This is a preprint of a critical notice that is forthcoming in *Analysis Reviews*

Extended Epistemology
Edited by J. ADAM CARTER, ANDY CLARK, JESPER KALLESTRUP, S. ORESTIS PALERMOS AND DUNCAN PRITCHARD
Oxford University Press, 2018. x + 362 pp.

Socially Extended Epistemology
Edited by J. ADAM CARTER, ANDY CLARK, JESPER KALLESTRUP, S. ORESTIS PALERMOS AND DUNCAN PRITCHARD
Oxford University Press, 2018. viii + 318

1. Introduction

Content externalism ala Tyler Burge (1979) and Hilary Putnam (1975) has informed and influenced many epistemological debates. Consider, for instance, the implications of content externalism for the idea that we have privileged knowledge of our own mental states (for an overview see Parent 2017). In contrast, *active* externalism ala Clark and Chalmers (1998) has enjoyed less of a reception within epistemology. Perhaps this is merely due to the fact that content externalism dates back to at least the 1970s, whereas it was Clark and Chalmers' paper that really started the ball rolling on active externalism. Or maybe it is due to a difference in perceived plausibility: content externalism is the majority view, whereas active externalism is perhaps more of a minority concern.[1] Regardless, these two volumes, *Extended Epistemology* and *Socially Extended Epistemology,* aim to (start to) fill this lacuna, and prompt further reflection on the potential epistemological (and other) implications of active forms of externalism.[2]

Both volumes are products of the "Extended Knowledge" project at the Eidyn research centre at the University of Edinburgh, which involved a mix of epistemologists and philosophers of mind. Each volume is split into two sections, "Foundational Issues" and "Applications and New Directions". While this split is crude, it is an effective organising principle. The contributors to the volumes include a mix of researchers from the epistemology and philosophy of mind sides, as well as a few who primarily do research in other fields. Many of these contributions are first-rate. Taken together, they constitute an advert for what the editors call an "extended epistemology": an epistemology that takes the possibility of active externalism seriously, and considers its implications for current epistemological debates, as well as the new research questions that may open up as a result. Topics and issues discussed include (but are not limited to):

- Is active externalism compatible with epistemological internalism (Declan Smithies, J. Adam Carter, Andy Clark and Orestis Palermos)?
- How plausible are Clark and Chalmers' "trust and glue" conditions for extended cognition (Kenneth Aizawa, Isaac Record and Boaz Miller)?
- Are there epistemological objections to extended cognition (Declan Smithies, J. Adam Carter and Jesper Kallestrup, Kenneth Aizawa, Isaac Record and Boaz Miller, Ben Kotzee)?
- Does extended cognition pose a problem for any views in epistemology, like virtue epistemology (K. Brad Wray)?

---

[1] https://philpapers.org/surveys/results.pl. I say "perhaps" because no question about active externalism is included in this survey.
[2] I will refer throughout to *Extended Epistemology* as the "first volume" and *Socially Extended Epistemology* as the "second volume".

- What implications might extended cognition have for theorising about collective intentionality and group mental attitudes (Fred Adams, Jeroen de Ridder, Olle Blomberg, Orestis Palermos and Deborah Tollefson, Joelle Proust)?
- How might taking extended cognition seriously lead us to modify epistemological assumptions (Michael Wheeler, Heather Battaly, Mark Alfano and Joshua August Skorburg, Sabine Roeser, Paul Smart)?
- Are there any links between extended cognition and the Confucian tradition (Chienkuo Mi and Shane Ryan, Eric L. Hutton)?
- New avenues of research, such as collective amnesia (Alessandra Tanesini), emotions (Sabine Roeser), epistemic injustice (Mark Alfano and Joshua August Skorburg), and the potential for interdisciplinary research and collaboration (Ben Kotzee, Holly Arrow and Alexander Garinther, Harry Halpin, Paul Smart).

In this critical notice I draw out some broad themes from each volume. In the process I will, of necessity, focus on some contributions at the expense of others. Taken together, there are 29 contributions I could have discussed. Covering each would have left no time for the sort of general discussion that I presume the reader is looking for. Those wanting an indication of the contents of each contribution can consult the editors' introductions to the volumes, which summarise each contribution in reasonable detail.

Before getting to these broad themes, I want to introduce a few terms I will use throughout. My use of these terms doesn't always match the articles I discuss, but no attempt has been made to impose terminological uniformity between or within the volumes.

2. Terminology and Clarifications

We can start with content and active externalism. Content externalism is the view that the contents of our mental states are partly individuated by our external environments. Thus, for content externalists, Hilary's (on Earth) watery thoughts are about the compound H_2O, whereas twin-Hilary's watery thoughts are about the compound XYZ because the stuff Hilary interacts with (drinks from taps, swims in) is H_2O whereas the stuff twin-Hilary interacts with is XYZ. Active externalism is the different—and, many think, more radical—view that an agent's cognitive processing can extend beyond the body and brain and into their external environment. Thus, for active externalists, when Otto (who has Alzheimer's) consults his notebook to look up some information he needs, his notebook is part of a process that is functionally analogous to storing and retrieving information in memory. Because of this analogy, we should be as willing to regard this process as cognitive as we are to regard memory as cognitive, and the notebook as much part of it as our biological memory is part of the analogous process. (This is the "parity principle" from Clark and Chalmers 1998, and it is discussed in many contributions).

It is important to make two distinctions here. The first is between three different versions of active externalism. The first, which I will call the *hypothesis of extended cognition* (HEC), is essentially just the view stated above: cognitive processes can extend into the external environment beyond individual organisms. The second, which I shall call the *hypothesis of distributed cognition* (HDC), is that cognitive processes can be distributed between several individuals. While HDC entails HEC, they are not equivalent: one may hold that cognitive processes can involve external artefacts without holding that they can involve several individuals (or vice versa). The third, which I shall call the *hypothesis of extended mind* (HEM), is that mental states (beliefs, desires, intentions) can (partly) supervene on elements of the world beyond individual organisms. HEM is different from HEC because one can hold that, for instance, when Otto consults his notebook the notebook is part of a cognitive process without holding that the beliefs Otto forms from looking at the notebook supervene on the notebook or what is written in it. In what follows I will

sometimes talk about "extended cognition" or "the hypotheses of extended cognition" without specifying which of these specific hypotheses I have in mind. But the reader should bear in mind that the differences between them do matter for certain purposes.

The second distinction is between all of these hypotheses and what (following several contributors) I will call *embedded cognition*. Defenders of the view that cognition is embedded think that cognition often draws on the individual organism's external environment (perhaps including other people, in which case we can call cognition "socially embedded"). Otto's reliance on an external resource in figuring out how to get around the world isn't particularly unusual; we all, to varying degrees, rely on external resources in our thinking. But, where defenders of HEC (and other forms of active externalism) hold that these external resources can be a *constitutive* part of our cognitive processes, defenders of embedded cognition hold that they merely play a *causal* role. For defenders of HEC, Otto's notebook is *part* of his cognitive process, in the same way that biological memory is part of memorial processes. For defenders of embedded cognition, Otto's notebook is merely a *causal trigger* for a cognitive process. It provides him with input that he can use, in much the same way that seeing a tiger with our eyes provides us with input we can use (run!). It is hard to see how anybody could dispute the claim that cognition is embedded in this way, so the real question is whether any of HEC, HEM and HDC are true.

With all this in hand, I now turn to some of the broad themes of the volumes. While these themes are generally more prominent in one volume than the other, in some cases there is a degree of overlap. In each case I indicate where the interested reader should look to find relevant discussion.

3. Epistemological Arguments against Extended Cognition

The first theme I want to highlight is most prominent in the first volume. This theme is the possibility of there being epistemological grounds for rejecting the hypotheses of extended cognition. This line is most prominent in the articles by Declan Smithies ("Access Internalism and the Extended Mind"), J. Adam Carter and Jesper Kallestrup ("Extended Circularity") and Isaac Record and Boaz Miller ("Taking iPhone Seriously"). I will summarise each article, and then draw some broad conclusions.

Smithies argues that HEM is incompatible with (access) internalism, and rejects HEM on the basis that the case for access internalism is more compelling than the case for HEM. The basic idea is that, while Otto's notebook may play much the same causal role as an ordinary belief, it plays a different normative role. When Otto's "notebook system" goes wrong, this doesn't constitute a breakdown in rationality. Otto's use of his notebook may become erratic over time, but that doesn't make him epistemically irrational. In contrast, failures in our ordinary memorial systems can constitute breakdowns in rationality.

Carter and Kallestrup introduce the problem of "cognitive bloat" (this objection is discussed in several contributions). We need criteria for distinguishing between external resources that are (or can be) part of cognitive architecture and external resources that aren't. They consider (following Clark 2010)) the following three conditions (p. 45):

1. That the resource be reliably available and typically invoked.
2. That any information thus retrieved be automatically endorsed.
3. That information contained in the resource be easily accessible as and when required.

These criteria provide an explanation why Otto's notebook is part of his cognitive architecture: they are all met. The worry is that these criteria are too easy to meet. Take, for instance, my

reliance on the phone book circa 2000. It was reliably available and typically invoked (I regularly used it to find my friends' numbers), I never questioned it, and it was presented in a format that was easy to use. What stopped the phone book being part of my cognitive architecture? Many find this result unwelcome.

Carter and Kallestrup go on to argue that a plausible fourth criterion is problematic on epistemological grounds. This is that the reliability of the resource be endorsed.[3] But this commits the defender of HEC to an objection form of epistemic circularity. Take Otto. If he has to endorse his notebook as reliable, then he is going to make use of the notebook in the process of investigating its reliability ("was it reliable yesterday? Well, let's look at what the notebook says…"). This puts the defender of HEC who uses the fourth criterion in the same position as someone who holds that we need to endorse the reliability of our basic faculties (perception, memory) in order for them to produce knowledge. There is thus a potential epistemic circularity problem for proponents of HEC.

Record and Miller address the issue of new and emerging technology and extended cognition via an illustrative example: our use of the GPS feature on smartphones. They argue (if I understand correctly) that Clark's "trust and glue" conditions (summarised above) mischaracterise the epistemic responsibilities of smartphone users—and, by extension, of users of other instruments and resources. On this basis they conclude that the hypotheses of extended cognition should be rejected. The basic point is that the epistemically responsible GPS user does *not* automatically endorse what their GPS tells them. Indeed, doing so is often a recipe for disaster. The information one gets from GPS is useful, provided one makes use of it in a discerning way. The more general point is that Clark's trust and glue conditions for cognitive extension require the agent to implicitly trust the relevant external resource. But, argue Record and Miller, this stance is often an irresponsible one to take with regard to many technologies. This leaves the defender of extended cognition with a dilemma: either severely limit the scope of cognitive extension, or accept that you often need to be epistemically irresponsible for your cognitive processes to extend.

I would submit that, of these arguments, Record and Miller's should be most troubling to defenders of extended cognition. Carter and Kallestrup arguably succeed only in demonstrating the inadequacy of their suggested fourth condition for cognitive extension. What about conditions that require a form of constant two-way interaction between agent and external resource (discussed in several other contributions)? Smithies may well succeed in establishing his incompatibility claim, but many epistemologists reject access internalism (though I should note that Smithies defence of it is about the most capable I am aware of). Further, his argument targets HEM rather than HEC. What Record and Miller show, I think, is that there may be a crucial disanalogy between ordinary internal cognitive processes (like retrieving information from memory) and any candidate external cognitive process. When Inga consults her biological memory for the location of MoMA the status of this process as cognitive does not depend on whether she is epistemically responsible or anything of the sort. So, if we are to preserve parity, the status of the process involving Otto and his notebook (or a frustrated driver and their GPS) as cognitive surely should not depend on whether Otto makes epistemically responsible use of his notebook. If—as Record and Miller think—it intuitively does so depend, then there is a crucial disanalogy between internal cognitive processes and supposed external cognitive

---

[3] As Aizawa notes in his contribution ("Extended Cognition, Trust and Glue, and Knowledge"), there is a slight oddity in the dialectic here, as this condition doesn't seem to help with the phone book case (I consciously endorsed the reliability o the phone book). It rather (arguably) helps with a variant on the familiar (to epistemologists) TrueTemp case, in the process demonstrating that this variant on the TrueTemp case raises a different problem to the phone book case.

processes. This may cast doubt on the "parity principle" that is crucial in Clark and Chalmers' argument for extended cognition.

This leads on to my second broad theme, which concerns a distinction between the metaphysical question of when cognition extends and the epistemological question of when a cognitive process that involves an external artefact (or other person) produces knowledge.

4. Cognitive Integration and Extended Knowledge

What is extended knowledge? You might think that you get a sort of extended knowledge when knowledge is produced via extended cognitive processes (for another sort see below). But a recurring theme in the first volume (and in parts of the second) is that things are perhaps not as simple as this. In his contribution ("Extended Knowledge") Duncan Pritchard discusses what he calls the problem of "cognitive integration". This is, put roughly, the problem of giving an account of when a process involving an external artefact (or other person) that produces beliefs should be regarded as sufficiently integral to the agent's character for the beliefs it produces to be creditable to the agent. The need for such an account is clear against the backdrop of a virtue epistemology, on which, also roughly, an agent has knowledge when they have a true belief that is sufficiently creditable to their cognitive agency. If it can be shown that a process is part of the agent's cognitive character, then the beliefs it produces can be regarded as creditable to the agent; if it can't, then the beliefs cannot be so regarded.

It is an interesting question how the problem of cognitive integration relates to the problem of cognitive bloat, which I discussed in the previous section. This was the problem of specifying criteria to distinguish between genuine cognitive extension (e.g. Otto's notebook) and spurious cases (e.g. my use of the phone book). I think it is fair to say that, in some (though not by no means all) contributions, these problems are often run together, if not treated as equivalent (e.g. Pritchard's "Extended Knowledge", Carter, Clark and Palermos' "New Humans? Ethics, Trust, and the Extended Mind"). Thus, for some contributors, a process involving an external artefact seems to count as an extended cognitive process to the extent that the agent (or, in some versions, the agent-plus-external-resource) deserves credit for the beliefs produced by the process.

In their contributions Kenneth Aizawa ("Extended Cognition, Trust and Glue, and Knowledge") and Ben Kotzee ("Cyborgs, Knowledge, and Credit for Learning") argue that this is a mistake. Put briefly, their point is that the problem of cognitive bloat is a metaphysical problem whereas the problem of cognitive integration is an epistemological problem. As Kotzee puts it:

> [T]he metaphysical and epistemic statuses in question can come apart. One can imagine a person/device combination that is well enough integrated in the metaphysical sense to be called a cyborg, while at the same time not conducting themselves creditably epistemically speaking (p. 227).[4]

Following Kotzee (who is following Clark), we can call a person/device combination that satisfies whatever criteria we need to deal with cognitive bloat a "cyborg". Kotzee's point is that it should always be possible for the beliefs a cyborg has to fail to count as knowledge because they are not sufficiently creditable to their cognitive agency. Aizawa images someone who designs an ingenious cycling computer that keeps track of all sorts of cycling data, and who

---

[4] Note that, while Kotzee seems to agree with this, he is actually reporting what he takes Aizawa's view to be here.

satisfies Clark's trust and glue conditions with respect to the computer. But it is surely a further question whether they can gain knowledge by simply trusting what the computer tells them.

Whether Aizawa and Kotzee are right here depends in part on the outcome of the issue I discussed in the previous section. Do the right criteria for genuine cognitive extension (the solution to the cognitive bloat problem) ensure that an agent cannot but be justified in forming any beliefs they arrive at via extended cognitive processes, because the criteria for genuine cognitive extension in effect ensure that the agent's beliefs are in epistemic good standing? I want to set this issue aside here and focus on another point Kotzee makes.

I think it is very plausible that we can have (something like) cognitive integration without genuine extended cognition. Take an example discussed by Kotzee (p. 231). Imagine a skilled user of Google. They can use Google to find out all sorts of things. Because they are skilled at using Google, they don't just click on the first link they find and trust what they read implicitly. They compare and contrast the information they find on different webpages. They are aware that, often, the first few results are actually paid for ads, and they scroll down the page to the "good stuff". They know certain very popular websites are disreputable, and they either don't click on links to them, or exercise appropriate discretion when they do. It seems clear that this person's use of Google does not satisfy Clark's trust and glue conditions (or approximations thereof). But it is plausible that they do deserve credit for the beliefs that they form as a result of their searches, and that they gain knowledge from them.

What this illustrates is that there are all sorts of ways in which we can use artefacts (including technology) to enhance our body of knowledge that don't require genuine cognitive extension. They merely require recognition that cognition is deeply embedded. This sends us back to where we started: what is extended knowledge? In what we can call the "narrow" sense, extended knowledge is knowledge produced by extended cognitive processes. But, in what we can call the "broad" sense, it is knowledge produced by extended *or* embedded cognitive processes. It is worth noting that, in several contributions to both volumes (especially in the respective "Applications" sections), it is unclear whether it matters if we understand "extended knowledge" in the broad or narrow sense. This suggests (at least to me) that there might be less of a link between the metaphysical issue of when cognition extends and the epistemological issue of how (if at all) we can extend what we know using artefacts than one might expect. This prompts a question: from an epistemological perspective, how important is the distinction between extended and embedded cognition?

5. What is "Socially Extended Epistemology?"

While it is a little unclear what makes an epistemology "extended", it is (I submit) even less clear what makes an epistemology "socially extended". In this final section I want to explore this issue in a little more detail.

We can start with the subfield of *social* epistemology. Where epistemology has traditionally studied the epistemic agent in splendid social isolation, socially epistemology investigates the epistemological consequences of social interaction and social systems (see Goldman and Blanchard 2018 for an overview). Key topics in social epistemology include (but are not limited to): testimony, disagreement, epistemic injustice, scientific knowledge, group belief and knowledge. Social epistemology thus investigates (among other things) the epistemological consequences of the social embeddedness of human cognition. This prompts a question: what, exactly, is the difference between social epistemology and socially extended epistemology? The difference, presumably, is that, while social epistemologists agree that human cognition is socially

embedded, they don't generally treat it as socially extended. A socially extended epistemology would treat human cognition as socially extended and not just socially embedded.

There are some contributions that do this. Several contributors address issues in collective epistemology such as group belief and group knowledge. Take, for instance, Orestis Palermos and Deborah Tollefson's "Group Know-How". In the literature on individual know-how, it is standard to distinguish between intellectualism (the view that knowledge-how is a species of knowledge-that) and anti-intellectualism (the view that it isn't: to know-how to x is to possess some sort of ability to x). The intellectualist (following Stanley and Williamson 2001) thinks that knowing how to x is a matter of knowing that some way W is the way to ride a book. The anti-intellectualist (following Ryle 1945) thinks that knowing how to x is a matter of x-ing in a way that exhibits a form of responsibility for one's x-ing that allows one to take ownership of it. Palermos and Tollefson develop a version of the anti-intellectualist view, on which a group knows how to x when responsibility for successful x-ing attaches to the group rather than its individual members.

Palermos and Tollefson think making this account work requires taking on board an idea from the literature on distributed cognition. When groups of individuals interact (continuously and reciprocally) to perform a task that we would classify as cognitive (e.g. making a decision), we can view the group as a cognitive system in its own right. Take the example of a jazz band doing some improv. If one player is out of rhythm, another will try to compensate for it. It is the group as a whole that is responsible for the smooth overall process. This seems to require distributed rather than merely socially embedded cognition. The idea is that the various players form a cognitive system, rather than that each player contributes to the cognitive processing of the other players.

However, other contributions don't need to take a stand on this issue. In her "Socially Extended Moral Deliberation about Risks" Sabine Roeser argues that emotions can contribute to socially extended knowledge of ethical aspects of risk. Specifically, she argues that narrative and other works of art can contribute to collective emotional-ethical reflection about risk. She puts special focus on some forms of art not often discussed by philosophers: bio-art, techno art and the like. Consider the risks of climate change, which has an ethical aspect (e.g. the issue of intergenerational justice). Roeser argues that, in general, emotions are an important source of ethical knowledge, and so they are an important source of knowledge about ethical aspects of risk. But individuals can often go badly wrong when they rely on their emotions. This is why emotion-based reflection about risk must be *collective*. As she puts it:

> [E]motional moral reflection on risky technologies should involve a collective deliberation in which emotions and values are central. By engaging in emotional deliberation with others, people can share others emotions with each other, allowing them to broaden their point of view, expand their perspectives, and achieve better understanding of the values, emotions, and concerns of others, and gain a better and more complete understanding of the evaluative aspects of risky technologies. Hence, such a form of 'emotional deliberation on risk' … can be understood as an emotional form of socially extended knowledge (p. 162).

I find this very plausible. But notice that, for her purposes, Roeser need take no stand on whether this collective deliberation is a form of distributed cognition or (merely) a form of socially embedded cognition. That is, she need take no stand on the question whether we have genuine cognitive extension.

While Roeser isn't explicitly neutral on this question, in her contribution ("Collective Amnesia and Epistemic Injustice"), Alessandra Tanesini is. Tanesini discusses collective amnesia, which

she construes as a sort of motivated (collective) ignorance of the past. It is a common idea in psychology that memories are constructed: they are highly partial and selective rather than a faithful reproduction of past events. But this is adaptive: memory serves the purpose of creating and strengthening social bonds, as well as recording what happened. Sometimes these purposes conflict, and accuracy is sacrificed. The same thing goes at the collective level. Collective memory is also highly partial and selective, and it sometimes sacrifices accuracy for social cohesion. Thus collective amnesia is sometimes the result of collective memory working as it was intended rather than malfunctioning. Tanesini connects this with the literature on epistemic injustice. The basic idea is that collective memory and amnesia enhances the intellectual self-trust of some community members, while undermining the intellectual self-trust of others. Because collective memories converge on the shared narrative promoted by the powerful, and these narratives are often self-serving, they often enhance the intellectual self-trust of the powerful, while undermining the intellectual self-trust of the subordinate. Because our degree of trust in our intellectual abilities and standing concerns us qua epistemic agents, this is a hitherto unrecognised form of epistemic injustice.

Tanesini frames collective remembering as a form of *scaffolded* cognition, where "a scaffold is an epistemic tool, such as a notebook, or an informational niche (like a path or a classroom or a supermarket) which modifies and structures cognition" (p. 204). As she notes, one can view these scaffolds as either constitutive of cognitive processing (extended cognition), or as mere causal triggers to organism-level cognition (embedded cognition). For her purposes, this difference doesn't much matter: they don't change her account of collective amnesia, or impact on the consequences she draws from it vis-à-vis epistemic injustice. Thus, like Roeser, Tanesini need not take a stand on the issue of extended vs. embedded cognition.

I have suggested that, in both Roeser's and Tanesini's contributions, the difference between socially extended cognition and socially embedded cognition is not particularly important. In my view, this applies to other contributions to both volumes (e.g. Paul Smart's "Mandevillian Intelligence", Mark Alfano and Joshua August Skorburg's "Extended Knowledge, the Recognition Heuristic and Epistemic Injustice"). Putting this together, while there are contributions that require genuinely socially extended cognition, others merely require that we view cognition as socially embedded, and draw out the epistemological consequences of this social embeddedness. Thus, many contributions (especially to the second volume) seem better read as contributions to social epistemology than as to "socially extended epistemology", at least if we construe socially extended epistemology in what I called the narrow sense above.

What are we to take from this? I don't put too much stock in labels, and I certainly don't want to downplay the importance or interest of any of the articles summarised above. Many of the articles in the second volume strike me as genuinely novel, exploring issues that are not much discussed in contemporary epistemology. But it is reasonable to expect some rationale for collecting several articles under the umbrella of "socially extended epistemology", and it is not clear that this rationale has much to do with the metaphysics of social cognition. So where is it to be found? One might pursue the following line. While social epistemologists study the epistemological consequences of social interaction, it is fair to say that many practitioners (at least in analytic social epistemology) don't think too much about the deeply social nature of cognition, whether we cash this out in terms of social embeddedness or social extension. One might pursue an epistemological project that does take the social nature of cognition seriously under the banner of a "socially extended epistemology" without worrying too much about how to cash out the metaphysics of social cognition, but in order to signal that one is departing from the norm in much of analytic social epistemology. If this is right, though, then there is further reason to question the existence of a deep connection between the metaphysical issue of when (if

at all) cognition extends and the epistemological issues that are the primary concern of these two volumes.

ROBIN MCKENNA
University of Liverpool
Walnut House, Mulberry Court, Liverpool, L69 7ZY, UK
rbnmckenna@gmail.com
r.j.mckenna@liverpool.ac.uk

References

Burge, Tyler. 1979. "Individualism and the Mental." *Midwest Studies in Philosophy* 4 (1): 73–122.
Clark, Andy. 2010. "Memento's Revenge : The Extended Mind Extended." In *The Extended Mind*, edited by Richard Menary, 43--66. MIT Press.
Clark, Andy, and David J. Chalmers. 1998. "The Extended Mind." *Analysis* 58 (1): 7–19.
Goldman, Alvin, and Thomas Blanchard. 2018. "Social Epistemology." In *The Stanford Encyclopedia of Philosophy (Summer 2018 Edition)*, edited by Edward N. Zalta. https://plato.stanford.edu/archives/sum2018/entries/epistemology-social/.
Parent, Ted. 2017. "Externalism and Self-Knowledge." In *Stanford Encyclopedia of Philosophy (Fall 2017 Edition)*, edited by Edward N. Zalta. <https://plato.stanford.edu/archives/fall2017/entries/self-knowledge-externalism/>.
Putnam, Hilary. 1975. "The Meaning of 'Meaning.'" *Minnesota Studies in the Philosophy of Science* 7: 131–93.
Ryle, Gilbert. 1945. "Knowing How and Knowing That." *Proceedings of the Aristotelian Society* 46 (n/a): 1–16.
Stanley, Jason, and Timothy Williamson. 2001. "Knowing How." *Journal of Philosophy* 98 (8): 411–44.