



System Reliabilism and basic beliefs: defeasible, undefeated and likely to be true

Spyridon Orestis Palermos¹

Received: 9 October 2020 / Accepted: 16 February 2021 / Published online: 1 April 2021
© Crown 2021

Abstract

To avoid the problem of regress, externalists have put forward defeaters-based accounts of justification. The paper argues that existing proposals face two serious concerns: (i) They fail to accommodate related counterexamples such as Norman the clairvoyant, and, more worryingly, (ii) they fail to explain *how* one can be *epistemically responsible* in holding basic beliefs—i.e., they fail to explain how basic beliefs can avoid being *arbitrary* from the agent's point of view. To solve both of these problems, a new, externalist, defeaters-based account of justification is offered—viz., *System Reliabilism*. The core message of the view—and the way it deals with both (i) and (ii)—is the claim that the justificatory status of justified basic beliefs originates from being the undefeated outputs of a reliable, *cognitively integrated* system that is *capable* of defeating them. Simply put, to be candidates for being justified, basic beliefs must be epistemically responsible and to be so they must be undefeated *while being defeasible*. The paper also offers a detailed, naturalistic analysis of the notion of cognitive integration. This long-due, mechanistic account of cognitive integration is then used to argue that an additional advantage of System Reliabilism is its unique position to account for the as yet unexplained intuition that *responsible beliefs are also likely to be true*.

Keywords Cognitive integration · Epistemic responsibility · Defeaters · Justification · Agent reliabilism · Proper functionalism

1 Introduction

Many of my beliefs, for example my empirical beliefs, appear to be *basic*. I believe there is a laptop in front of me, that the room in which I am sitting is warm, that noise is coming from the street... *not* because I infer them from other justified

✉ Spyridon Orestis Palermos
palermoss@cardiff.ac.uk

¹ School of English, Communication and Philosophy, Cardiff University, 1.44, John Percival Building, Colum Drive, Cardiff CF10 3EU, UK

beliefs, but because I entertain corresponding perceptual experiences. Such beliefs do not just appear basic to me—they need to be basic. Otherwise, to be justified in believing anything (at all), I would have to be in a position to engage in an infinite regress of providing reasons for my reasons and so on.

But then, two immediate worries follow from this: First, (A) if such beliefs are basic, does my lack of reasons in their support mean that, *from my point of view*, they are unjustified and thus *arbitrary*? The answer to this question is hopefully negative: Even though I have no reasons to offer in their support, I can still be, in some way, justified in holding basic beliefs.

Now, intuitively, being able to establish that we are in a position to avoid arbitrariness with respect to at least some of our basic beliefs would be a good thing, no doubt. But how useful would it be, in practice? That's the second worry. Specifically, (B) are beliefs that are, from my point of view, non-arbitrarily held more likely to be true? If the answer to this question is 'no' or 'I don't know', then showing that some basic beliefs can be held non-arbitrarily would, at least in an important sense, be practically useless. What's the point of non-arbitrarily holding certain basic beliefs, if such beliefs are not likely to accord with reality?

In attempting to answer (A), i.e., explain how basic beliefs can be non-inferentially justified, externalists have offered defeaters-based accounts of justification. Nevertheless, existing attempts face two serious concerns:

- (i) They fail to accommodate relevant counterexamples, such as Norman the clairvoyant (Bonjour, 1980).
- (ii) They fail to explain *how* one who is, on their accounts, justified is also *epistemically responsible* in holding their basic beliefs. As I will argue, this amounts essentially to a failure of explaining how non-inferentially justified basic beliefs can avoid being *arbitrary* from the agent's point of view.

With respect to (B), indeed there is the strong intuition that responsibly (i.e., non-arbitrarily) held beliefs are, normally, likely to be true. Bonjour notes, for example, that “if epistemic justification were not conducive to truth [...], if finding epistemically justified beliefs did not substantially increase the likelihood of finding true ones, then epistemic justification would be irrelevant to our main cognitive goal and of dubious worth.”¹ Yet, so far, there has been no explanation as to what might make this intuition true: *Why* are responsible beliefs likely to be true?² This is a particularly important

¹ Bonjour's definition of justification does not require that a belief be reliably formed (where a reliably formed belief is the product of a process that generates a preponderance of true over false beliefs)—otherwise justified beliefs would likely be true by definition. By 'justified' belief Bonjour means what I here refer to as 'responsible' belief (for more on the relation between responsibly formed and fully justified belief, see Sect. 1). The intuition that responsible beliefs are normally likely to be true is also noted by Cohen (1984) and its truth seems to be at the centre of a recent resurgence of the debate between permissivism and impermissivism—though the authors of this debate use the term 'rationality' instead of 'responsibility' [see, for example, Schoenfeld (2019), Horowitz (2019)].

² I do not count Littlejohn's (2012) argument for the unorthodox claim that only true beliefs can be justified, as this is an entirely different claim. On Littlejohn's account, justified beliefs are guaranteed to be true. The intuition I am interested in allows for justified beliefs to be false.

question, because, as Bonjour further notes, “it is only if we have some reason for thinking that epistemic justification constitutes a path to truth that we as cognitive beings have any motive for preferring epistemically justified beliefs to epistemically unjustified ones (Bonjour 1985, p. 8).”

My aim is to provide a negative answer to (A)—in a way that deals with both (i) and (ii)—and a positive response to (B), by offering, for the first time, a mechanistic explanation as to why responsibly held beliefs are likely to be true. To do so, I offer a new, externalist, defeaters-based account of justification—viz., *System Reliabilism*.

System Reliabilism is built around the, as yet undisclosed, tripartite realisation that: (i) to avoid being arbitrary from the agent’s perspective, the agent’s basic beliefs need to be not just *blameless* but *epistemically responsible*; (ii) mere lack of defeaters can only guarantee epistemic blamelessness; (iii) to also entertain epistemic responsibility, basic beliefs need to lack defeaters *while* the agent is actually capable of having them. In other words, one of the main insights System Reliabilism offers and is specifically designed to accommodate is that basic beliefs are epistemically responsible if and only if they are undefeated *while being defeasible*.

Another important aspect of System Reliabilism is the claim that to be defeasible, basic beliefs must be the outputs of a *cognitively integrated system*. The notion of cognitive integration, however, is notoriously underdeveloped within mainstream epistemology. Drawing on empirical considerations from cognitive science, I develop a detailed *structuralist* understanding of the notion that goes significantly beyond available treatments (Greco, 1999, 2010; Breyer & Greco, 2008). The advantage of building this improved account of cognitive integration is twofold. Not only can it help account for how beliefs produced by integrated systems are responsible by being defeasible but undefeated; it further offers the first mechanistic explanation of why responsible beliefs are likely to be true.

In outline, in Sect. 1, I discuss the way externalists have employed the notion of defeaters to account for non-inferential justification. In Sects. 2 and 3, I argue that, generally, this defeaters-based approach to justification faces the following two (aforementioned) concerns: (i) It fails to accommodate relevant counterexamples such as Norman the clairvoyant and (ii) it fails to explain how lack of defeaters (i.e., the absence of something) may confer epistemic responsibility to the agent’s beliefs (and thereby prevent them from being arbitrary from her point of view). In Sect. 4, I consider two recent and improved, externalist, defeaters-based accounts of justification that attempt, but eventually fail, to tackle (i) and (ii). Thus, in Sect. 5, I introduce System Reliabilism and show how it successfully deals with both (i) and (ii). In Sect. 6, I draw on empirical considerations from cognitive science to offer a sophisticated and naturalistically motivated account of the notion of cognitive integration. Finally, in Sect. 7, I employ this improved understanding of cognitive integration to offer a mechanistic explanation of the important intuition that responsibly formed beliefs are also likely to be true.

2 Infinite regress, arbitrariness and externalism

To avoid the problem of regress, epistemic foundationalism holds that all justification ultimately rests on *basic beliefs* that are *non-inferentially* justified. That is, the justification for basic beliefs does not depend on being derived from other justified beliefs: Basic beliefs are—in the complete absence of any reasons in their support—in *some* way justified.

The problem with this (and one of the main objections against epistemic foundationalism) is that, in the absence of any available supporting reasons, basic beliefs appear to be *arbitrary* from the agent's perspective. As Klein (1999, p. 297) puts it, “foundationalism is unacceptable because it advocates accepting an arbitrary reason at the base, that is, a reason for which there are no further reasons making it even slightly better to accept than any of its contraries.”

This sounds worrying indeed. However, epistemic externalists who deny that justification is always a matter of having reflectively accessible reasons in support of our beliefs are not particularly concerned. They simply insist that, on their view, the absence of available reasons is insufficient for rendering beliefs arbitrary from the agent's point of view (Bergmann, 2004).

This denial of the externalist that the availability of reasons is necessary for justification may sound like asserting that justification has nothing to do with the agent's subjective perspective of the situation. It would certainly be consistent with externalism to insist that the mere fact that a basic belief *b* is the product of a reliable or properly functioning process suffices for *b* to be justified. Crude forms of reliabilism may hold, for example, that a belief is justified just in case it is the product of a process that is *de facto* reliable—independently of whether the agent is in any way aware of this.

Of course, such externalist approaches to the way basic beliefs come to be justified are rather implausible: Obviously, such beliefs would still be *arbitrary from the agent's point of view*. No surprise then that most externalists, including Bergmann, are not satisfied with them. Indeed, while externalists deny that the reflective availability of supportive reasons is necessary for justification, they do not necessarily divorce justification from the agent's perspective entirely. Externalists have come to recognise that justification is also a matter of *epistemic responsibility* in the following sense: Even if the way *S* comes to believe that *p* is in fact reliable, *S* is not justified unless she is also epistemically responsible by being, in *some* way, sensitive to the fact that her belief is properly formed.³

Now, one way to interpret this is to say that the externalist acknowledges the force of the internalist's intuition: The agent's perspective is, in some way, integral to the agent's justificatory status.⁴ Even so, this does not mean that externalists concede to internalism. Claiming that the agent's perspective plays an important role to her doxastic justification does not automatically lead to the claim that justification is a matter of having available reasons in support of one's beliefs. Instead, externalists can

³ See, for example, Greco (1999, 2010), Bergmann (2006), Goldberg (2018).

⁴ See Bergmann (1997) for example.

assign to the agent's point of view *an inhibitory role*: They can claim that epistemic responsibility, on the part of the agent, is a matter of *lacking any negative reasons or evidence against* her beliefs. On this weaker condition on epistemic responsibility, the agent's perspective is involved in a 'stand-by' manner, expected to come to the fore only when something is amiss.

To account for this sense of epistemic responsibility that requires no awareness of supporting reasons, externalists usually invoke the notion of 'mental state defeaters.'⁵ As Bergmann (2006, pp. 175–176) notes, "a no-believed-defeater condition isn't an awareness requirement. A no believed-defeater condition doesn't require awareness of anything. At best, it requires the absence of awareness of something (i.e., defeaters)." Thus, by employing the terminology of defeaters, externalists aspire to account for epistemic responsibility by involving the agent's perspective in a negative fashion:

Externalist Epistemic Responsibility (EER):

S is epistemically responsible in holding p iff S lacks any mental state defeaters against p or the way it was formed.

According to such externalist approaches, then, justification consists in reliability plus epistemic responsibility, where the latter is accounted for in terms of the absence of mental state defeaters. Importantly, on such views, which attempt to

⁵ See, for example, Goldman (1979), Bergmann (2006), and more recently (Goldberg, 2018). *Mental state defeaters* are propositional attitudes or experiences or a combination thereof that the agent possesses, and which prevent the agent's belief from counting as justified. They should be contrasted with *propositional defeaters*, which are true propositions that are not believed by the agent, but the mere existence of which prevents the agent's belief from counting as warranted. Also, *normative defeaters* is another important notion in the literature on defeaters, which Lackey (2008, p. 45) defines in the following way: "A normative defeater is a doubt or belief that S ought to have and that indicates that S's belief that p is either false or unreliably formed or sustained. Defeaters in this sense function by virtue of being doubts or beliefs that S *should have* (whether or not S does have them) given the presence of certain available evidence." Normative defeaters can be true or false as well as possessed or unpossessed, thus, cutting across the mental state/propositional defeaters distinction. When normative defeaters (either true or false) are possessed by the agent, they will obviously be a subset of the agent's mental state defeaters thus undermining her justification for the relevant belief. The interesting question is whether *unpossessed* normative defeaters undermine one's justification/epistemic responsibility, given that these are defeaters that one *ought* to have, given the available evidence. Following Goldberg's (2018) recent diagnosis, I am inclined to respond in the negative. Briefly, according to Goldberg (2018, chs. 1, 2, 5 & 6), the defeaters one ought to have (i.e., normative defeaters) are informed by the 'general expectations' of one's epistemic community. Goldberg then notes that lacking such defeaters does not undermine one's epistemic responsibility (since by lacking them, the agent has to herself no evidence that her belief is false or unreliably formed). Instead, lacking normative defeaters disqualify the agent from being assessed as an epistemically responsible agent altogether—in Goldberg's terminology, lacking normative defeaters does not undermine one's '*prima facie*' justification, but it undermines one's '*ultima facie*' justification. (Among the many examples that Goldberg gives to explain what he has in mind is the parallel of a job candidate who fully satisfies all the advertised essential and desirable criteria of a job description, but fails to comply with general expectations, such as being friendly or bathing regularly (2018, pp. 55–56)). Given Goldberg's analysis, then, unpossessed normative defeaters (either true or false) *disqualify* one from being assessed as epistemically responsible/justified, but they are not part of the core epistemic criteria involved in assessments of justification/epistemic responsibility.

capture the significance of the agent's perspective on her beliefs, awareness of supporting reasons figures nowhere. We can distill this general approach to justification in the form of the following principle⁶:

No Awareness Justification (NAJ):

S is justified in believing p iff (i) p is the product of a reliable and/or properly functioning process and (ii) S has no mental state defeaters against her belief that p (or the way it was formed).

So, to return to the issue concerning the arbitrariness of basic, non-inferential beliefs: Can the way externalists attempt to account for justification—and especially the way they account for the responsibility component—help us explain how certain basic (i.e., non-inferentially derived) beliefs can avoid being arbitrary from the agent's point of view? To answer, we can mirror NAJ to get the following principle of non-inferential justification:

Non-Inferential Justification (NIJ):

A basic belief b is non-inferentially justified for S iff b (i) is the product of a reliable and/or properly functioning process and (ii) is not defeated by any of S's other mental states.

Condition (i) accommodates the externalist's requirement on reliability and/or proper function, and condition (ii) is meant to capture the intuition that justification involves epistemic responsibility. Accordingly, basic beliefs are not arbitrary (from the agent's point of view)⁷ if they are the product of a reliable and/or properly functioning process, while also being undefeated by the agent's other mental states.

Is this sufficient for rendering basic beliefs non-inferentially justified? The answer, I am about to argue, is 'not quite.' While NIJ is a step forward, it does not go far enough. Specifically, the view faces two problems. The first, which we are about to explore by reference to a well-known counterexample, is that satisfaction of NIJ does not guarantee that the agent's belief is epistemically responsible. The second problem is that it fails to explain how lack of defeaters (i.e., the absence of something) may confer epistemic responsibility to the agent's beliefs.

⁶ It is impossible to capture in a single definition the intricacies of all available externalist approaches to justification that invoke the notion of defeaters. This rough definition, however, should be sufficient for the purpose of the present argument.

⁷ An anonymous referee has pointed out that one may conflate the meaning of 'arbitrary' (which implies lack of support, *from the agent's point of view*) with the meaning of 'random' (which implies lack of systematic behaviour, *independently of the agent's point of view*). To prevent this, in most cases, I have so far conjoined 'arbitrary' with the phrase 'from the agent's point of view/perspective'. In the rest of the paper, I continue to do so as much as possible, but when, for stylistic reasons, I don't, this is how I intend the concept to be interpreted.

3 Problem I: Norman the clairvoyant

Consider Norman the clairvoyant, Bonjour's famous thought experiment:

Norman

Norman, under certain conditions that usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact the belief is true and results from his

clairvoyant power, under circumstances in which it is completely reliable. (Bonjour 1980, p. 62)

What should we say about the case? Norman's basic belief is the product of a process that is reliable. Ghijzen (2016) also suggests that if we further assume it has been passed on to him as the product of natural selection, then it is functioning properly.⁸ Therefore, condition (i) of NIJ is satisfied. Norman also lacks mental state defeaters against his belief, or the way it is formed. Thus condition (ii) of NIJ is also satisfied. According to NIJ, then, Norman's belief is non-inferentially justified and, thereby, not arbitrary. Intuitively, however, this is the wrong verdict.

Of course, an easy way to explain away this awkward situation is to follow Moon (2018) who insists that, despite the description of the thought experiment, it is not possible to imagine Norman as lacking mental state defeaters. Therefore, some account along the lines of NAJ should explain why Norman is not justified in his belief.⁹ Essentially, however, this amounts to refusing to engage with the actual thought experiment—a thought experiment that internalists have levelled against externalism for decades and which most externalists have accepted as a fair challenge. My aim here is to accommodate all of the internalists' dialectical demands. So, in what follows, I will keep with the tradition by assuming that Norman does, indeed, lack defeaters for his clairvoyant belief. However hard it is to imagine this, it is worth the try, because taking the experiment at face value can reveal important insights regarding the nature of epistemic responsibility.

⁸ According to Ghijzen (2016), in order for a process to function properly it must reliably produce true beliefs, because of what Graham (2012, 2014) calls a 'consequence etiology': a causal sequence that "explain[s] why something exists or continues to exist in terms of its consequences, because of a feedback mechanism that takes consequences as input and causes or sustains the item as output" (Graham 2014, p. 18). One, though not the only, example of such a consequence etiology is evolution by natural selection. Note that, in his paper, Ghijzen (2016) makes the point that clairvoyance could be the product of natural selection by considering the example of Norbert—another clairvoyant whose story is fashioned out of Norman's story—but there is no reason we cannot claim the same of Norman.

⁹ Moon (2018) endorses the following necessary condition on justification: "S justifiably believes that p only if it's not the case that (in virtue of her background beliefs or knowledge), S should withhold or disbelieve p*"; where p* stands for the proposition 'S's' belief that p was formed reliably.'

Just above, we noted that Norman's belief satisfies both conditions of NIJ and yet it should not count as non-inferentially justified: Even though Norman has no defeaters, his belief is still arbitrary from his point of view. But what makes us think this way? The reason, I submit, is that, given the way the case is set up, not only does Norman have no defeaters against his clairvoyant belief, but—crucially—it appears he is in no position of having any at all. Specifically, his clairvoyance seems to be entirely disconnected from the rest of his mind. There is no obvious way in which Norman's vision, hearing, touch and so on have an effect on his clairvoyance. Contrast this with the interplay between his vision, say, and his hearing. If Norman walks down a dark alley and seems to hear a sound, he might turn around to check whether someone is following him. If he sees no one, then he will assume that there was no sound and that it was just his impression after all. On the contrary, his clairvoyance ability and beliefs are not like this. Instead, they are entirely disjointed from his perceptual faculties. The same goes with regards to the relation between his clairvoyance and his memory. We can tell this much from the description of the case: Norman is supposed to have “no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for or against the thesis that he possesses it.” But the complete absence of reasons for the general possibility of clairvoyance and for the possibility that one possesses it means that one's memory will represent clairvoyance as a strange and unfamiliar process. Therefore, if Norman's mind were anything like ours, this memory-generated representation of unfamiliarity would automatically serve as an undercutting defeater to his clairvoyant belief. The fact that Norman has no such defeater indicates that his clairvoyant power does not interact with his memory either.¹⁰

But why should Norman's inability to possess *any* defeaters have an effect on his justificatory status? Recall that externalists require the absence of mental state defeaters as a way of ensuring that the agent's perspective is involved—at least as a latent inhibitor—in rendering her beliefs epistemically responsible. But when the agent is unable to have any defeaters, her perspective is *blind* (or epistemically inert, so to speak) such that the absence of mental state defeaters cannot confer epistemic responsibility to the agent's beliefs. Thus, while NIJ's condition (ii) on epistemic responsibility is satisfied, Norman's basic belief is *not* epistemically responsible.

To be sure, his belief is *blameless*. Having no defeaters against it, he is not irrational in holding it.¹¹ But a blameless basic belief is not the same as an epistemically responsible belief.¹² Blameless basic beliefs are mere beliefs; from the agent's point of view, they can be neither good nor bad. Put another way, nothing is going for

¹⁰ At least not with his declarative memory. Declarative memory is the main kind of memory we usually refer to when we use the term memory in an unqualified way (standardly, declarative memory is further subdivided into semantic and episodic memory but does not include procedural, or implicit, memory).

¹¹ I am assuming here that one's belief is epistemically irrational and thus epistemically blameworthy iff one is holding it despite having mental-state defeaters against it.

¹² The distinction between epistemically blameless belief and epistemically responsible belief is not entirely new with me. For instance, Pryor (2001, p. 117) mentions this distinction in relation to the New Evil Demon problem. Pryor's example of blameless belief, however, is not about a basic belief, but a belief derived from faulty reasoning, which the agent is not capable of recognising.

them and nothing could possibly go against them either.¹³ Thus, even if a blameless basic belief is also reliably formed, as in Norman's case, it is arbitrary all the same. To avoid arbitrariness, basic beliefs need to be more than just blameless and reliably formed; they must be epistemically responsible.

4 Problem II: Ex nihilo epistemic responsibility

As the above suggests, Norman's case reveals a structural inadequacy built into externalist, defeaters-based approaches to epistemic responsibility and justification. It demonstrates that mere satisfaction of EER won't guarantee that a belief is epistemically responsible. If the agent has no mental states that are inconsistent with her belief that *p*, simply because she is in no position of having any such mental state defeaters, EER is only *trivially* satisfied. In such cases, though blameless, the belief is not epistemically responsible. Epistemically responsible requires, instead, that EER be *substantively* satisfied: The agent needs to lack mental state defeaters against her belief that *p* while *being capable* of having such mental state defeaters.

Thus, mere satisfaction of EER is insufficient for ensuring that the agent is going to be epistemically responsible. But why assume that mere lack of defeaters would be sufficient for ensuring that a belief is epistemically responsible in the first place? Are there any reasons for this assumption? After all, as the above indicates, mere lack of defeaters may prevent a belief from being *irresponsible* (which only entails trivial satisfaction of EER), but it is by no means obvious how or whether it could further raise the belief from the status of epistemic blamelessness to the status of epistemic responsibility (which entails substantive satisfaction of EER). Put simply, the question is: How can the mere *absence* of something like defeaters *add* epistemic responsibility to an agent's belief?

Granted: *prima facie*, it does sound correct to claim that lack of defeaters is somehow involved in being epistemically responsible. But it is not at all obvious how this can be so in a positive rather than negative way. That is to say, it is easy to see how the presence of defeaters can remove epistemic responsibility and justification from a belief; but how can a basic belief (which is non-inferentially derived, and which thereby has, by way of reasons, nothing positive in its support, to begin with) *acquire* epistemic responsibility via the absence of negative reasons (or experiences) against it?

Remarkably (though perhaps unsurprisingly), no explanation is available in the literature. Rather, the only motivation for the view relies, to my knowledge, on intuitions alone. Bergmann (2006, pp. 176–177), for example, notes his agreement with Bonjour (Bonjour and Sosa 2003, p. 32) that the only reason in support of a defeaters-based condition on justification is its "intuitive obviousness": "*it just seems intuitively* (perhaps after considering many examples) that a belief isn't justified if one has either a reason for thinking it false or a reason for doubting the reliability

¹³ Perhaps some metaphysical assumptions, such as believing in the afterlife, or that we are not the victims of a skeptical scenario are also like this.

of its source—and that is so whether or not the belief is in fact reliably formed.”¹⁴ And notice, as Bergmann’s quote suggests, that the intuitive support may only go so far as to motivate the simple idea that defeaters can *remove* justification, while falling short of supporting the fancier idea that their absence somehow *adds* epistemic responsibility. And yet, this is precisely what externalist defeaters-based accounts of justification require—otherwise they have to accept that non-inferentially derived beliefs are either arbitrary (even if blameless), or that their epistemic responsibility arises *ex nihilo*. To put the worry another way, externalists who invoke the notion of defeat are yet to explain *how*—if at all—lack of defeaters *confers* epistemic responsibility to the agent’s beliefs.

A central objective of any externalist defeaters-based account of justification, then, is to explain *how* lack of defeaters can instil epistemic responsibility.¹⁵ Failure to do so could only mean three things: (i) The account is importantly incomplete, because to deny that basic beliefs are arbitrary, it heavily relies on—but fails to explain—the puzzling idea that lack of defeaters ensures that beliefs are not just blameless, but also epistemically responsible; (ii) the account does not rely on the idea that lack of defeaters confers epistemic responsibility, therefore it accepts that basic, non-inferentially derived beliefs are arbitrary; (iii) the account neither relies on the idea that lack of defeaters confers epistemic responsibility, nor accepts that basic, non-inferentially derived beliefs are arbitrary; rather the account assumes that epistemic responsibility can arise *ex nihilo*.

(ii) and (iii) are obviously unacceptable, and while externalists may have so far lived with (i), having a solution to it would constitute a significant advancement to the understanding of our epistemic nature. In Sects. 5, 6 and 7, I introduce a defeaters-based account of justification that explains how and when lack of defeaters confers epistemic responsibility to the agent’s beliefs (i.e., how and when EER is substantively satisfied). Before I do so, however, it will be helpful to consider how a couple of alternative proposals attempt to deal with Norman’s case. If the alternative, externalist defeaters-based accounts of justification that I have in mind can accommodate Norman’s case—by successfully explaining why Norman’s lack of defeaters does not suffice for him to be epistemically responsible—then, perhaps, they can also explain how lack of defeaters can, in other cases and under different conditions, confer epistemic responsibility to our basic

¹⁴ Emphasis in the original.

¹⁵ Regarding this point, an anonymous referee wonders whether proponents of alternative externalist defeaters-based accounts may complain that my understanding of epistemic responsibility is too demanding. Perhaps, that is, they could insist that the notion of epistemic responsibility does not refer to anything positive on the part of the agent but only to the absence of anything negative. In response, recall that the previous and present sections have argued at length that an agent who only lacks defeaters might be *blameless* but not epistemically responsible, which is further required for justification. Additionally, it was further argued, this explains why Norman is not justified, despite being reliable and lacking any defeaters (i.e., Norman is reliable and blameless but not epistemically responsible—therefore he is unjustified). Thus, even if, despite my analysis to the contrary, proponents of externalist defeaters-based accounts of justification insisted that mere blamelessness (i.e., lack of defeaters) and epistemic responsibility are the same, they would be in trouble explaining the intuition that Norman is not justified—an intuition which is widely shared by proponents of the externalist defeaters-based accounts I aim to improve on.

beliefs. Unfortunately, as I will argue, these accounts are not successful on either front, but their shortcomings can be most instructive.

5 Alternatives

5.1 Proper functionalist defeat

Ghijzen (2016) focuses on Norman's case and other variants of clairvoyance to convincingly criticize prominent externalist approaches to justification, including, among others, Process Reliabilism (Goldman, 1979), Inferentialist Reliabilism (Lyons, 2009), Proper Functionalism (Graham, 2012, 2014) and Evidentialist Reliabilism (Comesaña, 2010). Ghijzen then moves on to introduce what he considers to be a necessary clause on justification, which he calls 'Proper Functionalist Defeat':

PFD: S's belief in *p* at *t* is justified only if S does not have a defeater system *D* such that, had *D* been working properly, it would have resulted in S's not believing *p* at *t*.

The notion of a properly functioning defeater system plays a central role in Ghijzen's approach. To explain when a defeater system is working properly, Ghijzen (2016, p. 96) draws on the proper functionalist idea of 'consequence etiology': 'an etiology that "*explain[s]* why something exists or continues to exist in terms of its consequences, because of a feedback mechanism that takes consequences as input and causes or sustains the item as output" (Graham, 2014, p. 18).' Accordingly, in Ghijzen's account, "defeater systems are those systems that have the proper function of reliably preventing the formation or maintenance of false beliefs (without thereby preventing the formation or maintenance of true beliefs)" (*ibid.*, 105). In most cases, such defeater systems were created and continue to exist because of evolution by natural selection—a paradigmatic consequence etiology. Put simply, according to Ghijzen, organisms with defeater systems had fewer false beliefs, which contributed to the proliferation of organisms with such defeater systems. Additionally, Ghijzen notes, the consequence etiologies of defeater systems are inextricably tangled with those of our cognitive faculties:

Given that the outputs of our cognitive faculties are influenced by what the defeater systems have assessed as trustworthy, our cognitive faculties are no longer selected purely on the basis of their own merit, but in combination with the merit of the defeater system. Thus the consequence etiologies of our cognitive faculties by now seem to *include* our defeater systems in an important way. (Ghijzen, 2016, p. 105—emphasis in the original)

Now, with the above in mind, we may ask: How does Ghijzen employ his view to account for clairvoyance? His answer is to invite us to consider what the defeater systems of clairvoyant agents, such as Norman, would have done had they been functioning properly:

It certainly appears plausible that their monitoring mechanisms would have rejected their respective beliefs had they been functioning properly: the information presented by their special senses is not corroborated by any of their other senses, nor does the information stem from a recognizable trustworthy source. This should give their monitoring mechanisms sufficient cause to prevent the information rising to the status of belief (Ghijsen, 2016, pp. 107–108).

According to Ghijsen, then, when Norman has no defeaters, this is because his defeater system is not functioning properly, and had it been functioning properly it would have issued defeaters. Put simply, the way PFD explains why Norman is not justified is the following: Norman's defeater system fails to deliver defeaters, therefore it is not functioning properly. The problem is that not only is this way of accounting for Norman's case ad hoc, it is also inconsistent with Ghijsen's own account.

Recall that in order to explain why simple proper functionalism cannot account for Norman's case, Ghijsen noted that we can easily imagine that his clairvoyance ability is functioning properly: we may simply assume that it is a process, which has been inherited to Norman through evolution by natural selection. Moreover, in a previous quote, we saw Ghijsen noting that the consequence etiologies of our cognitive faculties *include* our defeater systems: If we have a properly functioning cognitive faculty, this is partly because our defeater systems have, over the course of evolution, assessed it as trustworthy. Putting the two together, if we accept that Norman's clairvoyant faculty exists because of evolution by natural selection, then so does a defeater system that has been considering it—again, over the course of evolution by natural selection—trustworthy. But, then, when Norman's defeater system issues no defeaters against his clairvoyant belief about the president's whereabouts, it is, in fact, functioning properly. Ghijsen's account, then, rules that Norman *is* justified—i.e., the exact opposite of what Ghijsen was aiming for.

Now, failing to rule out that Norman is justified is a significant problem for Ghijsen's account, but it is not the only one—especially not in the context of the present discussion. Recall that another important desideratum of externalist defeaters-based accounts of justification is to explain how lack of defeaters may render beliefs not just blameless but also epistemically responsible. This, we noted, is particularly important for explaining how basic beliefs may escape the charge of being arbitrary from the agent's perspective. Unfortunately, PFD does not offer any such explanation. Not only that, but towards the end of his paper, and without any apparent reasons, Ghijsen contends that “monitoring mechanisms are not necessary for perceptual justification” (2016, p. 108). But, then, assuming that PFD is part of a purely externalist account of justification, it would seem that Ghijsen would be content with accepting that basic beliefs, such as perceptual beliefs, are either arbitrary or that their epistemic responsibility arises *ex nihilo*.

5.2 Agent reliabilism

Though the connection may not be immediately apparent, another externalist account of justification that is associated with the notion of defeat is Greco's Agent Reliabilism (Greco, 1999).

To account for epistemic responsibility, Agent Reliabilism proposes to supplement process reliabilism with the *ability intuition* on knowledge.¹⁶ This is the intuition that in order for one's true beliefs to qualify as knowledge, they must be the product of a belief-forming process that counts as a cognitive ability. The motivating idea is that cognitive abilities seem to be the sort of reliable belief-forming processes that one can responsibly rely on for acquiring knowledge, even if one does not have any reasons to offer in their support. For example, no one needs to explain why their visual or auditory experiences are a reliable indication of reality when they come to acquire knowledge on their basis.

Of course, to account for epistemic responsibility in this way, Agent Reliabilism needs to explain when a process counts as a cognitive ability. Accordingly, in a number of places, Greco ((Greco, 1999, 2003, 2010), (Greco and Breyer 2010)) offers the following suggestion: A process counts as a cognitive ability only if it has been *cognitively integrated* into the agent's cognitive system. Thus, the notion of 'cognitive integration' constitutes Agent Reliabilism's cornerstone, and this, in turn, raises the following all-important question: *When* is a process cognitively integrated such that it can count as a cognitive ability that is knowledge-conducive?

While Greco ((Greco, 2003, 2010), (Breyer and Greco 2008)) has suggested several possible answers to this question, the following, general remarks are his preferred and most common approach (Greco, 2003, p. 474; Greco, 2010, p. 152):

One aspect of cognitive integration concerns the range of outputs—if the products of a disposition are few and far between, and if they have little relation to other beliefs in the system, then the disposition is less well integrated on that account. Another aspect of cognitive integration is sensitivity to counter-evidence, or to defeating evidence. If the beliefs in question are insensitive to reasons that count against them, then this too speaks against cognitive integration. In general, it would seem, cognitive integration is a function of cooperation and interaction, or cooperative interaction, with other aspects of the cognitive system.

Greco, however, does not elaborate on how to interpret these points on cognitive integration, and this has unfortunately left his view open to several criticisms (e.g., Bernecker, 2008; Ghijzen, 2016). In fact, Greco himself admits he has “hardly presented a clear and detailed account of cognitive integration” (Greco, 2003).

In the context of the present discussion, the problem is that lack of precision on what cognitive integration consists in leads to the worry that Agent Reliabilism (I) cannot appropriately accommodate Norman's counterexample and (II) cannot explain how and when lack of defeaters is sufficient for rendering a belief

¹⁶ The ability intuition on knowledge has also been proposed as a way to tackle a series of other problems facing process reliabilism, such as the problem of strange and fleeting processes (Greco, 2010) and the problem of the distinctive value of knowledge (Greco, 2002, 2009, 2010). The idea that knowledge must be grounded in cognitive abilities can be traced back to the writings of Sosa (1988, 1993) and Plantinga (1993). For more recent approaches to this intuition, see Greco (1999, 2004, 2007) and Pritchard (2009, 2010, 2012).

epistemically responsible. To see why, consider Breyer and Greco (2008, p. 177) who invite us to imagine two possible scenarios¹⁷:

Case 1: Norman shares our counterevidence against clairvoyance and our dispositions to respect that counterevidence.

Case 2: Norman shares neither our counterevidence against clairvoyance nor our disposition to respect such counterevidence when it arises. Moreover, his dispositions to form beliefs on the basis of clairvoyance, including his dispositions to override counterevidence when such arises, are well integrated with other of Norman's cognitive dispositions.

In the first case, Greco and Breyer note that Norman does not know. In the second case, however, they contend that he *does*.¹⁸

The problem with trying to account for Norman in this way is that it either fails to engage with the internalist challenge, or it speaks past the internalist. If we accept that Norman fails to know, because he has defeaters (i.e., case 1) then, like Moon (2018), we are not engaging with the original thought experiment. And if we assume that Norman neither has nor could he have had any defeaters (i.e., case 2), but we are willing to therefore accept that he can know, then we are ignoring the internalists' widely shared intuition on the case.¹⁹

What we need, instead, is an account of justification that explains why Norman fails to know, despite lacking any defeaters. More precisely, we need the account to explain why, despite lacking defeaters, Norman is not, as the internalist would point out, epistemically responsible in his belief. On the contrary, Greco and Breyer contend that, in the second case, Norman's clairvoyant ability is "a lot like visual perception in us" (2008, p. 178), meaning that Norman's way of forming beliefs is both reliable and epistemically responsible in the same way that our basic perceptual faculties are supposed to be. But, clearly, this claim is wrong. If we were disposed,

¹⁷ See also Greco (2010, p. 154).

¹⁸ Greco and Breyer also consider a third case: "Norman shares neither our counterevidence against clairvoyance nor our disposition to respect such counterevidence when it arises. Moreover, his dispositions to form beliefs on the basis of clairvoyance *are not well integrated* with other of Norman's cognitive dispositions" my emphasis). In this case, which is supposed to be very similar to the second case, Greco and Breyer deny that Norman knows. But it is unclear on what grounds cases 2 and 3 are judged to be epistemically different. After all, the only difference between the two cases is a difference on whether clairvoyance related beliefs and dispositions are "well integrated with other of Norman's dispositions". But by so stipulating whether the process is or is not integrated, Greco and Breyer beg the question, because their reason for offering their series of cases is precisely to test our intuitions on what's required for a process to count as cognitively integrated (and thereby knowledge-conducive).

¹⁹ If we read the second case in light of a related discussion that Greco offers in his (2010, p. 154), it is likely that Greco is trying to make a somewhat different point. According to this alternative interpretation, Greco is asking us to imagine of an agent whose clairvoyance is well integrated in the sense that his clairvoyance beliefs *can* be defeated—so when they are not, they qualify as responsible. Though I could, in principle, agree with this interpretation (depending on the details), it is difficult to accept that, in Greco and Breyer's above second case, Norman is epistemically responsible: On what grounds is his clairvoyance integrated, how does this confer epistemic responsibility to his belief and, most importantly, how can his clairvoyance be integrated all the while Norman is disposed to *disrespect and override* counterevidence to his clairvoyant beliefs?

like Norman in case 2, to ignore or override counterevidence against our perceptual beliefs when such arises—which thankfully we are not—then, even if reliable and true, such perceptual beliefs would fail to be epistemically responsible: Though blameless, such beliefs would be paradigmatically arbitrary, because, were they false, we’d have no way of telling. If agent reliabilism is happy to allow—incorrectly—for such ‘blindly’ undefeated beliefs to qualify as epistemically responsible then, clearly, it also fails the second desideratum on externalist defeaters-based accounts of justification that we laid out above²⁰: Externalists need an account that can explain *how and when* lack of defeaters can confer, not just blamelessness, but also epistemic responsibility to our beliefs, so that basic beliefs, such as our perceptual beliefs, can escape the charge of being arbitrary from our point of view.

Time then to start constructing such an account.

6 Structuralist integration

We noted above that Greco does not provide a detailed account of cognitive integration, and this, in effect, results in agent reliabilism’s failure to (I) accommodate Norman’s case and (II) provide an explanation of the way lack of defeat confers epistemic responsibility to our beliefs. This, however, in no way indicates that the intuition driving Greco’s Agent Reliabilism—i.e., the ability intuition on knowledge—is not on the right track. Indeed, in what follows, my aim is to spell out this intuition further by devising a naturalistically motivated account of cognitive integration that captures both the conditions under and the sense in which (i.e., when and how) lack of defeaters can confer epistemic responsibility to our beliefs. To do so, my starting point will be what Greco sometimes refers to as a ‘structuralist’ understanding of the notion of cognitive integration, according to which “cognitive integration is a function of cooperation and interaction, or cooperative interaction, with other aspects of the cognitive system” (Greco, 2003, p. 474; Greco, 2010, p. 152). While Greco does not develop this idea further, expanding on it can prove particularly helpful in a number of ways.

Let’s start by stipulating, then, in line with Greco’s structuralist remarks, that a reliable belief-forming process is cognitively integrated—and thereby counts as a cognitive ability—if and only if it mutually interacts with other aspects of the cognitive

²⁰ In other words, it seems that Agent Reliabilism does not distinguish between trivial and substantive satisfaction of EER and that it takes EER to be sufficient for epistemic responsibility even when this is trivially satisfied.

system.²¹ How can this approach to cognitive integration help with the main goal of explaining the way lack of defeaters may confer epistemic responsibility to our beliefs?

Since, by the above definition, cognitive abilities mutually interact with other aspects of the agent's cognitive system, their operation and outputs need to be in step with those other cognitive aspects. Thus, when discrepancies occur, the system halts; when no discrepancies occur, the system carries on, uninterrupted, with its automatic operations. In practice, this integrated set up results in the workings and deliverances of cognitive abilities being continuously monitored by other aspects of the cognitive system, in the background. From the agent's point of view, the resulting effect is that *if* there is something wrong with her cognitive abilities or their outputs, *then* she will be able to notice this and respond appropriately. Otherwise, if no alert signals are issued in the form of defeaters, the agent can be epistemically responsible in employing her cognitive abilities and accepting their results *by default*—even if she has absolutely no beliefs to offer as to whether or why her abilities and their results are reliable.

An immediate advantage of this approach to cognitive integration and the ability intuition on knowledge is that it can explain why lacking defeaters for beliefs produced by cognitive ability suffices for them to be, not just blameless, but epistemically responsible too. On this structuralist approach to cognitive integration, cognitive abilities are the reliable integrated components of an overall system, which continuously monitors the workings of its different parts against each other. In result, the agent is epistemically responsible in holding undefeated beliefs that are the product of her cognitive abilities, because these beliefs and the way they are formed are, *actually, being monitored* for being inconsistent with other aspects of the agent's cognitive system *and they successfully survive* the test. To use the terminology of Sect. 3, by thinking of cognitive integration in terms of ongoing mutual interactions with other aspects of the cognitive system, the present approach ensures that when an agent forms her beliefs on the basis of cognitive ability, EER is *substantively*, rather than *trivially*, satisfied.

²¹ From an empirical perspective, this is not an implausible demand. Scientific observations of the human and of many primates' brain indicate that brain areas hardly ever operate in isolation. Instead, during the performance of cognitive tasks, several parts of the brain work in tandem. Even as simple a cognitive task as reaching for and grasping an item sitting on the desk before you requires that neuro-anatomically and functionally distinct subsystems, each primarily responsible for processing either visually or proprioceptively received feedback, operate together [see, for example, Filimon et al. (2009), as well as Turella and Lingnau (2014)]. Similarly, predictive coding theories of the mind make apparent that there is a constant, bidirectional relationship between our basic perceptual processes, memory systems and other inference mechanisms. And cognitive psychologists note that widespread neural interactions between distinct cognitive modules contribute to perceptual awareness: "The integration of information has been considered a hallmark of human consciousness, as it requires information being globally available via widespread neural interactions" (Deroy et al., 2016). See also (Nagel, 2014) as well as Sect. 6, where I offer extensive remarks on how cognitive science lends support to the present approach to cognitive integration.

7 System Reliabilism

According to the above, the present, structuralist approach to cognitive integration offers an explanation of how and when EER is substantively satisfied, such that lack of defeaters can confer, in addition to blamelessness, epistemic responsibility to our beliefs. Let us then incorporate this approach to epistemic responsibility in an externalist defeaters-based account of justification:

System Reliabilism (SR):

S is justified in holding p iff (i) p is the product of a reliable process that is cognitively integrated* and (ii) S has no mental state defeaters against her belief that p (or the way it was formed).

*where a reliable process is cognitively integrated (i.e., counts as a cognitive ability) if and only if it mutually interacts with other aspects of S's cognitive system.

Condition (i) of SR ensures that p is the product of a reliable process that is integrated into the agent's cognitive system. Given (i), condition (ii) ensures that EER is substantively satisfied. Therefore, when SR is satisfied, the agent's beliefs will not be merely reliable and blameless, but also responsibly formed.

Now, we may recall that, due to its problems with the way EER was satisfied, NAJ could not be used to account for non-inferential justification. The problem was that its derivative, NIJ, did not ensure that basic beliefs are epistemically responsible, thus failing to safeguard against the objection that they are arbitrary from the agent's point of view. Can SR provide a better springboard for an account of basic beliefs that escapes this problem?

System Reliabilist Non-Inferential Justification (SR-NIJ):

A basic belief b is non-inferentially justified for S iff b (i) is the product of a reliable process that is cognitively integrated* and (ii) is not defeated by any of S's other mental states.

*where a reliable process is cognitively integrated (i.e., counts as a cognitive ability) if and only if it mutually interacts with other aspects of S's cognitive system.

Condition (i) ensures that b is the product of a reliable process that is integrated into the agent's cognitive system. Given (i), condition (ii) ensures that EER is *substantively* satisfied. So, on this approach, basic beliefs are not arbitrary if they are undefeated by the agent's other mental states and produced by a reliable process that is cognitively integrated. Is this sufficient for preventing basic beliefs from being arbitrary from the agent's perspective?

Let's think about Norman again. The diagnosis was that Norman has no way of telling whether his belief is likely to be false, so, intuitively, he is not, from his point of view, epistemically responsible in holding it. That is, his incapacity to have any mental state defeaters against his belief renders his perspective blind, such that the absence of defeaters is no sign of epistemic responsibility. And yet, if we were

to employ NIJ, we would have to conclude, against intuition, that Norman's basic belief is justified.

SR-NIJ, however, has no problem dealing with the case. On SR-NIJ Norman's lack of mental state defeaters is not enough for rendering his belief-forming process epistemically responsible. The reason is that his clairvoyant power is not integrated to his cognitive system: His special power does not cooperatively interact with any of his sensory modalities. There is no conceivable way, for example, in which Norman's vision, hearing, touch and other perceptual faculties interact with his clairvoyance and neither, it seems, does his memory system, which fails to communicate the strangeness of his clairvoyance power. Compare with the interplay between vision and touch, for example. If I clearly see a cup in front of me, but upon reaching for it, I receive no haptic feedback, I will assume that there is no cup after all and that I must be hallucinating.²² Norman's clairvoyance, by contrast, seems to be entirely disconnected from his other cognitive abilities. Thus, given the absence of interaction with other aspects of his cognitive system, Norman's clairvoyance power is not integrated into it. Therefore, condition (i) of SR-NIJ is not satisfied and, as a consequence, condition (ii) is only *trivially* satisfied. In this way, SR-VIJ rules correctly that Norman's belief fails to be non-inferentially justified.

SR-NIJ therefore accommodates Norman's case in a way that clearly respects the internalist's intuition. More importantly, it does so by offering an explanation of how and when basic (i.e., non-inferentially derived) beliefs can be justified, in (what seems to be, but is not) a mere lack defeaters. According to SR-NIJ, so long as undefeated basic beliefs are the products of reliable integrated belief-forming processes (i.e., cognitive abilities), they are non-inferentially justified in virtue of *being monitored* for being inconsistent with the rest of the agent's cognitive system *and successfully surviving* the test.

The general upshot of SR-NIJ, then, is that mere lack of defeaters won't do to prevent basic beliefs from being arbitrary (even if they are reliably formed); basic beliefs are justified if and only if they are the reliable outputs of a cognitively integrated machine that *is capable of defeating* them, yet they go *undefeated*.

8 Structuralist integration in cognitive science

The present approach capitalises significantly on the notion of cognitive integration, which it holds to result from cooperative interactions between different parts of the cognitive system. Given the theoretical weight assigned to this interactive process, it is important that we gain sufficient understanding of the underlying mechanism.

It should be welcome then that the claims that have been so far advanced with regards to the notion of cognitive integration are in agreement with a growing volume of studies within cognitive science, which can be used to thereby offer naturalist support to the present approach as well as elucidate it in crucial respects. In fact, listening to what cognitive science has to say about the process of cognitive

²² Similar examples can be given about all sorts of interactions between our normal perceptual faculties.

integration is necessary for appreciating the notion's full epistemic significance, and while this is not the right place to detail all of its empirical aspects, focusing on a few considerations can disclose much about its justificatory function. Specifically, it can reveal how cognitive integration, as the dynamic interplay between functionally and neuroanatomically different parts of our cognitive systems, contributes to the realisation of two core, but so far neglected, epistemic properties: The property to *epistemically self-organise* and the property to *epistemically self-regulate*.

8.1 Cognitive integration and epistemic self-organisation

Cognitive self-organisation is a diachronic, developmental process, which starts in early childhood and achieves optimal levels by late adolescence. During this process, different parts of the cognitive system interact until they evolve into a relatively stable configuration. Once the cognitive system has achieved this stable configuration, its component parts have mutually adapted by establishing interconnections that allow them to process information from a number of different modules in a coherent and reliable manner. As Stevens (2009) notes:

Age-related cognitive improvements are the result of how neural networks become increasingly more inter-connected and functionally specialized throughout development. [...] Greater anatomical and functional connectivity presumably permits more efficient communication among brain regions needed for task performance.

This diachronic aspect of integration is at the heart of most fundamental belief-forming capacities. Perceptual awareness, for example, appears inherently multisensory: “The integration of information has been considered a hallmark of human consciousness, as it requires information being globally available via widespread neural interactions” (Deroy et al., 2016). Think how many times we perceive the same property via different sense modalities (the vibrations of a guitar string), and how most times we perceive different properties that are delivered from distinct sense modalities as belonging to the same object (the shape and feel of the guitar on one's hands). Moreover, integration is not important just *across* sense modalities, but also *within* them. With respect to vision, for example, it is commonly held that the brain must bind together properties, such as colour and location, which are processed by different brain areas, into a unified object of perception (see, for example, Roskies, 1999, but also Ghose & Maunsell, 1999). While the exact mechanisms of sensory and multisensory integration remain to be understood, it has become evident that integration involves complex interactions between anatomically and functionally distinct brain areas (see, for example, Lamme & Roelfsema, 2000; Lamme et al., 1998; Hupé et al., 1998; Werner & Noppeney, 2009; Stein et al., 2014).

Additionally, such studies suggest that the interplay between different sense modalities is to a large extent responsible for the reliability of cognitive processing. Integration, Nagel notes, allows cognitive systems such as ours to subconsciously—yet optimally—process information received from different sensory channels:

As the human nervous system draws information from a variety of sensory modalities in making judgments about the environment, it assigns different weights to these channels of information as conditions change (for example, shifting to rely more on touch and hearing as darkness falls and the signal from vision is increasingly blurred). Strikingly, this process of integration is almost perfectly optimal, in the sense that the weight assigned to each sensory modality is an almost perfect reflection of the modality's relative current precision (e.g. Ernst & Banks, 2002). (Nagel 2014, p. 711)

Overall, what research like this indicates is that, diachronically, throughout ontogenetic development, neuroanatomically and functionally distinct cognitive modules interact to self-organise. The result is an integrated structure that is capable of supporting the automatic, yet reliable, production of true beliefs.

8.2 Cognitive integration and epistemic self-regulation

At the same time, research within cognitive psychology suggests that integration (in the form of interactions between different cognitive modules) also contributes to epistemic responsibility, by providing cognitive systems with the ability to self-regulate, synchronically.

An important way in which cognitive interactivity contributes to self-regulation is via noetic feelings. Nagel (2014) notes that as the above subconscious process of multisensory integration shifts the weights ascribed to the different channels of information, there is an accompanying feeling of 'uncertainty' on the part of the agent.²³ This feeling of uncertainty may not in itself have a direct effect on the automatic and subconscious process of multisensory integration. Its conscious, reportable nature, however, is held to have a very important effect on the way agents choose to process information. Within cognitive psychology, this feeling of uncertainty—also known as the feeling of lack of fluency—is associated with dealing

²³ Commenting on the nature of the feeling of uncertainty/lack of fluency, Nagel (2014, p. 13) notes: "In examining the metacognitive function of fluency, it is important to isolate the distinctive value of the metacognitive cue itself, as opposed to the underlying characteristics of the processing giving rise to it [...]. For example, when we experience shifts in our feelings of fluency as perceptual conditions change, while simultaneously undergoing corresponding shifts in our reliance on various sensory modalities, our feelings are epiphenomenal to those shifts in multisensory integration: they have a common cause, but the feelings play no part in the automatic regulation of the balance of power between sensory modalities." Following Nagel, I here assume that the feeling of uncertainty is the by-product of multisensory integration. It should be noted, however, that mathematical models (e.g., Ma et al., 2006) have been proposed that may suggest a somewhat different hypothesis: i.e., on the face of noise, uncertainty is represented in the brain by an activation pattern in the neural population dedicated to the relevant sensory cue that shows increased variance; this increase in the variance of the activation pattern, however, may *precede* shifts in multisensory integration, which might indicate that the feeling of uncertainty is independent of the process of integration. Nevertheless, it is unclear whether it is simple increases in the variance associated with a specific sensory cue over time or increases in its *relative* variance, as *compared* to that of other sensory cues, that leads to the feeling of uncertainty and the shifts in integration [for the role of *relative variance* and uncertainty in multisensory integration, see Heron et al. (2004) as well as Ernst and Banks (2002)]. The latter hypothesis does lend considerable support to the claim that the feeling of uncertainty results from the interactive processes underlying multisensory integration.

with suboptimal cognitive processing. Specifically, Oppenheimer (2008) and Alter et al. (2007) suggest that subjects use the feeling of lack of fluency as a metacognitive cue for choosing between two different processing styles—a fast and a slow one. When subjects experience this feeling, they shift from fast, automatic and effortless processing that makes no requirements on working memory to a slower, analytic and deliberate processing style that places considerable demands on working memory. Moreover, it is widely held that lack of fluency does not only present itself as a metacognitive cue in cases of perception, but it is also found across a spectrum of different kinds of cognitive processing. “When people perceive, process, store, retrieve, and generate information, they experience the ease or difficulty of these cognitive operations” [(Unkelbach & Greifeneder, 2013, p. 11); see also (Alter & Oppenheimer, 2009)]. In epistemological terms, then, this metacognitive cue, which is generated as the result of interactions between different parts of the cognitive system, acts as an undercutting mental state defeater, alerting subjects that there might be something wrong with the way they automatically form their beliefs.

Of course, another important way that integration by means of interactions between different cognitive abilities contributes to self-regulation, beyond experiential metacognitive cues, is on the basis of contrasting beliefs: No psychological experiments are required to observe how the interconnectedness of our memory, reasoning and perceptual systems can deliver undercutting and rebutting mental state defeaters in the form of beliefs. On a cloudless day, the water in the sea looks blue, but the memory of my primary school teacher explaining that the sea reflects the colour of the sky presents me with a rebutting defeater against my perceptual belief. Similarly, if I were to find myself lost in the desert and I visually perceived what seems to be a lake, my knowledge of the existence of ‘oasis mirages’ would present me with an undercutting defeater against my perceptual belief.²⁴

8.3 Structuralist integration and epistemic responsibility

In summary, interactivity between cognitive modules contributes to positive epistemic standing both in a diachronic and a synchronic manner. Developmental processes allow cognitive systems to self-organise into a structure that can automatically, yet reliably, bring about true beliefs. On the basis of this self-organised structure, cognitive modules are then poised to interact in such a way so as to allow the overall system to self-regulate during the time of performance.

Of course, the distinction between diachronic reliability and synchronic responsibility as well as the corresponding distinction between the underlying mechanisms

²⁴ There is the possibility that an agent has previously encountered defeating propositions, which, if sufficiently pressed, they would be able to recall, but, currently, they have no easy access to them. It is unclear whether such propositions could count as either dispositional beliefs or disposition to believe [for the distinction between the two, see Audi (1994) but also Palermos (2018)]. The way the present account would explain this is by pointing out that they do not interact with the rest of the agent’s cognitive system and thus they are not cognitively integrated. As such, they may only count as propositional defeaters, even though the agent may possess a dim representation of them or, if pressed, a (second-order) representation that they have encountered them in the past.

of self-organisation and self-regulation are, to a large extent, theoretical artefacts. Given brains' plasticity, every instance of epistemic self-regulation has the capacity to contribute to the process of epistemic self-organisation and vice versa. The diachronic structure of our cognitive systems is continuously shaped by, and at the same time shapes, its ongoing performance.

For example, Shea et al. (2014) argue that a central function of the reportable nature of metacognitive cues, such as the feeling of lack of fluency, is to allow agents to *socially* construct epistemic strategies for employing such cues. That is, even though metacognitive feelings are the by-product of subconscious cognitive control (which relies on our diachronically shaped neural architecture) and even though, when conscious, such feelings can be employed for synchronic self-regulation, they can also have a further diachronic epistemic effect, which is socially mediated:

Metacognition can also be used diachronically, for example, making it possible for people to discuss how metacognitive representations should be deployed, affecting their own cognitive control [(Job et al., 2010)]. Control strategies based on metacognitive representations, for example, what to do when memory fails, can be the subject of explicit instruction (Shea et al., 2014).

Overall then, cognitive science suggests that the interconnected nature of our cognitive systems allows them to:

- (1) diachronically self-organise into a structure that can support the automatic, yet reliable, production of true beliefs;
- (2) synchronically self-regulate in an epistemically responsible manner by issuing defeaters in the form of metacognitive cues such as the feeling of lack of fluency; and
- (3) do so in a way that (1) and (2) can continuously affect each other to the extent that (1) may be shaped by socially mediated epistemic norms, the need for which is prompted by (2).

Evidently, this picture from cognitive science vindicates and significantly elucidates the claim that cognitive integration in the form of interactions between different cognitive modules is necessary for a process to count as a cognitive ability capable of generating beliefs that are reliable and epistemically responsible.

9 The responsibility-truth connection

The above points offer naturalist support to System Reliabilism while significantly contributing to our understanding of cognitive integration and its epistemic significance. In this final section, I want to demonstrate how they can add to the appreciation of cognitive integration even further. Specifically, on the basis of the preceding remarks, it is possible to understand the central role that cognitive integration plays in making it objectively likely that responsible beliefs are true beliefs.

As noted in Sect. 1, there is the important intuition that responsible beliefs are likely to be true beliefs (alternatively, responsibly produced beliefs will also be reliable).²⁵ To give one more example of an author citing with this intuition, Bergmann (2004, p. 49) has noted that a virtue of his defeaters-based account of justification is the consequence that, according to it, there will be a “high objective probability that a justified belief will be a true belief if the properly functioning faculties that produce it are operating in the environment for which they were ‘designed.’”²⁶

Bergmann’s insight that responsible beliefs are likely to be true when produced in the environment for which one’s faculties were ‘designed’ is a welcome specification. Unfortunately, however, Bergmann offers no explanation as to *why* any of this might be true: How is it that the absence of mental state defeaters makes it likely that the resulting belief will be true, when this is produced by a cognitive process that operates in its ‘normal’ environment? To adequately motivate this claim, we need a mechanistic explanation of the threefold relation between (α) a belief’s lack of mental state defeaters (i.e., its responsibility), (β) its objective high probability to be true (i.e., its reliability) and (γ) the environment within which it is produced.

On the present account, and given the foregoing structuralist understanding of cognitive integration, we can say the following. Our cognitive systems are bundles of modular cognitive abilities. Though modular, on the basis of cognitive integration, cognitive abilities are phylogenetically, ontogenetically and socially co-calibrated to synergistically detect any possible shortcomings—provided they operate in the environments they were ‘designed’ for.²⁷ This is because the part-modular-part-integrated, cross-monitoring structure of our cognitive systems makes it highly unlikely for underperforming abilities or faulty outputs to go undetected—unless, of course, the cognitive system operates under particularly abnormal conditions, where several of its faculties may be compromised at once. In normal environments, if one of the operating faculties, or its resulting beliefs, happens to be problematic, it is highly likely that at least one of the other reliably operating faculties will issue an alert signal in the form of a (rebutting or undercutting) mental state defeater. If no such mental state defeaters are presented, the cognitive agent can be confident that the resulting (responsible) beliefs are highly likely to be true (i.e., reliable).

In other words, given cognitive integration, within normal environments, the mental state defeaters that integrated cognitive systems can deliver are highly likely to track most true propositions that could defeat one’s beliefs. It follows that the

²⁵ A reliable process is standardly defined as a process that gives rise to a preponderance of true over false beliefs. Thus, a reliable belief—i.e., a belief that is the product of a reliable process—is one that is likely to be true.

²⁶ According to Bergmann, a belief is justified if it lacks mental state defeaters. Thus, using the terminology adopted in this paper, Bergmann’s justified beliefs translate into beliefs that are (at best) merely responsible (on the present account, to be justified, undefeated beliefs further need to be reliable and the products of an integrated cognitive system).

²⁷ Phylogeny’s contribution to this calibration is only indirect: Cognitive systems that are not genetically pre-disposed to integrate in the right way, throughout their ontogeny, will be unlikely to pass on their genes to subsequent generations. For example, in the wild, animals whose brains are not hard-wired to integrate audio and visual information are highly unlikely to make it to the age of reproduction.

undefeated (i.e., responsible) beliefs of such integrated systems will also be reliable beliefs (i.e., highly likely to be true). To see how this works, recall that the synchronic self-regulation of integrated cognitive systems continuously shapes their self-organized structure and it is, at the same time, diachronically affected by it. In this way, integrated cognitive systems continuously adapt their ability to self-regulate on the basis of an iterated process of trial and error that is being and has been fine-tuned phylogenetically, ontogenetically and socially in response to the forces that occur in the agent's normal environment.²⁸ The result is a cognitive architecture that is poised to detect nearly all facts that might be relevant to its belief-forming processing, when this is exercised in its normal environment. That is, most times, there will be a near perfect overlap between the set of mental state defeaters that integrated systems can deliver and the set of true propositions that could defeat one's beliefs. Therefore, in normal environments, it will be highly unlikely for beliefs that lack mental state defeaters to be false.²⁹

Thus, with this mechanistic explanation in its disposal, System Reliabilism is in a unique position to account for the intuition that responsible beliefs are also likely to be true beliefs.

10 Conclusion

I have been concerned with (A) demonstrating that basic, non-inferentially derived beliefs can avoid being arbitrary from the agent's point of view and (B) explaining why such non-arbitrary beliefs are likely to be true. My starting point was to argue that, while existing, externalist defeaters-based accounts of justification have been on the right track, they fail (A), because (i) they fail to accommodate related counterexamples such as Norman the clairvoyant and (ii) cannot explain how one can be epistemically responsible in holding basic beliefs—which essentially, amounts to a failure of explaining how basic beliefs can escape the charge of being arbitrary from the agent's perspective. To solve these problems, I introduced System Reliabilism, a new externalist defeaters-based account of justification.

If the foregoing is correct, System Reliabilism compares favourably to alternative, externalist defeaters-based approaches to justification for two main reasons. Firstly, it offers a detailed, mechanistic explanation of how A is possible, i.e., of how basic beliefs can avoid being arbitrary (while, of course, successfully dealing with both (i) and (ii)). Secondly, this mechanistic explanation—which is naturalistically motivated on the basis of empirical considerations from cognitive science—is the only available means for explaining the truth of (B)—i.e., the claim that non-arbitrary basic beliefs are likely to be true.³⁰

²⁸ For discussion of the same point from the point of view of evolutionary epistemology, see Palermos (2015).

²⁹ Obviously, I am referring to beliefs that are the products of cognitive abilities—i.e., cognitively integrated belief-forming processes.

³⁰ Given its affinity to virtue reliabilism, an anonymous referee wonders how System Reliabilism compares with Sosa's prominent virtue reliabilist alternative. I cannot overstate the importance of Sosa's

Briefly, with respect to (A), System Reliabilism holds that basic beliefs are epistemically responsible and thereby non-arbitrary, when they are the undefeated products of cognitively integrated belief-forming processes (i.e., a cognitive abilities). The reason why such undefeated beliefs are responsible is that they are being monitored for being inconsistent with the rest of the agent's cognitive system and they successfully survive the test. Simply put: They are undefeated while being defeasible.³¹ With respect to (B), System Reliabilism holds that cognitively integrated systems have been phylogenetically, ontogenetically and socially calibrated to detect most true proposition that could defeat their outputs. Thus, any resulting undefeated beliefs are, within normal environments, likely to be true.

In this way, System Reliabilism has the means to explain not only how certain basic beliefs can avoid being arbitrary, but why, normally, these beliefs are also likely to be true.

Acknowledgements I am thankful to my colleagues in the Philosophy Unit at Cardiff University for feedback during a work-in-progress seminar as well as to Mona Simion and Brent Madison for offering comments on a previous draft. I am also thankful to two anonymous referees for *Synthese*, whose comments helped me make significant improvements. The paper has also benefited immensely from several discussions I had with John Greco over the years.

Footnote 30 (continued)

work and his insightful account of competence, but it is difficult to interpret his view on knowledge and justification in a way that I can clearly compare it with my own. In *Judgment and Agency*, Sosa (2015) spends a good deal of his time discussing responsibilist virtues, but he does not discuss epistemic responsibility as examined in the present paper. From context, it appears that Sosa would be prepared to ascribe epistemic responsibility in the sense I am interested in only in cases of 'knowing full well.' If this is correct, the problem is that Sosa defines 'knowing full well' as "aptness of *epistemic affirmation* in the first order [that] is attained *through the guidance* of apt second-order awareness that such affirmation would then be apt" (p. 86). So, if this is the only way to be epistemically responsible (and thus avoid arbitrariness)—through apt second-order awareness of the aptness of one's first-order epistemic performance—then the worry is that we are led back to the infinite regress (i.e., is the agent aware that their second-order awareness is apt, and so on?). Perhaps this is why Sosa appears to conclude his discussion on the problem of infinite regress (ch. 9) by considering a pragmatic solution to it—i.e., that even from a purely intellectual perspective we may have better things to do than attempting to infinitely defend our beliefs (p. 212). Yet, in other places, Sosa seems to be, at least in principle, open to a solution similar to the one offered by System Reliabilism—i.e., that epistemic responsibility may be attained by lacking defeaters against a belief that is, as a matter of fact, defeasible. For example, on p. 84, Sosa notes: "But the second order stance required need not *in general* take the form of a conscious judgement. It might just be a pre-supposition, an implicit awareness that all's well enough for first-order judgment. And in certain basic cases this might even be a default stance properly sustained absent defeaters. No special rational basing is required, since it suffices that one be sensitively ready to detect defeaters." Unfortunately, however, Sosa does not pursue this thought any further. I would therefore be very curious (and fortunate should I ever have my curiosity satisfied) to know what Sosa would make of the present account.

³¹ An anonymous referee has expressed the worry that System Reliabilism might be too demanding, because it holds that justification and epistemic responsibility require that the agent be capable of having defeaters. Since knowledge entails justification, this requirement may incorrectly prevent unsophisticated cognisers from counting as having knowledge. As it happens, I agree with the referee's intuition; unsophisticated cognisers do seem capable of having some knowledge. But this should be no problem for System Reliabilism, since the view is in a position to easily accommodate this intuition: Though perhaps not infants, young children and even animals are plausibly capable of having defeaters, at least in the form of noetic feelings.

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

- Alter, A. L., & Oppenheimer, D. M. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychology Review*, 13(3), 219–235.
- Alter, A. L., Oppenheimer, D. M., Epley, N., & Eyre, R. N. (2007). Overcoming intuition: metacognitive difficulty activates analytic reasoning. *Journal of Experimental Psychology: General*, 136(4), 569.
- Bergmann, M. (1997). Internalism, externalism and the no-defeater condition. *Synthese*, 110, 399–417.
- Bergmann, M. (2004). What's NOT wrong with foundationalism. *Philosophy and Phenomenological Research*, 68(1), 161–165.
- Bergmann, M. (2006). *Justification without awareness: A defense of epistemic externalism*. Oxford University Press.
- Bernecker, S. (2008). Agent reliabilism and the problem of Clairvoyance. *Philosophy and Phenomenological Research*, 76(1), 164–172.
- BonJour, L. (1980). Externalist theories of empirical knowledge. *Midwest Studies in Philosophy*, 5, 53–73.
- Bonjour, L. (1985). *The structure of empirical knowledge*. Harvard University Press.
- Bonjour, L., & Sosa, E. (2003). *Epistemic Justification*. Blackwell.
- Breyer, D., & Greco, J. (2008). Cognitive integration and the ownership of belief: Response to Bernecker. *Philosophy and Phenomenological Research*, 76(1), 173–184.
- Cohen, S. (1984). Justification and truth. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 46(3), 279–295.
- Comesaña, J. (2010). Evidentialist reliabilism. *Noûs*, 44(4), 571–600.
- Deroy, O., Faivre, N., Lunghi, C., Spence, C., Aller, M., & Noppeney, U. (2016). The complex interplay between multisensory integration and perceptual awareness. *Multisensory Research*, 29(6–7), 585–606.
- Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870), 429.
- Filimon, F., Nelson, J. D., Huang, R. S., & Sereno, M. I. (2009). Multiple parietal reach regions in humans: cortical representations for visual and proprioceptive feedback during on-line reaching. *Journal of Neuroscience*, 29(9), 2961–2971.
- Ghijssen, H. (2016). Norman and Truetemp revisited reliabilistically: a proper functionalist defeat account of clairvoyance. *Episteme*, 13(1), 89–110.
- Ghose, G. M., & Maunsell, J. (1999). Specialized representations in visual cortex: A role for Binding? *Neuron*, 24(1), 79–85.
- Goldberg, S. C. (2018). *To the best of our knowledge: Social expectations and epistemic normativity*. Oxford University Press.
- Goldman, A. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge* (pp. 1–23). D. Reidel.
- Graham, P. (2012). Epistemic entitlement. *Noûs*, 46(3), 449–482.
- Graham, P. (2014). Functions, warrant, history. In A. Fairweather & O. Flanagan (Eds.), *Naturalizing Epistemic Virtue* (pp. 15–35). Cambridge University Press.
- Greco, J. (1999). Agent reliabilism. In J. Tomberlin (Ed.), *Philosophical perspectives 13: epistemology* (pp. 273–296). Ridgeview Press.
- Greco, J. (2002). Knowledge as credit for true belief. In M. DePaul & L. Zagzebski (Eds.), *Intellectual virtue: Perspectives from ethics and epistemology*. Oxford University Press.
- Greco, J. (2003). Further thoughts on agent reliabilism: Replies to Cohen, Geivett, Kvanvig, and Schmitt and Lahroodi. *Philosophy and Phenomenological Research*, 66(2), 466–480.

- Greco, J. (2004). Knowledge as credit for true belief. In M. DePaul & L. Zagzebski (Eds.), *Intellectual virtue: Perspectives from ethics and epistemology*. Oxford University Press.
- Greco, J. (2007). The nature of ability and the purpose of knowledge. *Philosophical Issues*, 17, 57–69.
- Greco, J. (2009). The value problem. In A. Haddock, A. Millar, & D. H. Pritchard (Eds.), *Epistemic value* (pp. 313–321). Oxford University Press.
- Greco, J. (2010). *Achieving knowledge: a virtue-theoretic account of epistemic normativity*. Cambridge University Press.
- Heron, J., Whitaker, D., & McGraw, P. V. (2004). Sensory uncertainty governs the extent of audio-visual interaction. *Vision Research*, 44(25), 2875–2884.
- Horowitz, S. (2019). The Truth Problem for Permissivism. *The Journal of Philosophy*, 116(5), 237–262.
- Hupé, J., James, A., Payne, B., Lomber, S., Girard, P., & Bullier, J. (1998). Cortical feedback improves discrimination between figure and background by V1, V2 and V3 neurons. *Nature*, 394, 784–787. <https://doi.org/10.1038/29537>
- Job, V. et al. (2010.) Ego depletion – Is it all in your head? Implicit theories about willpower affect self-regulation. *Psychological Science*, 21, 1686–1693.
- Klein, P. D. (1999). Human knowledge and the infinite regress of reasons. *Philosophical Perspectives*, 13, 297–325.
- Lamme, V. A., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, 23(11), 571–579.
- Lamme, V., Supér, H., & Spekreijse, H. (1998). Feedforward, horizontal, and feedback processing in the visual cortex. *Current Opinion in Neurobiology*, 8(4), 529–535.
- Lyons, J. (2009). *Perception and Basic Beliefs: Zombies*. Oxford University Press.
- Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, 9(11), 1432–1438.
- Moon, A. (2018). How to use cognitive faculties you never knew you had. *Pacific Philosophical Quarterly*, 99, 251–275.
- Nagel, J. (2014). The meanings of metacognition. *Philosophy and Phenomenological Research*, 89(3), 710–718.
- Oppenheimer, D. (2008). The secret life of fluency. *Trends in Cognitive Science*, 12(6), 237–241.
- Palermos, S. O. (2015). Could reliability naturally imply safety? *European Journal of Philosophy*, 23(4), 1192–1208.
- Palermos, S. O. (2018). Epistemic presentism. *Philosophical Psychology*, 31(3), 458–478.
- Plantinga, A. (1993). *Warrant and proper function*. Oxford University Press.
- Pritchard, D. (2009). *Knowledge*. Palgrave Macmillan.
- Pritchard, D. (2010). Cognitive ability and the extended cognition thesis. *Synthese*, 175(1), 133–151.
- Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy*, 109, 247–279.
- Pryor, J. (2001). Highlights of Recent Epistemology. *British Journal for the Philosophy of Science*, 52, 95–124.
- Roskies, A. L. (1999). The binding problem. *Neuron*, 24, 7–9.
- Shea, N., Boldt, A., Bang, D., Yeung, N., Heyes, C., & Frith, C. D. (2014). Supra-personal cognitive control and metacognition. *Trends in Cognitive Sciences*, 18(4), 186–193.
- Sosa, E. (1988). Beyond skepticism, to the best of our knowledge. *Mind, New Series*, 97(386), 153–188.
- Sosa, E. (1993). Proper functionalism and virtue epistemology. *Nous*, 27(1), 51–65.
- Sosa, E. (2015). *Judgment and agency*. Oxford University Press.
- Stein, B. E., Stanford, T. R., & Rowland, B. A. (2014). Development of multisensory integration from the perspective of the individual neuron. *Nature Reviews Neuroscience*, 15(8), 520–535.
- Stevens, M. C. (2009). The developmental cognitive neuroscience of functional connectivity. *Brain and cognition*, 70(1), 1–12.
- Turella, L., & Lingnau, A. (2014). Neural correlates of grasping. *Frontiers in Human Neuroscience*, 8, 686.
- Unkelbach, C., & Greifeneder, R. (2013). A general model of fluency effects in judgment and decision making. *The experience of thinking: How the fluency of mental processes influences cognition and behaviour*, 11–32.
- Werner, S., & Noppeney, U. (2009). Superadditive responses in superior temporal sulcus predict audio-visual benefits in object categorization. *Cerebral Cortex*, 20(8), 1829–1842.