THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Extended self-knowledge

Citation for published version:

Pritchard, D & Carter, A 2018, Extended self-knowledge. in P Pedrini & J Kirsch (eds), *Third-Person Self-Knowledge, Self-Interpretation, and Narrative*. Contributions To Phenomenology, vol. 96, Springer, pp. 31-49. https://doi.org/10.1007/978-3-319-98646-3_3

Digital Object Identifier (DOI):

[10.1007/978-3-319-98646-3_3](https://doi.org/10.1007/978-3-319-98646-3_3)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Third-Person Self-Knowledge, Self-Interpretation, and Narrative

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

For *Third-Person Self-Knowledge, Self-Interpretation, and Narrative*,
(eds.) J. Kirsch & P. Pedrini, (Springer, *forthcoming*).

EXTENDED SELF-KNOWLEDGE

J. Adam Carter & Duncan Pritchard

University of Edinburgh & University of California, Irvine/University of Edinburgh

ABSTRACT. We aim to move the externalism and self-knowledge debate forward by exploring two novel sceptical challenges to the prospects of self-knowledge of a paradigmatic sort, both of which result from ways in which our thought content, cognitive processes and cognitive successes depend crucially on our external environments. In particular, it is shown how arguments from *extended cognition* (e.g., Clark & Chalmers 1998; Clark 2008) and *situationism* (e.g., Alfano 2012; 2014; Doris 1998; 2002; Harman 1999; 2000) pose hitherto unexplored challenges to the prospects of self-knowledge as it is traditionally conceived. It is shown, however, that, suitably understood, these apparent challenges in fact only demonstrate two ways in which our cognitive lives can be dependent on our environment. As such, rather than undermining our prospects for attaining self-knowledge, they instead illustrate how self-knowledge can be *extended* and *expanded*.

1. INTRODUCTION

Although much human knowledge is acquired by looking ‘outward’ to the world, one especially important variety of knowledge, self-knowledge, is characteristically attained by looking instead inward at oneself—i.e., toward one’s own mental states (e.g., beliefs, desires), with the objective of representing these states aright. The most famous challenge to the very possibility of self-knowledge adverts to a thesis that is now orthodoxy in the philosophy of mind and language—*viz.*, *content externalism* (e.g., Burge 1986; Putnam 1975). This is, broadly, the view that one’s physical or social environment can play a role in the individuation of mental contents, and hence can play a role in determining what our thoughts are in fact thoughts about.

There is a clear tension between content externalism and self-knowledge, in that if the content of one's thoughts can be in part determined by environmental factors, then how, in the relevant cases at least, can one come to know about one's thoughts by directing one's attention inwards (and hence away from environmental factors)? A great deal of recent philosophical discussions of self-knowledge have focused on this tension, which is in effect a tension between, on the one hand, the privileged 'first-person' perspective that seems to be a hallmark of the kind of self-knowledge that we claim of our own mental lives, and, on the other hand, the 'third-person' perspective which would, given content externalism, seem necessary to ascertain with accuracy what, in the relevant cases, our thoughts are about.

Rather than retrace the contours of this well-worn 'externalist' objection to self-knowledge, we want to move the self-knowledge debate forward by exploring two novel sceptical challenges to the prospects of self-knowledge—challenges which exploit two entirely different ways in which knowledge seems to depend on elements of the world beyond the (biological) agent. One such challenge arises from what has been called in the philosophy of mind the *extended cognition thesis* (e.g., Clark 2008; Clark & Chalmers 1998). This is the thesis that cognitive processes can supervene on parts of the world that are external to the biological agent. The other challenge—known as *situationism* (e.g., Doris 1998; 2002; Harman 1999; 2000)—has its roots in recent work in moral psychology. It insists that our cognitive successes depend to a surprising extent on various extra-agential factors specific to particular situations.

After setting out how these two challenges look initially very troublesome for the prospects of self-knowledge, we argue that in fact the *opposite* is the case. Rather than diminishing our self-knowledge, these two strands of argument in fact reveal thus far unexplored ways that this kind of knowledge can be realised. In particular, these two challenges highlight the importance of what we call *extended knowledge*, which is knowledge that depends in significant ways on an epistemic contribution from factors outwith the (biological) subject, but which is no less compatible, as a result, with self-knowledge.

Here is the plan. In §2, we outline the challenge posed to self-knowledge on the basis of extended (or 'active') cognition and articulate how this challenge differs importantly from the more familiar challenge posed by content (or 'passive') externalism. In §3, we highlight the key contours of the situationist critique and outline how a recent attempt (e.g., Alfano 2012; 2014) to argue from situationism to the falsity of virtue epistemology can be straightforwardly redeployed so as to threaten the scope of self-knowledge. In §4 we show how, on closer inspection, both of these novel

challenges to self-knowledge fail to generate the sceptical conclusions for the prospects of self-knowledge which they appear to threaten. Instead, both challenges ultimately reveal hitherto unforeseen ways in which self-knowledge can be extended and expanded. Finally, in §5, we offer some concluding remarks.

2. THE CHALLENGE FROM EXTENDED COGNITION

The sceptical challenge to self-knowledge posed by content externalism has puzzled epistemologists and philosophers of language, given that the following two independently plausible claims seem *prima facie* incompatible: (i) that the *content* of our thoughts is determined partly by the environment; and (ii) that a subject can know by reflection alone (i.e., from the armchair) the content of her own thoughts.¹

While (i) has been widely endorsed in mainstream philosophy of mind and language since at least the 1980s, widespread endorsement for (ii) goes back much further, at least to Descartes, who famously noted important disanalogies between our knowledge of our own mental states, and our knowledge of other things in the world—*viz.*, the things which mental states concerning the world, such as perceptual beliefs, aim to represent. There is no general consensus regarding which disanalogies are the most important in characterising self-knowledge. However, three important disanalogies which will be of interest to us in what follows involve *directness*, *authority* and *method*.²

First, regarding *directness*: we generally don't regard ourselves as becoming aware of our own mental states by being aware of something else first.³ In this way, it looks as though our mental states are thus distinctively *direct* in a way that knowledge of the world characteristically isn't. Second, regarding *method*: we gain knowledge of our own mental states by *introspection*, by looking inward rather than outward.⁴ Introspection seems idiosyncratic of self-knowledge because it is not the kind of method that we can use to gain knowledge of external things.⁵ Third, regarding *authority*: we generally should treat subjects as authoritative about their own states. If you tell someone you

¹ For a comprehensive overview, see Parent (2013). Some philosophers who have found (i) and (ii) at tension take the incompatibility at face value, though a further popular way to defend the (alleged) incompatibility is via appeal to *slow-switching arguments* (e.g., Boghossian 1989; Burge 1988; Ludlow 1995; 1997).

² There are a range of other claimed disanalogies. For discussion, see Gertler (2015).

³ Gertler (2015, §3.1) notes that such awareness is also taken to be *metaphysically* and not just epistemically direct 'in that no event or process mediates between my awareness and the mental state itself'.

⁴ Though, this is disputed by proponents of transparency accounts of self-knowledge (e.g., Dretske 1994).

⁵ See, however, Carter & Palermos (2014) for critical discussion on this point.

believe something, it's normal practice for that individual to take you at your word, rather than to challenge you.

While these disanalogies have been deployed in the service of demonstrating tensions between content externalism and self-knowledge, they can as well be put to work in the service of demonstrating how a more radical form of externalism poses threats to self-knowledge.

According to *active externalism*, parts of the world can do much more than individuate mental contents. Parts of the world—for example, a notebook, an iPhone, Google Glass, etc., and sometimes even other individuals—can partly *constitute* cognition.⁶ Active externalism is a bold thesis, though it is widely held in contemporary cognitive science, and it is quickly gaining traction in contemporary philosophy of mind and cognitive science.

A commonly discussed variety of active externalism is the hypothesis of *extended cognition*, according to which *cognitive processes* are realizable by brain, body and *world*. To bring this idea into focus, consider the following widely referenced case-pair from Clark and Chalmers (1998):

Inga: Inga has a normally functioning biological memory. When she learns new information, she stores it in her biological memory (as is typical and familiar) and it is her biological memory which she consults when she needs old information.

Otto: Otto suffers from Alzheimer's disease, and like many Alzheimer's patients, he relies on information in the environment to help structure his life. Otto carries a notebook around with him everywhere he goes. When he learns new information, he writes it down. When he needs some old information, he looks it up. For Otto, his notebook plays the role usually played by a biological memory.

The salient similarity between Inga and Otto is *functional*; they are both relying on something (a biological brain and a notebook, respectively) to play the functional role of information storage and retrieval. A salient dissimilarity is of course that what's playing this functional role for Otto is something in the world which he is physically interacting with.

One very natural response to what's going on in this case would be, of course, to grant that Otto is using his notebook in a way that is in fact very similar (perhaps even functionally isomorphic) to the way in which Inga is relying on her biological memory, but then to not draw any further conclusions. On this conservative diagnosis of the case, whatever functional similarities there are between Otto and Inga, they don't motivate any interesting *metaphysical* conclusions—*viz.*, we should simply hold firm that whilst Inga's biological memory is part of the cognitive process she

⁶ The three central varieties of active externalism are the *extended cognition*, *extended mind*, and *distributive cognition* theses. For a recent survey of the differences between these versions of active externalism, and how they interface with more familiar forms of externalism in epistemology and the philosophy of mind, see Carter *et al* (2014).

employs, Otto's (non-biological) notebook can't be part of his *cognitive* process. It is, after all, something in the world external to him, made of paper, not brain matter.

Clark responds that this kind of conservative diagnosis reveals a kind of 'bioprejudice'—i.e., an unprincipled inclination to regard the physical constitution of something as inherently interesting with respect to whether the thing in question is a component part of a cognitive process. As Clark and Chalmers (1998) have famously argued, our judgments about what counts as part of a cognitive process should be guided not by traditional bioprejudice but rather by what they call the *parity principle*:

Parity Principle: If, as we confront some task, a part of the world functions as a process which, were it to go on in the head, we would have no hesitation in accepting as part of the cognitive process, then that part of the world is part of the cognitive process. (Clark & Chalmers 1998, 8)

If Clark and Chalmers are right about this, then notice that, with reference to the parity principle, we should include Otto's notebook as part of his cognitive process insofar as we count Inga's biological brain as part of hers.⁷ Put more forcefully: what follows is that Otto's memory (i.e., his process of memory storage and retrieval) is located partly outside his head. And, furthermore, proponents of extended cognition typically reason that, to the extent that we attribute in the default case certain dispositional beliefs to Inga in virtue of having certain contents *stored* in her biological memory, so we should attribute dispositional beliefs to Otto in virtue having certain contents stored in his notebook.

Extended cognition has, of course, its critics.⁸ But given that the view is becoming increasingly mainstream, the time is ripe to consider whether and to what extent this species of active externalism might itself pose challenges to self-knowledge which mimic or even outstrip the challenges posed by content externalism.⁹

A natural starting point to this end will be to juxtapose the alleged *tension* between content externalism (CE) and self-knowledge (SK) with the pairing of extended cognition (EC) with self-knowledge.

⁷ While 'commonsense functionalism' is the primary way in which extended cognition has been philosophically motivated, the position also has been argued for by appeal to dynamical systems theory. See, in particular, Palermos (2011; 2014).

⁸ See, for instance, Adams & Aizawa (2001; 2008) and Rupert (2004).

⁹ See Carter *et al* (2016) for a collection of recent papers which highlight not only the recent popularity of extended cognition in the philosophy of mind and cognitive science, but also how this view is increasingly influencing positions in contemporary epistemology which have traditionally proceeded against a background of *epistemic individualism*.

Tension (CE + SK): (i) that the *content* of our thoughts is determined partly by the environment; and (ii) that a subject can know by reflection alone (i.e., from the armchair) the content of her own thoughts.

Tension (EC + SK): (iii) that the *processes driving* our thoughts are realized partly by the environment; and (iv) that a subject can know by reflection alone (i.e., from the armchair) the content of her own thoughts.¹⁰

Note that part of the philosophical basis for the alleged tension between CE and SK is that if content externalism is true, then it looks as though we lack the kind of knowledge that we take to be distinguished from third-person knowledge on the basis of the *directness*, *authority* and *method* disanalogies which are often used to characterize self-knowledge. And so to the extent that the tension between EC and SK is genuine, it will be helpful to ask: if extended cognition is true, then what import does this have for the kind of knowledge that we take to be distinguished from third-person knowledge on the basis of the directness, authority and method disanalogies? Consider to this end a ‘self-knowledge’ twist on Clark and Chalmers’ (1998) original case of Otto.

Otto_(SK): Otto* suffers from Alzheimer’s disease, and like many Alzheimer’s patients, he relies on information in the environment to help structure his life. Unlike some Alzheimer’s patients who have (in losing their biological memory) lost in the process significant self-knowledge (i.e., knowledge of their own mental states, beliefs, and feelings), Otto* carries a notebook around with him everywhere he goes, and this notebook plays an important role in Otto*’s preservation of his own mental narrative. For example, when he learns new information about his own mental states (i.e., beliefs, feelings, desires, etc.)—information about his mental states which would be lost in biological storage—he writes it down in the notebook. Likewise, when he needs some old information about his mental life, he looks it up. For Otto*, his notebook plays the role usually played by a biological memory in preserving a mental narrative.

The relevant contrast is now between Otto* and Inga*, where the latter is a parallel self-knowledge version of our original description of Inga.

An initial observation about this case is that Otto*’s beliefs about his own mental states seem at least to be straightforwardly *indirect*, rather than direct, in that Otto* is becoming aware of his own mental states by looking at something else first. This indirectness seems, following Brie

¹⁰ Note that there is a parallel literature which engages with the apparent tension between content externalism and internalist accounts of epistemic justification. Bonjour (1992, 136), for instance, remarks in an oft-cited passage that characterises one line of motivation for incompatibilism:

“The adoption of an externalist account of mental content would seem to support an externalist account of justification in the following way: if part or all of the content of a belief is inaccessible to the believer, then both the justifying status of other beliefs in relation to that content and the status of that content as justifying further beliefs will be similarly inaccessible, thus contradicting the internalist requirement for justification.”

Chase (2001), Vahid (2003), Kallestrup & Pritchard (2004), and Williamson (2007) have followed Bonjour in defending versions of incompatibilism. See, however, Brueckner (2002), Gerken (2008), Madison (2009), and Carter & Palermos (2016) for some arguments for compatibilism.

Gertler's (2015) distinction, to be both *epistemological* and *metaphysical*. Otto* after all appears to be taking the evidence he obtains via consulting his entries in the notebook as his epistemic grounds for his beliefs about his mental states. Furthermore, the act of consulting the notebook comes between his seeking knowledge of his mental states and his awareness of these states.

Second, regarding *authority*: as Crispin Wright (1989, 630) remarks, drawing inspiration from Wittgenstein, it seems very much that mental states are 'subject to groundless, authoritative self-ascription'.¹¹ As Sanford Goldberg (1999, 169) condenses the thrust of this line: 'it is part of our belief-attribution practices, that we treat as legitimate a person's self-attributions of belief even when these self-attributions are formed in an effortless manner and without any evidence'. Consider now Inga* and Otto*. Whereas it is part of our belief-forming practices that we treat Inga*'s self-attributions as authoritative, it's far less clear that we will be inclined to treat Otto*'s self-attributions as equally authoritative, in the sense intimated by Wright. To appreciate why, we can just run a 'high-stakes' case—i.e., where there is a lot of practical significance to whether Inga* and Otto* are correct in their respective self-belief reports. In such a circumstance, we will be *disinclined* to simply take Otto*'s self-attribution as a 'groundless, authoritative self-ascription', even if we take Inga*'s to be, in the very same practical environment. After all, one might be tempted to point out that Otto*'s attribution does not appear to be *groundless* at all, but rather, based on his consultation of the notebook. Rather than treat Otto*'s self-report as authoritative in a high-stakes case—as we might do with regard to Inga*'s self-report—we will thus be inclined to verify Otto*'s report specifically by verifying the report's basis (i.e., by looking in his notebook).

Thirdly, regarding *method*: Otto* hardly appears to be introspecting when ascertaining what he believes. After all, Inga can tell just as easily as Otto* can what Otto* believes by looking in Otto*'s notebook. Yet, as Eric Schwitzgebel (2014) puts it, introspection is generally viewed as a way of learning about your own mind which *no one else can*.¹²

Putting these points about directness, authority and method together, a special kind of threat to self-knowledge, borne from extended cognition, materialises. For lack of a better term, let's call the ensuing argument the argument from *distinctive knowability*. On first blush, call a mental state, M, *distinctively knowable* if M is knowable, for S, in a way that is at least one of the following: (i) direct, (ii)

¹¹ See Goldberg (1999) for a criticism of Wright's neo-Wittgenstenian view. Goldberg's view, in short, is that Wright's project of articulating an account of the epistemology of self-knowledge (i.e., by connecting our epistemic justification for our own mental states with their authoritative nature in socio-linguistic practice) cannot be reconciled with a plausible construal of the *psychology* of self-knowledge.

¹² For a helpful overview on various theories of introspection, see Schwitzgebel (2014).

authoritative; and/or (iii) via an introspective method. We use the term ‘distinctively *knowable*’ for the following reason: self-knowledge is supposed to be interestingly different from garden-variety, third-person knowledge not because each of our mental states is *actually known* in ways that third-person knowledge is not (i.e., directly, authoritatively, or introspectively). Consider after all, Freudian suppression cases. Rather, the idea is that if self-knowledge (as such) is interestingly different from third-person knowledge, it’s because our mental states are distinctively *knowable*—*viz.*, mental states are the sort of things that can be known in ways that are direct, authoritative and via an introspective method. Other kinds of knowledge simply aren’t knowable in these ways.

We can now state the argument from distinctive knowability as follows: for any mental state m , m is *distinctively knowable* only if there is a subject S such that m is knowable for S in a way that is *distinctive*—i.e., direct, authoritative, or via an introspective method. If self-knowledge is interesting as a kind of knowledge, then necessarily, mental states are the sort of thing which are distinctively knowable (regardless of whether they are known). However, proponents of extended cognition will be inclined to regard the case of Otto_{SK} as one where the contents of Otto^* ’s notebook (which include entries about his beliefs, feelings, etc.,) are such that it’s *not* the case that there is a subject for whom these states are distinctively knowable (e.g., especially once the beliefs Otto has about his mental life have eroded entirely from biological storage). Therefore, as the argument goes, if extended cognition is true, then self-knowledge, as such, is not an interesting kind (or, at least, a much less interesting kind) of knowledge.¹³

¹³ A further interesting implication of the argument is that to the extent that (in ordinary biological cases) some individuals have direct, authoritative, introspective access to their mental states, this will be merely an *accidental* feature of their self-knowledge, one that is an artifact of the material constitution of the kind of cognitive process they happen to use to store their beliefs about their mental life.

3. THE CHALLENGE FROM EPISTEMIC SITUATIONISM

The challenge posed by extended cognition to self-knowledge was *undercutting* in the sense that it calls into doubt that self-knowledge is as *distinctive* as we'd been inclined to think. In this section, we want to consider a very different kind of challenge to our self-knowledge, one which does not take issue with self-knowledge as distinctive, but rather which appeals to our dependence on our environments so as to motivate a special sort of sceptical challenge about self-knowledge. This is not a form of radical scepticism about self-knowledge, but rather a moderate sceptical challenge according to which we plausibly have much *less* self-knowledge than we ordinarily suppose.

The challenge here stems from *epistemic situationism*, the view that our judgments are to a surprising extent influenced by factors in our environments, factors that are external to our cognitive agency and mental lives. Epistemic situationism can be organised into three key 'steps'. The first step is descriptive: there is a growing body of empirical literature which demonstrates that 'extra-agential' factors such as ambient light levels and sounds, ambient smells and mood depressors, the presence of bystanders, hunger levels and so on can influence our actions and judgments to a surprising extent.¹⁴ The second step is *analogical*. John Doris (1998; 2002) and Gilbert Harman (2000; 1999; 2003) have reasoned from these kinds of empirical observations to the conclusion that *virtue ethics* is empirically inadequate, insofar as the virtue ethicist's postulation of stable character traits—i.e., moral virtues and vices—in the service of explaining moral (and immoral) behaviour is empirically undermined. This is because of how it is in fact extra-agential factors, including ones of which the subject may be consciously unaware, that are playing much of the relevant explanatory in our ethical conduct, thereby precluding stable character traits from playing this explanatory role. The third step is that if these empirical observations can undermine virtue *ethics* by showing that our ethical virtues do not play the explanatory role in our ethical behaviour that the view proposes, then—as Mark Alfano (2012; 2014) has argued—in principle the very same situationist critique can be applied *mutatis mutandis* to *virtue epistemology* to show that our epistemic virtues do not play the explanatory role in our epistemic behaviour that the view proposes. Virtue epistemology holds that knowledge is acquired via the manifestation of *epistemic virtues*, where this is a class of belief-forming processes which includes cognitive abilities, cognitive faculties, and intellectual virtues. It follows that epistemic situationism can force a dilemma: *either* we deny that knowledge generally results from

¹⁴ For further discussion of these points, see Carter & Pritchard (2015; 2016)

epistemic virtues as virtue epistemology claims, *or* we concede that we have much less knowledge than we hitherto supposed.

Alfano (2014) himself has clarified that for those who embrace the virtue epistemologist's core insight, the epistemic situationist threat is best understood as a threat that targets *inferential* knowledge. This is not surprising, given that the most widely studied cognitive biases—i.e., the *availability heuristic* and the *representativeness heuristic*—are ones that manifest at the level of inference. Interestingly, though, there are a range of biases that look as though they have direct relevance to the scope of our self-knowledge, even on the assumption that self-knowledge in paradigmatic cases is *not* inferential knowledge. For even if self-knowledge is paradigmatically *direct*, in that we don't typically come to know the content of our mental states by reasoning through any premises, we might well still form beliefs about our mental lives in a way that *manifests* cognitive bias. Consider here a non-epistemic analogy to preference. I might state my preference directly, in a way that is not on the basis of anything else, and so state my preference directly, or non-inferentially. Nonetheless, my preference might *manifest* bias (e.g., a knee-jerk preference to hire a person of one ethnicity rather than another).

To appreciate the relevance of this point to self-knowledge, consider now three example biases that have been well-studied—the *hindsight*, *egocentric* and *confirmation* biases:

Hindsight Bias: the filtering memory of past events through one's present knowledge, so that those events are recalled as having been predicted by one; also known as the 'I-knew-it-all-along effect'.¹⁵

Egocentric Bias: recalling the past, including one's own actions, intentions and beliefs, in a self-serving manner.¹⁶

Confirmation Bias: the tendency to search for or interpret information, including about one's own mental life, in a way that confirms one's preconceptions.¹⁷

The hindsight, egocentric and confirmation biases are suggestive of various ways in which our beliefs about the content of our own minds might well be subject a range of biases. Here are three examples. First, the hindsight bias engenders in individuals false beliefs about what they knew. This is an error in self-knowledge; individuals who manifest the bias have an inaccurate picture of the past quality of their beliefs. And, furthermore, the hindsight bias can also engender in individuals

¹⁵ For an influential meta-analysis on studies on the hindsight bias, see Christensen-Szalanski & Willham (1991).

¹⁶ As Taylor & Brown (1988, 193) note, 'considerable research evidence suggests that overly positive self-evaluations, exaggerated perceptions of control or mastery, and unrealistic optimism are characteristic of normal human thought.'

¹⁷ See Nickerson (1998) for a representative overview.

false beliefs about what proposition the individual *did* in fact believe.¹⁸ Therefore, by the hindsight bias, you might thus err in thinking both that certain beliefs of yours were held when they were not, and that of certain beliefs that *were* held but not known, that they were known. Second, in the case of *egocentric bias*, an individual will be inclined to, for example, believe that they've had many praiseworthy beliefs and desires which the individual in fact did not have, thus generating for one an inaccurate (but self-serving) picture of one's doxastic (and cognitive) history. Third, when exhibiting confirmation bias, individuals will reflect on the contents of their minds in ways that are distorted by their preconceptions. If, for example, one has a self-serving preconceived impression that some description of one's beliefs and attitudes is correct, then when reflecting on which beliefs and attitudes one holds one is likely to (mistakenly) accept the preconception as accurate.¹⁹

The hindsight, egocentric and confirmation biases are by no means the only biases which can distort our self-knowledge, but their prevalence forms the basis for two distinct challenges with respect to the received thinking about the nature and scope of our self-knowledge. The first worry concerns scope and can be stated simply: if, as the empirical literature suggests, our judgments about the contents of our own minds often manifest bias (i.e., biases of the sort noted in this section), then it looks as though we have significantly *less* self-knowledge than we might have otherwise supposed. Put another way, the challenge can be understood as highlighting important ways in which even our non-inferential, direct beliefs about the contents of our own minds might well manifest (undetected by us) cognitive biases which undermine the epistemic status of these beliefs about the contents of our mind as knowledge.

A second, and associated, problem is that whatever distinctiveness self-knowledge is regarded as having on the basis of *epistemic privilege* is called into doubt. The directness, authority and introspective character of paradigmatic self-knowledge is generally taken to afford such knowledge with a kind of epistemically security (e.g., Horgan 2012; Gertler 2012). However, the supposition that the epistemic status of paradigmatic self-knowledge can be undermined by the kinds of biases which afflict our 'non-privileged' (third-person) beliefs accordingly challenges the presumption of epistemic security.²⁰

¹⁸ In the former case, one might, for instance, falsely believe that one knew that a previous prediction would come out true. In the latter case, one might, via the hindsight bias, falsely believe that one in fact made a prediction one did not make.

¹⁹ See, for instance, Forer (1949) for a classic study.

²⁰ For a detailed discussion of the challenge that the prevalence of epistemic bias poses to third-person beliefs, see Carter & Pritchard (2016).

4. SELF-KNOWLEDGE, EXTENDED (AND EXPANDED)

The aim of this section will be to show how, to the extent which the challenges to self-knowledge outlined in §§2-3 appear to threaten the nature and scope of self-knowledge, this may well be symptomatic of a tacit commitment to a traditional position in epistemological meta-theory called *epistemic individualism*. After outlining the view, and suggesting some problems for it, we will then suggest how, through the lens of *epistemic anti-individualism*, we have a way of reinterpreting the challenges so as to extend and expand rather than to limit self-knowledge.

Epistemic individualism is the widely embraced view within contemporary epistemology that the cognitive processes grounding knowledge are exclusively internal to the biological subject—indeed, typically, they will be supposed to take place exclusively within the brain (and possibly central nervous system) of the subject.²¹ While the epistemic individualist picture fits snugly with internalist approaches in epistemology, the view has also received explicit and implicit endorsement from paradigmatic epistemic externalists, who regard cognitive processes, the reliability of which matters for knowledge, as themselves resident in the agent's head.²²

If epistemic individualism is true, then the challenges articulated in §§2-3 do indeed appear to have negative import for the prospects of vindicating the presumed extent of self-knowledge, at least as it's construed in paradigmatic cases. For in both cases what these arguments seem to show is that a class of apparent cases of self-knowledge in fact depend in substantive ways on factors which are significantly external to the subject. Granted epistemic individualism, then, it seems to follow that such apparent self-knowledge in such cases is not *bona fide*.

There are, however, good reasons to reject epistemic individualism. In particular, there are independent reasons for supposing that knowledge exhibits a phenomenon known as *epistemic dependence*, such that it can substantively depend on factors external to the subject's cognitive agency. The notion of epistemic dependence in play here arises out of Duncan Pritchard's critique of a form of virtue epistemology—known as *robust virtue epistemology*—that attempts to exclusively understand knowledge in terms of the appropriate manifestation of cognitive ability or virtue.²³ If such a proposal were correct, then epistemic dependence would be ruled out: insofar as one's cognitive

²¹ See Kallestrup & Pritchard (2012) and Pritchard (2016a) for further discussion of epistemic anti-individualism.

²² See, for example, Goldman (1979, 13).

²³ For some of the key defences of robust virtue epistemology, see Sosa (1991; 2007; 2009; 2015), Zagzebski (1996; 1999), and Greco (2010).

success manifests a sufficient level of cognitive agency, then one thereby has knowledge. Conversely, Pritchard argues that the existence of epistemic dependence demonstrates that such a proposal is not correct. More precisely, the phenomenon of epistemic dependence shows that epistemic individualism is untenable, and hence that robust virtue epistemology, insofar as it is wedded to this thesis (as it standardly is), is also untenable.

This rationale for rejecting epistemic individualism can be most straightforwardly appreciated with reference to the ways in which it is due to factors entirely outwith an individual's skin and skull whether her true belief amounts to knowledge—i.e., the ways in which knowledge exhibits epistemic dependence. In particular, epistemic dependence comes in two varieties, positive and negative. *Positive epistemic dependence* features in cases where an agent manifests very little cognitive agency—much less than would normally suffice for knowledge—but where her cognitive success amounts to knowledge nonetheless because of factors in her environment. A good example of this is epistemically friendly testimonial environments, in which an agent can come to have knowledge by for the most part trusting the word of another.²⁴ In such cases the knowledge that is acquired is not primarily due to the agent's manifestation of cognitive agency, but rather substantially depends on the informant's manifestation of cognitive agency. *Negative epistemic dependence*, in contrast, is when an agent manifests a high level of cognitive agency (i.e., of a level that would ordinarily easily suffice for knowledge), but where the cognitive success does not amount to knowledge because of environmental factors. Paradigmatic examples here include barn façade cases.²⁵ In such case, the subject's cognitive success manifests a great deal of cognitive agency, and yet nonetheless fails to amount to knowledge on account of how purely environmental factors ensure that the belief so formed is subject to high levels of epistemic luck/risk.²⁶

Our goal is not to further defend epistemic dependence here, since that thesis has been extensively argued for elsewhere. Rather, we want to show that with this thesis in place, and thus with epistemic individualism rejected, we are in a position to reevaluate the putative challenges posed for self-knowledge by extended cognition and situationism.

²⁴ Note that the clause, 'for the most part', is important here. The claim is not that in the right epistemic conditions mere trust can suffice for testimonial knowledge (gullibility is never a route to knowledge), but rather that in the right epistemic conditions relatively low levels of cognitive agency can suffice for knowledge.

²⁵ See also the *epistemic twin earth* case offered by Kallestrup & Pritchard (2014).

²⁶ For Pritchard's initial critique of robust virtue epistemology, which turns on a distinction between intervening and environmental epistemic luck, see Pritchard (2009a, 2009b; 2012) and Pritchard, Millar & Haddock (2010, chs. 2-4). For a development of this critique in terms of the specific notions of epistemic dependence and epistemic anti-individualism, see Kallestrup & Pritchard (2012; 2013; 2014) and Pritchard (2016a). For further discussion of the notion of epistemic risk and its relevance in this context, see Pritchard (2016b).

4.1. EXTENDED SELF-KNOWLEDGE

Once we embrace epistemic anti-individualism, then the force of the argument from distinctive knowability (outlined in §2) can be significantly mitigated. First, consider the point about *directness*. It was supposed that since Otto* is becoming aware of his own mental states by looking at something else (i.e., the notebook), hence his mental states are not *directly* knowable, but rather mediated by the notebook. With epistemic anti-individualism in the background, however, there is a different way to interpret this case. So long as we take the proponent of extended cognition's functional analogy seriously, then Otto* will be consulting his notebook in a way that is both automatic and entirely uncritical, much as we uncritically embrace the deliverances of intracranial self-scanning. That is, in order for Otto*'s use of the notebook to count as an extended cognitive process, it will need to be cognitively integrated with his other cognitive processes to the same extent as his biological memory, and thus to a similar extent embedded within his cognitive character.²⁷ But if that is right, then there should be no significant difference between Otto* and Inga* in terms of the directness of the belief so formed—it is just that Otto* is making use of external technology while Inga* is relying exclusively on her biological memory. But isn't that simply a form of positive epistemic dependence, such that Otto*'s cognitive processes now in part depends upon external factors, albeit in ways that are favourable to his acquisition of knowledge? Moreover, insofar as Otto* and Inga* are forming their beliefs in analogously direct ways, then it follows that just as we should not regard Inga*'s self-knowledge as mediated by her biological brain, neither should we regard Otto*'s cognitive processes as mediated. Hence there is no reason for supposing, on this front at least, that the knowledge acquired by Otto* and Inga* is in any way different in kind.

Second, consider the point about *authority*. The crux of the matter here is that once we take on board the issue of cognitive integration just noted, then it is hard to see why authority in the relevant sense will not transfer to Otto*'s self-ascriptions which are mediated by the use of the notebook. Put another way, the idea that authority doesn't transfer in this way seems to assume that there is a functional difference between Otto*'s use of the notebook and Inga*'s use of her biological memory, and yet by hypothesis no such difference can in fact be present. In short: if Otto*'s use of

²⁷ For a discussion of the various conditions that should plausibly be satisfied for cognitive integration, see Pritchard (2010; *forthcoming*) and Palermos (2014a).

the notebook in making self-ascriptions is as seamless and integrated within his cognitive processes as Inga*'s use of her biological memory, then on what basis would we treat the former as less authoritative than the latter?

Relatedly, note that once we grant that there is a genuine functional equivalence in play here, then the related point about the groundlessness of self-ascriptions fails to gain any purchase. This point supposes, after all, that Otto* is self-consciously basing his belief on the deliverances of the notebook, while Inga* is simply reporting the content of her mental states directly. But as we noted above, taking the functional equivalence in play here seriously means taking both self-ascriptions to be equally direct, and hence equally unmediated by a conscious basing on rational grounds.

Third, consider the point about *method*. The objection that the entries in Otto*'s notebook fail to be recoverable by *introspection* is an objection with considerably less import on an anti-individualistic reading. The thrust of the original worry was that Otto*'s would-be self-knowledge is not introspectively recoverable because it is *also* recoverable by others. The underlying idea here was that introspection is a way of knowing the contents of your own mind which no one else has. Call this idea *epistemic privacy*.

If *epistemic privacy* is a necessary condition for introspection, then Otto* fails to be introspecting. Interestingly, however, the three most commonly embraced necessary conditions on introspection in the literature, while of course *compatible* with epistemic privacy, do not *entail* epistemic privacy. And, furthermore, Otto*'s situation can plausibly meet these other conditions. As Schwitzgebel (2014, §1.1) notes, introspection 'is aimed at generating, knowledge, judgments, or beliefs about *mental* events, states, or processes' (*mentality condition*), which are 'about *one's own mind only* and no one else's' (*first-person condition*), and which are 'part of one's *currently ongoing* mental life' (*temporal proximity condition*). Note that, on the face of it, neither the mentality, first-person nor the temporal proximity conditions on introspection imply that a process is not introspective if another individual *can* (as Inga can, for instance) have access to one's mind. We thus conclude that there is a perfectly respectable sense in which Otto* counts as satisfying the conditions for introspection, and thereby evades the point about method, even though we grant that he doesn't satisfy the more demanding conditions for epistemic privacy.

The point of the foregoing is that the distinctive knowability argument lacks bite, at least insofar as we are willing to embrace (positive) epistemic dependence, and thus reject epistemic individualism. One upshot of this point is that the scope of self-knowledge is now considerably

widened.²⁸ Indeed, given that we now occupy an age in which we are increasingly embedded within technologically rich environments, the opportunities for substantially extending our self-knowledge on this view are legion. For example, consider the ubiquitous use of smart-phone journal apps, such as the smartphone app ‘Diario’, which allows users to record their ‘experiences, thoughts and ideas’ throughout the day, and in a way that syncs to all of the user’s other devices.²⁹ Diario also tracks time and location, to better facilitate the user’s grasp of the chronology of their thought contents. To the extent that our philosophical theorising treats certain kinds of interactions with technology (such as Diario) on a par with Otto’s use of his notebook as cases of *extended* self-knowledge, it is accordingly clear how the apparent challenge to orthodox thinking about self-knowledge posed by extended cognition opens up the door to more varieties of self-knowledge than can be accounted for within the epistemic individualist’s framework.

The second key benefit regards the *quality* of our self-knowledge, and in particular, its *stability*. The point about stability can be made helpfully with reference to the *Von Restorff effect* and the *reminiscence bump*. The Von Restorff effect is the tendency for individuals, in recovering their own beliefs and attitudes, to more easily recover those that ‘stand out’ than other beliefs, which are more easily forgotten.³⁰ The reminiscence bump is the effect by which individuals are more inclined to recall the beliefs one has which are of a personal nature from adolescence and early adulthood than from other life periods.³¹ In each case, the cost of biological memory storage is *stability*; via effects such as these, some of our self-knowledge is lost because (in each of these ways) some of the content of our minds is unrecoverable. Extended self-knowledge (as is recognised within the anti-individualist paradigm), by contrast, is susceptible to *neither* the von Restorff effect nor the Reminiscence bump, and in this respect, extended self-knowledge has a kind of stability which self-knowledge retrieved from biological storage lacks.³²

4.1. EXPANDED SELF-KNOWLEDGE

²⁸ Relatedly, we think it also follows that the line separating self-knowledge from other kinds of knowledge is now also much more blurred, though we will not be exploring this issue further here.

²⁹ See <https://play.google.com/store/apps/details?id=com.pixelcrater.Diario>.

³⁰ See Karis, Fabiani & Donchin (1984).

³¹ See, for example, Jansari & Parkin (1996).

³² See Carter & Kallestrup (2014) for more detailed discussion on the comparative stability of extended over biological storage.

The situationist challenge canvassed in §3 suggested that we have much *less* self-knowledge than we ordinarily take ourselves to have. This challenge, under closer consideration, might be overstated. After all, the hindsight, egocentric and confirmation biases surveyed were representative of the kind of cognitive biases which can influence our belief formation about our own mental states. If knowledge of our own mental states demands (as, for example, is implied by robust virtue epistemology) that our believing correctly about our mental states should be *primarily* explained by our exercising cognitive abilities or virtues, then the empirical data regarding the prevalence of our susceptibility to these biases would indeed suggest that knowledge is often lacking in these cases.

This line of reasoning goes through only against the background of supposing that if cognitive biases play *some* role in our belief formation, then we thereby fail to know the target belief. But this claim, as we noted at the beginning of §4, is highly suspect. As the arguments from epistemic dependence showed, propositional knowledge does not demand that our cognitive abilities *primarily* explain our cognitive successes. Rather, our cognitive abilities must at least play some weaker but nonetheless significant role in explaining cognitive success in cases of knowledge. But this weaker claim does not preclude that cognitive biases (including the egocentric, hindsight and confirmation biases) and other situational factors can play *some* role in our acquisition of knowledge about the contents of our minds.

Against an epistemic anti-individualistic background—i.e., one which countenances epistemic dependence—there is thus room to vindicate some of our beliefs about our own mental states as *bona fide* knowledge even if they are to some degree afflicted by cognitive bias or other situational factors, so long as these extra-agential factors are not the overarching explanation for the belief's correctness. The crux of the matter is that so long as the explanatory load borne by the cognitive bias is relatively small and doesn't significantly undermine one's cognitive performance, then it can be entirely compatible with knowledge. Furthermore, even when it does play a significant explanatory role in one's cognitive performance, so long as the bias promotes the cognitive goal—not all cognitive biases lead to error, after all—and so long as the subject's cognitive agency is also playing a significant explanatory role, then this can also be compatible with knowledge. In such case, the cognitive bias would rather be a kind of positive epistemic dependence. In these ways, we can vindicate a significant part of our self-knowledge against what would otherwise, on an epistemic individualistic picture, appear to be a more worrying empirically driven sceptical challenge to the

scope of our self-knowledge. The situationist challenge to our self-knowledge may be real, but it is not as extensive as proponents of this challenge suppose.³³

5. CONCLUDING REMARKS

Traditionally, the primary challenge for the prospect of vindicating self-knowledge of a paradigmatic sort in contemporary epistemology and philosophy of mind has been the thesis of content externalism, according to which one's mental contents are individuated by features of her physical or social environment. This paper has considered two other broadly externalist challenges to self-knowledge, both of which advert to ways in which our thought content, cognitive processes and cognitive successes depend crucially on our external environments. In particular, we've shown how extended cognition and epistemic situationism stand, in various ways, in *prima facie* tension with received thinking about the nature and scope of self-knowledge. After outlining these challenges, we've then shown how, to the extent that these purported challenges should be worrying, this is symptomatic of a tacit but we think mistaken background commitment to epistemic individualism. We've concluded by showing how, once we leave epistemic individualism behind, the purported challenges to our extended cognition and epistemic situationism can be attenuated, such that we are left in a better overall epistemic position with regard to the contents of our own minds.³⁴

³³ For more on this point about the limits of the epistemic situationist challenge, see Pritchard (2014). See also Carter & Pritchard (2017) which argues for this point in a specifically educational context.

³⁴ This paper was written as part of the AHRC-funded 'Extended Knowledge' (#AH/J011908/1) project which is hosted by the University of Edinburgh's *Eidyn* Philosophical Research Centre, and we are grateful to the AHRC for their support of this research. Special thanks to Jesper Kallestrup, Julie Kirsch, Patrizia Pedrini, and Orestis Palermos.

REFERENCES

- Adams, F. & Aizawa, K. 2001. The bounds of cognition. *Philosophical Psychology* 14: 43–64.
- Adams, F. & Aizawa, K. (2008). *The bounds of cognition*. Oxford: Blackwell.
- Alfano, M. 2012. Expanding the situationist challenge to responsibilist virtue epistemology. *Philosophical Quarterly* 62: 223–49.
- Alfano, M. 2014. Expanding the situationist challenge to reliabilism about inference. *Virtue Epistemology Naturalized*, (ed.) A. Fairweather, 103–22. Dordrecht, Holland: Springer.
- Boghossian, P. 1989. Content and self-knowledge. *Philosophical Topics* 17: 5–26.
- BonJour, L. 1992. Externalism/internalism. *A Companion to Epistemology*, (eds.) J. Dancy, E. Sosa, & M. Steup, 132–36. Oxford: Blackwell.
- Brueckner, A. 2002. The consistency of content-externalism and justification-internalism. *Australasian Journal of Philosophy* 80: 512–15.
- Burge, T. 1986. Individualism and psychology. *Philosophical Review* 95: 3–45.
- Burge, T. 1988. Individualism and self-knowledge. *Journal of Philosophy* 85: 649–63.
- Carter, J. A., Clark, A., Kallestrup, J., Pritchard, D. & Palermos, S.O. (eds.) 2018. *Extended epistemology*. Oxford: Oxford University Press.
- Carter, J. A. & Kallestrup, J. 2014. Extended cognition and propositional memory. *Philosophy and Phenomenological Research*. doi: 10.1111/phpr.12157.
- Carter, J. A., Kallestrup, J., Pritchard, D. H. & Palermos, S. O. 2014. Varieties of externalism. *Philosophical Issues* 24: 63–109.
- Carter, J. A. & Palermos, S. O. 2016. Epistemic internalism, content externalism and the subjective/objective justification distinction. *American Philosophical Quarterly* 53: 231–44.
- Carter, J. A. & Palermos, S. O. 2015. Active externalism and epistemic internalism. *Erkenntnis* 80: 753–72.
- Carter, J. A. & Pritchard, D. H. 2017. Epistemic situationism, epistemic dependence, and the epistemology of education. *Epistemic Situationism*, (eds.) M. Alfano & A. Fairweather, pp. 168–191 Oxford: Oxford University Press.
- Carter, J. A. & Pritchard, D. H. 2016. Cognitive bias, scepticism and understanding. *Explaining Understanding: New Perspectives from Epistemology and the Philosophy of Science*, (eds.), S. Grimm & C. Baumberger. London: Routledge.
- Chase, J. 2001. Is externalism about content inconsistent with internalism about justification? *Australasian Journal of Philosophy* 79: 227–46.
- Christensen-Szalanski, J. & Willham, C. 1991. The hindsight bias: a meta-analysis. *Organizational Behavior and Human Decision Processes* 48: 147–68.
- Clark, A. 2008. *Supersizing the mind: embodiment, action, and cognitive extension*. Oxford: Oxford University Press.
- Clark, A. & Chalmers, D. 1998. The extended mind. *Analysis* 58: 7–19.
- Doris, J. M. 1998. Persons, situations, and virtue ethics. *Noûs* 32: 504–30.
- Doris, J. M. 2002. *Lack of Character: Personality and Moral Behavior*. Cambridge: Cambridge University Press.
- Dretske, F. 1994. Introspection. *Proceedings of the Aristotelian Society* (new series) 94: 263–78.
- Forer, B. R. 1949. The fallacy of personal validation: a classroom demonstration of gullibility. *Journal of Abnormal and Social Psychology* 44: 118–23.
- Gerken, M. 2008. Is internalism about knowledge consistent with content externalism? *Philosophia* 36: 87–96.
- Gertler, B. 2012. Renewed acquaintance. *Introspection and Consciousness*, (eds.) D. Smithies & D. Stoljar. Oxford: Oxford University Press.

- Gertler, B. 2015. Self-knowledge. *Stanford Encyclopedia of Philosophy*, (ed.) E. N. Zalta, <http://plato.stanford.edu/archives/sum2015/entries/self-knowledge/>.
- Goldberg, S. 1999. The psychology and epistemology of self-knowledge. *Synthese* 118: 165–99.
- Goldman, A. 1979. What is justified belief? *Justification and Knowledge*, (ed.) G. Pappas, 1–23. Dordrecht, Holland: Springer.
- Greco, J. 2010. *Achieving knowledge*. Cambridge: Cambridge University Press.
- Greco, J. 2012. A (different) virtue epistemology. *Philosophy and Phenomenological Research* 85: 1–26.
- Harman, G. 1999. Moral philosophy meets social psychology: virtue ethics and the fundamental attribution error. *Proceedings of the Aristotelian Society* 99: 315–31.
- Harman, G. 2000. The nonexistence of character traits. *Proceedings of the Aristotelian Society* 100: 223–26.
- Harman, G. 2003. No character or personality. *Business Ethics Quarterly* 13: 87–94.
- Horgan, T. 2012. Introspection about phenomenal consciousness: running the gamut from infallibility to impotence. *Introspection and Consciousness*, (eds.) D. Smithies & D. Stoljar, 405–22. Oxford: Oxford University Press.
- Jansari, A. & Parkin, A. 1996. Things that go bump in your life: explaining the reminiscence bump in autobiographical memory. *Psychology and Aging* 11: 85–91.
- Kallestrup, J. & Pritchard, D. H. 2004. An argument for the inconsistency of content externalism and epistemic internalism. *Philosophia* 31: 345–54.
- Kallestrup, J. & Pritchard, D. H. 2012. Robust virtue epistemology and epistemic anti-individualism. *Pacific Philosophical Quarterly* 93: 84–103.
- Kallestrup, J. & Pritchard, D. H. 2013. Robust virtue epistemology and epistemic dependence. *Knowledge, Virtue, and Action: Essays on Putting Epistemic Virtues to Work*, (eds.), T. Henning & D. Schweikard. London: Routledge.
- Kallestrup, J. & Pritchard, D. H. 2014. Virtue epistemology and epistemic twin earth. *European Journal of Philosophy* 22: 335–57.
- Kallestrup, J. & Pritchard, D. H. *Forthcoming*. Epistemic supervenience, anti-individualism and knowledge-first epistemology. *Knowledge First: Approaches in Epistemology and Mind*, (eds.) J. A. Carter, E. C. Gordon & B. Jarvis. Oxford: Oxford University Press.
- Karis, D., Fabiani, M. & Donchin, E. 1984. P300 and memory: individual differences in the von restorff effect. *Cognitive Psychology* 16: 177–216.
- Lackey, J. 2007. Why we don't deserve credit for everything we know. *Synthese* 158: 345–61.
- Lackey, J. 2009. Knowledge and credit. *Philosophical Studies* 142: 27–42.
- Ludlow, P. 1995. Externalism, self-knowledge, and the prevalence of slow switching. *Analysis* 55: 45–49.
- Ludlow, P. 1997. On the relevance of slow switching. *Analysis* 57: 285–86.
- Madison, B. J. C. 2009. On the compatibility of epistemic internalism and content externalism. *Acta Analytica* 24: 173–83.
- Nickerson, R. 1998. Confirmation bias: a ubiquitous phenomenon in many guises. *Review of General Psychology* 2: 175–220.
- Palermos, S. O. 2011. Belief-forming processes, extended. *Review of Philosophy and Psychology* 2: 741–65.
- Palermos, S. O. 2014a. Knowledge and cognitive integration. *Synthese* 191: 1931–51.
- Palermos, S. O. 2014b. Loops, constitution, and cognitive Extension. *Cognitive Systems Research* 27: 25–41.
- Parent, T. 2013. Externalism and self-knowledge. *Stanford Encyclopedia of Philosophy*, (ed.) E. N. Zalta, <http://plato.stanford.edu/archives/sum2013/entries/self-knowledge-externalism/>.
- Pritchard, D. H. 2009a. Apt performance and epistemic value. *Philosophical Studies* 143: 407–16.

- Pritchard, D. H. 2009*b*. Knowledge, understanding and epistemic value. *Epistemology (Royal Institute of Philosophy Lectures)*, (ed.) A. O’Hear, 19-43. Cambridge: Cambridge University Press.
- Pritchard, D. H. 2010. Cognitive ability and the extended cognition thesis. *Synthese* 175: 133–51.
- Pritchard, D. H. 2012. Anti-luck virtue epistemology. *Journal of Philosophy* 109: 247–79.
- Pritchard, D. H. 2014. Re-evaluating the epistemic situationist challenge to virtue epistemology. *Naturalizing Epistemic Virtue*, (eds.) A. Fairweather & O. Flanagan, 143-54. Cambridge: Cambridge University Press.
- Pritchard, D. H. 2016a. Epistemic dependence. *Philosophical Perspectives* 30: 1-20.
- Pritchard, D. H. 2016b. Epistemic risk. *The Journal of Philosophy*, 113(11), 550-571.
- Pritchard, D. H. *Forthcoming*. Extended knowledge. *Extended epistemology*, (eds.) J. A. Carter, A. Clark, J. Kallestrup, S. O. Palermos & D. H. Pritchard. Oxford: Oxford University Press.
- Pritchard, D. H., Millar, A., & Haddock, A. 2010. *The nature and value of knowledge: three investigations*. Oxford, Oxford University Press.
- Putnam, H. 1975. The meaning of “meaning”. *Minnesota Studies in the Philosophy of Science* 7: 131–93.
- Rupert, R. 2004. Challenges to the hypothesis of extended cognition. *Journal of Philosophy* 101: 389–428.
- Schwitzgebel, E. 2014. Introspection. *Stanford Encyclopedia of Philosophy*, (ed.) E. N. Zalta, <http://plato.stanford.edu/archives/sum2014/entries/introspection/>.
- Sosa, E. 1991. *Knowledge in perspective: selected essays in epistemology*. Cambridge: Cambridge University Press.
- Sosa, E. 2007. *A virtue epistemology: apt belief and reflective knowledge*. Oxford: Oxford University Press.
- Sosa, E. 2009. *Reflective knowledge: apt belief and reflective knowledge*. Oxford: Oxford University Press.
- Sosa, E. 2015. *Judgment and agency*. Oxford: Oxford University Press.
- Taylor, S. & Brown, J. D. 1988. Illusion and well-being: a social psychological perspective on mental health. *Psychological Bulletin* 103: 193–210.
- Vahid, H. 2003. Content externalism and the internalism/externalism debate in justification theory. *European Journal of Philosophy* 11: 89–107.
- Williamson, T. 2007. On being justified in one’s head. *Rationality and the good: critical essays on the ethics and epistemology of robert audi*, (eds.), M. Timmons, J. Greco & A. Mele, 106–22. Oxford: Oxford University Press.
- Wright, C. 1989. Wittgenstein’s later philosophy of mind: sensation, privacy, and intention. *Journal of Philosophy* 86: 622–34.
- Zagzebski, L. 1996. *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge: Cambridge University Press.
- Zagzebski, L. 1999. What is knowledge? *The blackwell guide to epistemology*, (eds.) J. Greco & E. Sosa. Oxford: Blackwell.