⁷ As recently defended, for example, in Brown (2018).

normative defeater for her belief whether or not she in fact considers matters. (ibid.)

Following Lackey, normative defeaters are described as propositions that *ought to be* believed by the subject, regardless of whether the subject does so.⁸ In other words, given S*'s epistemic predicament, she ought to believe that (RC) fails to hold (or suspend judgment on whether it does), and that is enough to defeat her (propositional) justification to believe p—even if she, for whatever reason, does not form the belief that (RC) fails to hold.

I sympathize with the appeal to normative defeaters, and I am persuaded that it applies in the way outlined to a wide range of cases. However, accepting that failure to meet (RC) constitutes a normative defeater for reflective subjects while maintaining that mere animal knowledge is impossible for reflective subjects raises questions in relation to whether some beliefs are justified. In the next two sections I will explore some of these questions with the help of toy cases. Doing so will help to motivate the challenge to premise 4 of Goldberg and Matheson's argument, which I will raise in Sects. 6–7.

4 What it Takes to Meet (RC): Norman, Jake, Cristiano

Let us begin with the example of Norman the clairvoyant considered by Goldberg and Matheson (832).

Norman: Norman happens to be a clairvoyant but lacks any reason to believe that he is one. He then finds himself believing that the president is in New York—a belief that seems to merely have “popped into his head”.

Given the very unclear and unusual genesis of his belief, and the lack of reasons to think that he has any special cognitive power, arguably Norman ought to (at least) suspend judgment on whether the truth of his belief was accidental—i.e., on whether (RC) holds. And this is true regardless of whether he does in fact wonder about the epistemic credentials of the belief. I agree with Goldberg and Matheson that, in this scenario, Norman's epistemic perspective provides him with a normative (undermining) defeater for his belief that the president is in New York.⁹

⁸ In Goldberg and Matheson's (834) words: “[n]ormative defeaters do their defeating work in virtue of the subject being such that she is justified in believing or doubting the relevant proposition.” Normative defeaters are contrasted with doxastic defeaters, which are attitudes held by the subject. Graham and Lyons (2021) raise powerful arguments against Lackey's way of drawing the distinction between doxastic and normative defeaters. Their critical discussion of Lackey, however, does not affect a more minimal characterization of the distinction by analogy with the dimension of propositional and doxastic justification. Roughly, just like one can have (propositional) justification to believe p without realizing it, one can have a (normative) defeater for p without realizing it. On the other hand, just like doxastic justification for p entails that one believes p, a doxastic defeater d (for p) requires that one believes d. The relation between propositional justification and normative defeaters will be discussed in Sects. 5–6.

⁹ For a dissenting voice on the suggestion that having justification to suspend judgment about reliability of one's belief-forming methods gives one an undermining defeater for one's beliefs produced by those

Goldberg and Matheson's suggestion may also contribute to explain why some followers of conspiracy theories are blameworthy in embracing conspiratorial beliefs. Consider the following case.

Jake: Jake is an adult with ordinary cognitive and social abilities who happens to be in the thrall of conspiracy theories. Upon following the discussion thread of a web-forum, he forms the belief that several CNN journalists are Satan worshippers who eat children.

Arguably, Jake should have paid more attention to the way in which he formed his belief: *qua* human adult of ordinary cognitive and social abilities (i.e. a reflective agent), he has the epistemic capacity to realize that obscure web-forums are not reliable sources of information, and should have done so. In other words, Jake's epistemic perspective offered him a normative defeater for his belief that CNN journalists are Satan worshippers. That belief does not meet (RC), and Jake ought to have noticed.¹⁰

methods, see Hazlett (2012).

¹⁰ I propose Jake's example as a paradigmatic real-life case of failure to meet (RC) generating a defeater. Yet, one might be tempted to argue that Jake's belief meets (RC), since Jake has an epistemic perspective according to which his belief is non-accidentally true. But if things were so, then virtually every belief of a reflective subject who is unwilling to question her belief-forming methods would meet (RC). This would have the result that, *for such subjects*, Goldberg and Matheson's claim that failing to meet (RC) provides a normative defeater would be vacuous. If we want to take seriously the normative significance of Goldberg and Matheson's suggestion, we have to rule that Jake does not meet (RC). That said, some clarifications on Jake's epistemic predicament may be in order. We might note that it is unlikely that not meeting (RC) is the only normative defeater in Jake's possession which speaks against believing that several CNN journalists are Satan worshippers who eat children (call this belief *Evil Journalists*). While we can suppose—if only for the sake of the argument—that Jake's overall evidence does not include an *overriding defeater* to the effect that no CNN journalists are Satan worshippers who eat children, it may well include *undermining defeaters* which decrease the likelihood of the truth of *Evil Journalists*, given Jake's evidence. Examples might be the considerations that police tend to investigate founded allegations of harm to children and that there are no reports of any such investigations affecting CNN journalists, or that eating babies and worshipping Satan are not practices commonly taken up by members of mainstream institutions, and more. One might think that such considerations are sufficient to defeat Jake's belief without bringing failure of meeting (RC) into the picture. Maybe so. The point here is that realizing that (RC) is not met is a perfectly legitimate way of appreciating that there is something epistemically amiss with believing *Evil Journalists*, which is available to Jake—not that it is the only one. Moreover, arguably, to realize that the undermining defeaters just mentioned speak against believing *Evil Journalists*, Jake has to appreciate that something is wrong with his putative grounds for believing *Evil Journalists* (cf. Melis, 2014; Sturgeon, 2014 on the higher-order dimension of undermining defeaters). That, in turn, would be a way of appreciating that (RC) is not met. Similarly, it may be pointed out that the underminers noted are part of the broader body of evidence which puts Jake in the position to realize that (RC) is not met, and which warrants criticism of his epistemic behaviour. This larger body of evidence includes what an ordinary adult of common cognitive and social abilities—someone who has not been brainwashed, interacts regularly with others, normally acquires information from ordinary sources, and retains some capacity for independent judgment—would be, perhaps with a little effort, in the position to acquire. On top of common-sense claims like the poor likelihood that any journalist eats babies and the link between grounded allegations and ongoing police investigations, such larger body of evidence would include considerations regarding which sources of information are more or less reliable. Presumably, Jake would have had the opportunity to observe that several other claims propagated by obscure web-forums turned out to be false (even if he, at least in his current conspiratorial phase, might be quick to explain them away with the help of additional far-fetched hypotheses). Jake is in the (epistemic) position to notice all this but fails to do so. I thank two anonymous reviewers for urging me to make these clarifications.

Of course, Jake's belief is false, and the case is thus not perfectly analogous to Norman's. Nevertheless, what the two cases have in common is that the presence of a normative defeater determines the epistemic status of the doxastic attitude. Goldberg and Matheson are directing our attention to an important feature of normative theories of belief-revision. Now consider Cristiano.

Cristiano: Cristiano is an ordinary adult human who has never wondered about philosophical problems. In particular, he has never considered scepticism about the external world.

Call *Hands* the proposition that Cristiano has two hands. Does Cristiano know *Hands*? Presumably so. Does he have *reflective knowledge* of *Hands*? The answer is 'no', if by 'reflective knowledge' we mean something in the vicinity of Sosa's notion of knowing full well,¹¹ or anything that might underlie claims of knowledge which take seriously the challenge to one's intellectual conscience posed by scepticism. Yet, since Goldberg and Matheson argue that mere animal knowledge is impossible for reflective subjects, if they wish to allow that Cristiano knows *Hands*, they must ascribe him reflective knowledge. Thus, they must have in mind a fairly undemanding notion of reflective knowledge and corresponding conditions to satisfy (RC).

One way to see how Cristiano's belief may satisfy (RC) starts by observing that, if asked, Cristiano would find it obvious that he knows *Hands*. He would dismiss the idea that he might not know that, and perhaps would even advance some initial considerations in support of his answer. While Cristiano would not be able to rehearse any sceptical argument and possible replies, it does seem reasonable to admit that, unlike Jake, his epistemic perspective on *Hands* satisfies (RC). Goldberg and Matheson might sympathise with this suggestion, as they explicitly deny that their account has any far-reaching sceptical consequences and maintain that "most reflective subjects do meet (RC) most of the time" (839). Taking the plausible view that the conditions for (RC) are easily met, however, raises further issues.

5 What it Takes to Meet (RC): Nora

Consider that Jake, if asked, would reply in a way similar to Cristiano: he would dismiss the suggestion that his conspiratorial belief may have originated from an unreliable source, and would advance some considerations in support of his answer.¹² So, what does warrant the judgment that Cristiano's belief, but not Jake's, meets (RC)?

¹¹ Without going into the details of Sosa's view, roughly, knowing full well requires that the belief in question is formed through the guidance of one's higher-order knowledge that the relevant conditions for (first-order) knowledge are met (cf. Sosa, 2015: 84–87). In discussing another case, Sosa (2015: 79) notes: "in order to know full well, Barney must know that if in his conditions he affirmed that he faces a barn, not easily would he be wrong". This is precisely the sort of knowledge, or even belief, that Cristiano lacks.

¹² It is common for average consumers of conspiracy theories to conceive of their beliefs as being in good epistemic standing (e.g. as being the result of reports of honest truth-seekers, while official sources are epistemically corrupt).

An initially plausible answer is that Jake is wrong in thinking that his belief was formed through an epistemically good method. By contrast, Cristiano is correct in thinking that his belief that he has hands is (not accidentally) true¹³—his rather unarticulated reply to the challenge notwithstanding. This may be too crude a suggestion (see below), but what matters now is that, however we might articulate the details, one's beliefs count as meeting (RC) by conforming to some standards of appropriateness that are not ultimately settled by the agent's own investigation on the matter. That is why we can say that Cristiano's belief satisfies (RC) while Jake's does not.

Presumably, the standpoint from which assessments on whether someone's beliefs satisfy (RC) are made is the standpoint of epistemic normativity: it encompasses what the subject, *given her specific epistemic predicament*, ought to figure out upon reflection or careful investigation. Never mind any *non-epistemic* obstacles that might get in the way of one's cognitive capacities, such as emotional and personal factors that prevent one from reasoning well (as, we suppose, it happens with Jake); what counts are the epistemic features of one's situation. So understood, the relevant standpoint is the standpoint of propositional justification: it determines the propositions one has justification to believe, disbelieve or doubt, and it is the home of normative defeaters.¹⁴ The appeal to the dimension of propositional justification is, I take it, in line with Goldberg and Matheson's proposal. Now consider Nora:

Nora: Nora was inadvertently served a coffee containing a hallucinatory drug, and no one noticed what happened—not even the waiter whose clumsiness resulted in Nora's misfortune.

The perceptual beliefs that Nora forms after drinking the coffee are *not* produced by a reliable method, and yet Nora is not in the *epistemic* position to realize that. Even if she were to launch a full-fledged investigation, she would find no reason to believe

¹³ Assuming that scepticism is false, that is.

¹⁴ Some philosophers (e.g. Conee and Feldman 1985: 86–7, Ichikawa & Jarvis 2013 : 163) characterize propositional justification (or, more modestly, a priori propositional justification, see e.g. Smithies 2015: § 2) as determining the propositions one has justification to believe, *regardless of the cognitive limitations of the specific agent*. This characterization rules that, in some cases, the reflective subject may simply lack the capacity to appreciate whether she meets (RC) in relation to a specific belief. As we have read the cases of Norman, Jake, and Cristiano, their beliefs do *not* fall into such category. This matches with how Goldberg and Matheson characterize reflective subjects as having an epistemic perspective on their beliefs and thereby being in the position to take a view on whether (RC) is met (see, e.g., the following passage at p. 831: “since S* is a reflective subject (at t), S* (at t) does have a reflective perspective on her own belief-forming methods, including those belief-forming methods responsible for her belief that p”). It is also worth noting that the characterization of propositional justification just recalled is not a forced choice. Some philosophers have distinguished a notion of propositional justification that is unconstrained by the subjects' cognitive abilities from one that is so constrained (Coliva 2014: 254, Ichikawa & Jarvis 2013: *ibid.*, Melis 2018: 371), some use the term “propositional justification” to refer exclusively to the notion that is constrained by one's cognitive abilities (Boghossian 2014, Turri 2010), and others explicitly argue against the good-standing of the unconstrained notion (e.g. Goldman 1979, Kornblith 2017). The introduction of Nora's case and the discussion of the importance of the dimension of the subject's epistemic activity which is to follow offer some indirect support to the view that the most fundamental notion of propositional justification is the one constrained by the subject's cognitive abilities.

that she was served anything other than regular coffee: as it happens, there was no CCTV in the café, and no one is in the position to reconstruct what happened.¹⁵

If asked about the epistemic good standing of her perceptual beliefs, Nora would reply along lines very similar to Cristiano and Jake. She would dismiss the very suggestion that her beliefs may not be reliably formed and might bring some *prima facie* considerations in support of her reaction. Yet, Nora's epistemic predicament is different from Cristiano's and Jake's. It differs from Cristiano's because, unlike Cristiano, Nora would be wrong in thinking that her perceptual beliefs are reliably formed; it differs from Jake's because, unlike Jake's, Nora's failure to figure out that her beliefs are not reliably formed is a feature of her *epistemic* circumstances. Does Nora meet (RC)?

Someone who wished to hold on to the earlier suggestion that what determines that Cristiano meets (RC) while Jake does not is that the former is right and the latter is wrong in thinking that the relevant belief-forming method is reliable, should also say that Nora does not meet (RC). This, however, would make actual reliability the standard for assessing whether one's belief meets (RC) and would make it impossible for one's belief to meet (RC) when one is mistaken about the good-standing of the relevant belief-forming method. But if not meeting (RC) is supposed to provide one with a normative defeater (i.e. justification to believe that one's belief may be only accidentally true or just false), this would mean that any misleading evidence to the effect that the relevant belief-forming method is in good epistemic standing be automatically ruled out by one's overall epistemic perspective and thereby excluded from to the set of considerations which determine what propositions one has justification to believe.¹⁶ This comes close to saying that it's impossible to have misleading evidence about the reliability of one's methods, and it strikes me as a problematic result.¹⁷ Moreover, it would make meeting (RC) harder than it needs to be if, as Goldberg and Matheson contend, "most reflective subjects meet (RC) most of the time". Nora and the notorious brain in a vat, for example, would lack justification for their visual beliefs.¹⁸

¹⁵ Detailing the exact conditions that determine that no one is in the position to reconstruct what happened would lead us astray, but I trust it that cases where no ordinary person is in the position to discover an empirical truth are not unusual.

¹⁶ That set of considerations may be called one's *total evidence*: that is, the set of all relevant first-order and higher-order evidence. See Sliwa and Horowitz (2015) for discussion.

¹⁷ Some philosophers (e.g. Ichikawa & Jarvis 2013 ; Smithies, 2015; Titelbaum, 2015) have argued that it is impossible to have misleading evidence about the requirements of rationality. Among other things, these views face the objection of relying on idealized notions of rationality which are too far apart from the dynamism of our epistemic lives (cf. Field, 2020 for some recent discussion). It seems to me that the impossibility of misleading evidence about the good standing on one's belief-forming methods would be even more implausible than the impossibility of misleading evidence about the requirement of rationality: after all, when we're talking about the reliability of one's belief-forming methods, we are focusing on specific fallible agents, and not on demands of rationality that could in principle abstract away from specific agents.

¹⁸ Views according to which justification is factive would happily endorse this result. I disagree, but can't engage with that here. See Brown (2018: ch. 4) for a comprehensive discussion of the main issues arising in relation to these views. I take it that Goldberg and Matheson's characterization of (RC) as the possession of an epistemic perspective *of the subject* is hospitable to the claim that Nora and the brain in a vat are

In the light of the foregoing, I suggest we should concede that Nora *does* meet (RC). Doing so is in line with the suggestion that it is fairly easy, yet not trivial, for one's beliefs to meet (RC). Yet, it carries the burden of having to explain the different verdicts on Nora and Jake vis-à-vis (RC). I think that this can be achieved satisfactorily and consistently with Goldberg and Matheson's suggestion that *not meeting* (RC) provides a normative defeater. However, accounting for the difference between Nora and Jake will also reveal why *meeting* (RC) is *not* sufficient to make one's knowledge or justified belief reflective. This last point will be crucial for my criticism of Goldberg and Matheson's argument.

6 The Importance of Epistemic Activity

Let's focus on the difference between Jake's predicament and Nora's. As we have read Jake's case, his incapacity to realize that his belief about CNN journalists was not formed through reliable means is due to *non-epistemic* factors which got in the way of the full view that his epistemic perspective allowed. By contrast, Nora's not being in the position to realize that her vision is no longer reliable is determined by what is encompassed by the full view of her epistemic perspective. If Nora were to use her cognitive and epistemic powers at her best, she still would find no reason to think that the relevant beliefs may not be formed through methods in good epistemic standing.¹⁹ Nora is held back by her epistemic predicament; Jake is held back by *non-epistemic* factors.²⁰

Now, if we let the standpoint of epistemic normativity—from which it is determined what propositions one has justification to believe, disbelieve and doubt—be constrained by what is within the epistemic reach of the agent, it straightforwardly follows that Jake has justification to believe that (RC) *is not* met, while Nora has justification to believe that (RC) *is* met. And that in turn means that while Jake has a normative defeater for his first-order belief, Nora, like Cristiano, does not. Nora, like Cristiano, has justification to believe that her beliefs are reliably formed: it is just that the reasons she has in support of that happen to be misleading.²¹

justified in holding their first-order beliefs. Indeed, Goldberg (2018: 237) explicitly expresses sympathy for the view that the beliefs of the brain in a vat are justified.

¹⁹ This includes what a highly idealized version of Nora would believe. The problem is that Nora is not in the position to access any evidence suggesting that her belief does not meet (RC).

²⁰ Just to be clear: I don't mean to deny that some advocates of conspiracy theories may be held back by their epistemic predicaments too (someone brainwashed in a conspiracy theory as a child might be an example). I'm only taking it that some of them fit the broad characterization of Jake's case. As real stories of folks who have managed to free themselves of conspiracy theories attest, people do occasionally slip into predicaments like Jake's. Indeed, Jake's case is loosely inspired by the story of Jitarth Jadeja, former follower of QAnon, who publicly discussed his endorsement of views propagated by QAnon. See, e.g. <https://www.washingtonpost.com/technology/2020/10/24/qanon-believer-conspiracy-theory/>.

²¹ Moving at the level of doxastic attitudes, Cristiano may form a justified *true* belief that (RC) is met, while Nora would form a justified *false* belief that (RC) is met.

I do sympathise with this view, and I take it to be broadly in line with Goldberg and Matheson's proposal.²² Yet, it helps to see why the suggestion that mere animal knowledge is impossible for reflective subjects is mistaken. The considerations proposed in support of the claim that Nora meets (RC) while Jake does not suggest that one can only have justification to believe *p* if one is in the (epistemic) position to form the belief that *p* on the basis of available evidence. The picture that emerges is one where the demands of epistemic normativity take into account the dimension of the agents' *use*—or *potential use*—of the reasons or evidence available to them when they form and manage their doxastic attitudes. It is the dimension of doxastic justification, or potential doxastic justification, where agents engage in the activity of forming and revising doxastic attitudes.

But if the dimension of epistemic activity and doxastic justification matters for establishing the demands of epistemic normativity in specific scenarios considered in isolation, it is even more important in understanding the relation between reflective and unreflective *knowledge* (and justified belief). After all, knowledge, unlike (propositional) justification, entails belief-formation. We are now in the position to see what is wrong with Goldberg and Matheson's assumption that meeting (RC) is incompatible with mere animal knowledge.

7 Animal Knowledge for Reflective Subjects

The foregoing discussion reveals that the notion of having an epistemic perspective that figures in (RC) is a static one: it finds its place in the dimension of propositional justification and the demands of epistemic normativity, but it is independent of how an agent acts epistemically. However, the distinction between reflective and unreflective knowledge (just like that between reflective and unreflective justified belief), properly understood, pertains to the level of epistemic activity. This is in line with Sosa's own characterization of the distinction, which has the agents' *performances* at its centre, but it is a more general point. To put it bluntly: something must have gone awry somewhere if verdicts as to whether one has reflective knowledge or reflective justified belief turn out to be indifferent as to whether one does any reflection.

An account of the difference between reflective and unreflective knowledge or justification which considers the dimension of epistemic activity would put the way in which agents form and revise—and ought to form and revise—their attitudes at its

²² I take the proposal just outlined to be an attractive development of Goldberg and Matheson's view because it lines up with: (i) their appeal to normative defeaters (which brings the dimension of propositional justification into the picture); (ii) their characterization of reflective subjects as subjects who have an epistemic perspective on their beliefs and are thus capable of taking a view on whether (RC) is met (which suggests that the verdicts as to whether one has the relevant normative defeater depend on the cognitive abilities of the subject); (iii) their suggestion that (RC) has no far-reaching sceptical consequences (which would emerge if we said that Cristiano fails to meet (RC) and accept the claim that animal knowledge is impossible for reflective subjects). On top of that, I find the proposal outlined appealing because it avoids the problematic consequence of ruling out the possibility of misleading evidence about the reliability of one's methods (which we would have to accept if we endorsed the claim that actual reliability is the standard for assessing whether one's belief meets (RC)). Thanks to an anonymous reviewer for pushing me to clarify this.

centre. Such a view would allow that reflective subjects have unreflective knowledge and justification by holding that the epistemic perspective that they *have* on their beliefs is inert until they go on to reflect on the matter.

It is only when agents actually consider the question whether some belief of theirs is in good epistemic standing that they take advantage of the epistemic perspective that they enjoy and may acquire reflective knowledge (as in Cristiano's case), or reflectively justified belief (as in Nora's case). In other words, the distinction between reflective and unreflective knowledge is inherently diachronic and cannot be flattened at the level of time-slice epistemology of (some characterizations of) propositional justification, where the normative demands concerning what attitudes would be epistemically justified are set.

That, of course, is not to say that the static dimension of propositional justification is not important for assessing whether a subject knows or has formed a justified belief, reflectively or otherwise. Goldberg and Matheson are right to point out that one's (animal or reflective) knowledge may be defeated by unnoticed normative defeaters, as in Norman's case. Meeting (RC) is indeed necessary for reflective justified belief and reflective knowledge.

But in assessing the claim that mere animal knowledge is impossible for reflective subjects we also need to consider whether meeting (RC), in addition to the conditions for animal knowledge, is sufficient for reflective knowledge. And meeting (RC), together with the conditions for animal knowledge, is not sufficient for that: to achieve reflective knowledge, one also needs to reason and appreciate why she meets (RC), if only in some pre-theoretical form. If so, mere animal knowledge for reflective subjects is possible in cases where (RC) is met and the agent has not considered the question whether it is met.

Goldberg and Matheson are right to say that not reflecting cannot preserve one's animal knowledge in cases where one's belief does *not* meet (RC), but they overlooked the importance of reflecting for acquiring reflective knowledge in cases where one's belief *does* meet (RC).

To sum up: I think that Goldberg and Matheson's basic point that failure to meet (RC) provides one with a normative defeater for the relevant (animal or reflective) justification is a good one. This may well be problematic for Sosa's own version of the distinction between animal and reflective knowledge,²³ but it does not generalize to all ways of carving out the distinction. Goldberg and Matheson's insightful observation about the role of normative defeaters does not show that unreflective knowledge is impossible for reflective subjects. Once we appreciate the importance of the dimension of epistemic activity for a full account of epistemic normativity and the difference between reflective and unreflective knowledge or justification, we can see that reflective subjects may have plenty of unreflective or animal knowledge.

²³ Sosa allows that even some *noticed* defeaters may not defeat animal knowledge while they defeat the reflective one, See his take on the Pyrrhonist in Sosa (2015: 208–210; 242).

8 Moving Forward

I have argued that mere animal knowledge and justified belief are possible for reflective subjects. At the same time, I have acknowledged that meeting (RC) is necessary for reflective knowledge and justified belief, but not for their unreflective counterparts. Does this difference mark a dualism about knowledge or a difference in kinds of rational or cognitive achievements?

Unreflective and reflective knowledge and justified belief differ in the way they are acquired or sustained. Roughly, the latter, but not the former, involve an explicit endorsement of one's belief, which might take the form of a positive assessment of the way in which the belief was formed, or of the strength of one's grounds for belief. On the face of it, this difference is compatible with the view that unreflective and reflective justification and knowledge are two manifestations of the same epistemological category. If we keep the focus on knowledge, we can say that both reflective and unreflective knowledge are cases of justified true belief, plus some anti-Gettier condition. If we keep the focus on justified belief, we can say that both unreflective and reflective justified belief instantiate epistemic justification. In this picture, the difference lies in the character of justification. Unreflective knowledge relies on justification whose good standing is not something that the subject evaluates on reflection, be it because she's not in the position to do so as an unreflective subject, or because she simply hasn't done it yet. By contrast, reflective knowledge is sustained by justification which is assessed by the subject. What suggests that the two justifications are instances of the same epistemological category is that both, in different ways, have a link with truth, or with likelihood of truth.²⁴ After all, philosophers who have endorsed internalist and externalist theories of justification along lines that broadly match the unreflective/reflective distinction have taken themselves to be theorizing about the very same thing: *epistemic* justification.

However, it might be pointed out that lack of reflection prevents unreflective subjects from ever understanding why they believe the things they believe. In turn, that may be taken to suggest that *all* epistemic achievements of reflective subjects—whether or not they involve reflection—are radically different from those of unreflective subjects. Philosophers who sympathise with the view that reflection is essential or transformative for rationality such as, e.g., Boyle (2018), Korsgaard (1996), Brandom (2010), McDowell (1994), or Davidson (1982) would likely advocate a position along these lines.²⁵ If they are right, meeting (RC) may be read as a sign of a profound hiatus between reflective and unreflective subjects even if, as I have argued, the former are capable of mere animal knowledge. In this picture, the mere animal knowledge of reflective subjects may be seen as also *potential* reflective knowledge,

²⁴ This is, in rough outline, the underlying monism that some commentators have found in the apparent epistemic pluralism of philosophers like Goldman, Burge, and Alston. For more discussion see, for example, Pedersen (2017: §§ 2–4).

²⁵ Indeed, these philosophers would be reluctant to use the term “rationality” in relation to unreflective agents, and would prefer terms like “intelligence”, thereby explicitly marking the difference in kinds of cognitive achievements. Cf., e.g., Boyle (2018: § 2).

and hence quite unlike the mere animal knowledge of unreflective subjects.²⁶ It thus seems fair to say that, even if Goldberg and Matheson's argument isn't successful, the relation between reflective and unreflective knowledge or justified belief remains an open issue. Among other things, addressing it requires investigating in what exactly animal knowledge which is not accompanied by the satisfaction of (RC) differs from that which is so accompanied.

It seems to me that, in addressing the question of the relation between unreflective and reflective knowledge or justification, it would be advisable to build on the main lesson which emerged from the present discussion. Namely: we should keep the dimension of one's first-person activity—or potential activity—of acquisition and maintenance of a system of justified or rational beliefs at the centre of investigation. And since responding to counterevidence is an essential component of rational belief-revision, it would be a good idea to keep the focus on how subjects acquire and process epistemic defeaters.

Aside from being integral to rational belief-revision, there is another reason why we might want to invest energies in the study of the way in which agents respond to epistemic defeaters. At least if we keep the attention on all-out beliefs—which is itself central to account for the first-person epistemic activity of ordinary agents—there is a potentially significant disanalogy between epistemic defeaters and ordinary positive evidence. While ordinary positive evidence for a proposition *p* always speaks in favour of believing *p*, there are two different ways in which counterevidence may speak against believing *p*. Counterevidence may suggest that one should believe the negation of a proposition *p* that one previously accepted (a case of *overriding* defeat), or it may suggest that one should give up one's belief in *p* without thereby moving towards accepting its negation (a case of *undermining* defeat). Thus, the study of epistemic defeat can be expected to provide a distinctive contribution to the general understanding of epistemic rationality which goes beyond what an exclusive focus on responsiveness to positive evidence might deliver.²⁷

Someone who sympathises with the suggestion of keeping agents' epistemic activity at the centre of the investigation would concur that a focal question for understanding the relation between unreflective and reflective knowledge and justified belief is the following: how does one become a reflective subject? Slightly more precisely: does one's capacity for reflective responsiveness to evidence emerge seamlessly from one's capacity to respond to evidence unreflectively (say, in a way that conserves the latter and adds to it), or does it involve a complete transformation of one's ability to respond to evidence simpliciter? These questions are especially poignant when we consider that there are subjects who make the transition from unreflective to reflective rational agency; namely, human children.

The last observation brings us to what seems to me another important methodological point: developmental psychology and other disciplines studying the mental

²⁶ Note that this is different from what Goldberg and Matheson have argued. According to them "in every case, a reflective subject will either have both animal knowledge and reflective knowledge that *p* or will lack both animal knowledge and reflective knowledge that *p*." (833) They are talking about *actual*, not potential knowledge.

²⁷ See Piazza (2021) for a recent overview of the literature on epistemic defeaters.

processes of presumed unreflective subjects may help. It's not just that, to the extent that the criteria for reflective or unreflective justified beliefs are supposed to apply to the beliefs of ordinary subjects, they'd better be sensitive to what ordinary subjects can actually do. It's also that the philosophical task of identifying the criteria for reflective and unreflective justified belief (together with their relation) may be clarified by familiarity with current studies on the psychological mechanism underlying each. Of course, verdicts on whether one's beliefs are rational or justified are normative evaluations which fall within the remit of epistemology. Yet, empirical research may help to understand the nature of the achievement that needs to be evaluated.²⁸

Since this is largely unexplored territory, I'll resort to a quick example from my own work in progress to illustrate the point. Engaging with some developmental work on different levels of abstraction involved in the belief-revision of young children (e.g. Kimura & Gopnik 2019) has alerted me to the possibility that basic forms of reflective thinking may consist in some generalizations involved in the acquisition of some undermining defeaters. While I plan to articulate the details and significance of this suggestion on another occasion, the point for present purposes is just that there are areas of empirical research that may contribute to understanding the boundary between unreflective and reflective justified belief and knowledge.²⁹

On the other hand, empirical disciplines like developmental psychology and animal cognition may draw on epistemological research on reflective and unreflective rationality to design experiments aimed at testing what subjects may be capable of either form of cognitive achievements. That would help to make progress on understanding the extent of reflective and mere animal knowledge in living subjects. An interaction between philosophical and empirical studies of reflective and unreflective epistemic states may benefit both strands of research.

Acknowledgements This work was supported by a UKRI Future Leaders Fellowship [grant # MR/T042249/1]. Thanks to Tommaso Piazza, Hilary Kornblith, and Crispin Wright for comments on earlier drafts, the participants at a Cogito Work in Progress Seminar at the University of Glasgow for discussion, and four anonymous reviewers for the care with which they engaged with the article.

Authors Contribution For the purpose of open access, the author(s) has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

Availability of Data and Materials Data sharing is not applicable to this article as no new data were created or analysed in this study.

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

²⁸ Cf. Harman's (1986: 7) suggestion that, in a theory of "reasoned change in view", the descriptive and normative dimensions are intertwined.

²⁹ Areas of empirical research that may be relevant include: metacognitive research on the relation between cognitive monitoring and the representation of the epistemic states monitored (e.g. Kornell and Terrace 2007, with critical discussions in Carruthers 2008: § 3.6 and Perner 2012: 106-8), work on inferential reasoning in non-human animals (e.g. Call 2006), research on dogs' sensitivity to the accuracy of informants (e.g. Takaoka et al., 2015), the study of selective trust in children (e.g. Harris et al., 2018), and more.

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

- Beddor, B. (2015). Process reliabilism's troubles with defeat. *The Philosophical Quarterly*, 65(259), 145–159.
- Bermudez, J. L. (2003). *Thinking without words*. Oxford: Oxford University Press.
- Boghossian, P. A. (2014). What is inference? *Philosophical Studies*, 169(1), 1–18.
- Bonjour, L. (2010). Recent work on the internalism—externalism controversy. In J. Dancy, E. Sosa, and M. Steup (Eds.), *A companion to epistemology* (2nd edn, , pp. 33–43). Oxford: Blackwell.
- Boyle, M. (2016). Additive theories of rationality: A critique. *European Journal of Philosophy*, 24(3), 527–555.
- Boyle, M. (2018). A different kind of mind? In K. Andrews and J. Beck (Eds.), *The Routledge handbook of philosophy of animal minds* (pp. 109–118). Routledge: London and New York.
- Brown, J. 2018. *Fallibilism: Evidence and knowledge*. Oxford: Oxford University Press.
- Brandom, R. (2010). Conceptual content and discursive practice. *Grazer Philosophische Studien*, 81(1), 13–35.
- Call, J. (2006). Inferences by exclusion in the great apes: The effect of age and species. *Animal Cognition*, 9, 393–403.
- Carruthers, P. (2008). Meta-cognition in animals: A skeptical look. *Mind and Language*, 58–89.
- Coliva, A. (2014). Moderatism, transmission failures, closure, and humean scepticism. In D. Dodd, and E. Zardini (Eds.), *Scepticism and perceptual justification* (pp. 248–271). Oxford: Oxford University Press.
- Conce, E., & Feldman, R. (2001). Internalism defended. *American Philosophical Quarterly*, 38(1), 1–18.
- Conce, E., & Feldman, R. (1985 (2004)). Evidentialism. *Philosophical Studies* 48/1: 15–34. Reprinted in *Evidentialism: Essays in Epistemology*, Oxford: Clarendon Press, 83–107.
- Constantin, J. (2020). Replacement and reasoning: A reliabilist account of epistemic defeat. *Synthese*, 197(8), 3437–3457.
- Davidson, D. (1982). Rational animals. *Dialectica*, 36, 317–327.
- Dretzke, F. (2006). Minimal rationality. In S. Hurley & M. Nudds (Eds.), *Rational animals?* (pp. 107–116). Oxford: Oxford University Press.
- Field, C. (2020). Anti-exceptionalism about requirements of epistemic rationality. *Acta Analytica*, online first: 1–19.
- Glock, H. J. (2019). Agency, intelligence and reasons in animals. *Philosophy*, 94, 645–671.
- Goldberg, S. (2018). *To the best of our knowledge*. Oxford: Oxford University Press.
- Goldberg, S., & Matheson, J. (2020). The impossibility of Mere Animal Knowledge for reflective subjects. *Erkenntnis*, 85(4), 829–840.
- Goldman, A. (1979). What is justified belief? In G. S. Pappas (Ed.), *Justification and knowledge: New studies in epistemology* (pp. 1–25). Dordrecht: Reidel.
- Graham, P., & Lyons, J. (2021). The structure of defeat: Pollock's evidentialism, Lackey's framework, and prospects for reliabilism. In J. Brown & M. Simion (Eds.), *Reasons, justification and defeat* (pp. 39–68). Oxford: OUP.
- Harmann, G. (1986). *Change in view: Principles of reasoning*. Cambridge, MA: MIT Press.
- Harris, P. L., Koenig, P. L., & Corriveau, M. A. K. H., Jaswal V. K (2018). Cognitive foundations of learning from testimony. *Annual Review of Psychology*, 69, 251–273.
- Hazlett, A. (2012). Higher-order epistemic attitudes and intellectual humility. *Episteme*, 9(3), 205–223.
- Ichikawa, J. J., & Jarvis, B. W. (2013). *The rules of thought*. Oxford: Oxford University Press.
- Kimura, K., & Gopnik, A. (2019). Rational higher-order belief revision in Young Children. *Child Development*, 90, 91–97.
- Kornblith, H. (2014). *A naturalistic epistemology*. Oxford: Oxford University Press.
- Kornblith, H. (2017). Doxastic justification is fundamental. *Philosophical Topics*, 45, 63–80.

- Kornell, N., Son, L. K., & Terrace, H. S. (2007). Transfer of metacognitive skills and hint seeking in monkeys. *Psychological Science*, 18, 64–71.
- Korsgaard, C. (1996). *The sources of normativity*. Cambridge: Cambridge University Press.
- Lackey, J. (1999). Testimonial knowledge and transmission. *The Philosophical Quarterly*, 49(197), 471–490.
- Lord, E. (2018). *The importance of being rational*. Oxford: Oxford University Press.
- McDowell, J. (1994). *Mind and world*. Oxford: Oxford University Press.
- Melis, G. (2014). Understanding undermining defeat. *Philosophical Studies*, 170, 433–442.
- Melis, G. (2018). The intertwining of propositional and doxastic justification. *Australasian Journal of Philosophy*, 96(2), 367–379.
- Pedersen, N. J. L. L. (2017). Pure epistemic pluralism. In A. Coliva & N. J. L. L. Pedersen (Eds.), *Epistemic pluralism* (pp. 47–92). Cham: Palgrave MacMillan.
- Perner, J. (2012). Minimeta: In search of minimal criteria for metacognition. In M. J. Beran, J. Brandl, J. Perner, & J. Proust (Eds.), *Foundations of metacognition* (pp. 94–114). Oxford: Oxford University Press.
- Piazza, T. (2021). Epistemic defeaters. *Routledge encyclopedia of philosophy*. Taylor and Francis. <https://doi.org/10.4324/9780415249126-P080-1>. <https://www.rep.routledge.com/articles/thematic/epistemic-defeaters/v-1>.
- Sliwa, P., & Horowitz, S. (2015). Respecting all the evidence. *Philosophical Studies*, 172, 2835–2858.
- Smithies, D. (2015). Ideal rationality and logical omniscience. *Synthese*, 192, 2769–2793.
- Sosa, E. (2015). *Judgment and agency*. Oxford: Oxford University Press.
- Sturgeon, S. (2014). Pollock on defeasible reasons. *Philosophical Studies*, 169, 105–118.
- Takaoka, A., Maeda, T., Hori, Y., & Fujita, K. (2015). Do dogs follow behavioural cues from an unreliable human? *Animal Cognition*, 18, 475–483.
- Titelbaum, M. (2015). Rationality's fixed point (or: In Defence of Right Reasons). *Oxford Studies in Epistemology*, 5: 253–294.
- Turri, J. (2010). On the relationship between propositional and doxastic justification. *Philosophy and Phenomenological Research*, 80(2), 312–326.
- Wedgwood, R. (2017). *The value of rationality*. Oxford: Oxford University Press.
- Williamson, T. Forthcoming. Justifications, excuses, and sceptical scenarios, J. Dutant, & F. Dorsch (Eds.), *The new evil demon*. Oxford: Oxford University Press.

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

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.