

Essay Review

Review Socially Extended Epistemology

J. Adam Carter, Andy Clark, Jesper Kallestrup, S. Orestis Palermos, and Duncan Pritchard (eds.), *Socially Extended Epistemology*, Oxford, UK: Oxford University Press, 2018, pp. 336. £55.00HB

Joshua Habgood-Coote

My bike needs its brake cables changed. Annoyingly, the cables are routed through the inside of the frame, and I don't know how to navigate the new cables from the tiny holes at the front of the frame to those at the back. I have a couple of options. I can take the bike to the shop, and pay someone to change them for me. I can look up instructions in my well-thumbed and annotated bike maintenance book, and follow them. Or I can ask my friend to work on the problem together. In the first case, I rely on someone to do something for me, for which they will take full credit. The second case involves a kind of cognitive outsourcing, allowing me to replace the question of how to change the cables to with the question of how to follow the instructions. In the third case, I build a team to solve the problem. Even if my friend doesn't know the solution, together we might possess more bike maintenance knowledge, and we can employ various interactive problem solving techniques that are not available to individuals.

Following Clark and Chalmers (1998), we might think that in the scenario in which I consult my maintenance book my cognitive processes are *extended*, in the sense that the book itself is part of my thinking about how to fix the bike. Even twenty years after this publication, this thesis is still provocative, and its implications have by no means been fully explored. Alongside the extension of cognitive processes, we might want to countenance the possibility of extended knowledge. We might think that if my relation to the book is appropriate — fulfilling Clark's 'trust and glue' conditions (Clark 2008) — then my knowledge stretches to the instructions written in the book. The book is in a real sense part of my mind.

The central question of this volume is how to think about the kinds of epistemic relations involved in pursuing a solution in the third scenario. Social relations like those that occur when I work on the bike with my friend involve extending my thinking so that it occurs in our interactions, creating a new agent with its own knowledge.

This volume is a partner to the collection *Extended Epistemology*, by the same editors, which focuses on general issues about the epistemology of extension. As with that volume, many of the contributions engage in some way with *virtue reliabilism*: the view that a subject knows some proposition when it is the product of a reliable process which is cognitively integrated into her epistemic character. That said, reading the first volume is not a prerequisite, and there is plenty in this volume to interest people who are sceptical about virtue reliabilism.

Brad Wray argues against the application of virtue reliabilism to scientific collaboration, contending that there is an important distinction between scientists' reliance on artefacts (which he concedes is an example of extension) and their reliance on other scientists. Wray observes a dilemma in Pritchard's version of virtue reliabilism (2010), which claims that knowledge requires one agent to possess a true belief and another agent to take responsibility for that belief. Either we think of cases of collaborative science as involving individual extended knowledge—so that all members of the team can legitimately claim knowledge which is the output of collective projects—or we think of them as involving collective knowledge of these results. The individual-level extension implausibly allows individual scientists to take credit for one another's successes. Ascribing collective knowledge cases creates problems for the responsibility condition, since scientific practice tends to involve individuals taking all responsibility. Wray's conclusion is that cases of scientific collaboration are a significant issue for Pritchard's virtue epistemology.

Cathal O'Madagain explores the connection between semantic externalism and cognitive extension. The semantic externalist thinks that the meaning of certain terms — like 'quaking aspen' — is determined not by descriptive beliefs in the head of a speaker, but by stuff going on outside the speaker's head. According to the Putnam-Burge version of semantic externalism, the external ingredient that determines the meaning of our terms is the discriminatory abilities of other people. So, 'quaking aspen' in my mouth refers to quaking aspen, because I defer to expert horticulturalists who are able to discriminate quaking aspen from other trees. O'Madagain suggests that we might think about this kind of semantic deference as a kind of social extension, meaning that a deferential speaker knows what her terms mean in virtue of other peoples' semantic knowledge.

To my mind, the difficulty with this proposal lies with the 'trust and glue' conditions on extended knowledge (Clark 2008), which are introduced in order to avoid a problem of cognitive bloat. Clark introduces four conditions for information to be part of a subject's extended mental states: (i) typical use, (ii) easy accessibility, (iii) automatic endorsement, and (iv) past endorsement. All of these conditions are in question in cases of semantic deference: as an ordinary user of a language I might be disposed not to invoke expert semantic knowledge (i), the experts might be difficult to get hold of (ii), I might have dogmatic false beliefs about what distinguishes quaking aspen (iii), and I may have never invoked or endorsed the expert claim (iv). O'Madagain argues that conditions (ii) and (iv) are not plausible requirements on extended semantic knowledge, but the failure of all four conditions does leave the analogy on somewhat shaky ground.

Jeroen de Ridder's essay responds to a concern about the reality of collective attitudes raised by Paul Thagard (2010). Thagard contends that attitudes require the presence of representations — which de Ridder understands as mental imagery, linguistic, or conceptual structure — points out that group agents do not possess representations, and concludes that groups do not have collective attitudes. De Ridder considers two ways to respond to the argument. The first is to endorse a dispositional account of belief, denying that attitudes require representations. The second is to concede that attitudes do require representations, but to contend that a group's mental life extends to that of its members. On this story, a group's attitudes *are* accompanied by mental representations, it's just that the representations take place in the minds of its members. De Ridder takes up the second response, meaning that

groups have mental states that extend into their members' mental lives.

Francis Heylighen and Shima Beigi argue for an extreme version of the extended mind thesis, endorsing the panpsychist view that “mind is a ubiquitous property of all minimally active matter,” (60). To develop this picture, they lay out an action ontology which thinks of matter not as a static substance, but as a set of dynamic processes involving hierarchically organised self-maintaining systems. From the perspective of this background metaphysics, all matter is already a-thinking, so it is no surprise that a human individual's thought should encompass parts of her environment, including social processes.

The next two pieces shift to collective knowledge-wh, focusing on knowledge of what one is doing, and how one is doing it. Olle Blomberg's contribution concerns whether a group has collective knowledge of *what* it is doing. According to a popular Anscombean paradigm in action theory, when an agent is acting intentionally she has non-observational knowledge of what she is up to. This knowledge is often thought of as practically grounded, although how to make sense of this idea is contested (see Ford Hornsby and Stoutland 2011). Blomberg's essay is admirably ambitious, bringing together the literatures on agent's knowledge and collective knowledge.

Blomberg considers three possible views: i) a distributive view which claims that a group knows what it is doing if each member knows what the group is doing, ii) a shared intention view which claims that group action involves all actors possessing a single token intention, and iii) a robustly collective view according to which a group knows jointly what it is collectively up to. Ultimately, he argues that each of these approaches ends up either endorsing unattractive claims, or failing to be fully general. Given the different ways in which an Anscombean view has been developed, I read this chapter as a call to action for defenders of Anscombean approaches to develop an account of how to think about agents' knowledge in the collective case.

Orestis Palermos and Deborah Tollefsen consider the nature of group knowledge-how. Bird (2010) points out that a group might know how to do something which none of its members individually know how to do. For example, NASA might know how to make a space shuttle, while all of its members only know how to perform some relevant subtask. (Palermos and Tollefsen claim that this divergence judgement is counterintuitive — see their discussion of knowing how to make a Corvette (115) — but survey evidence suggests that the folk are happy to endorse this judgement, see (Jenkins et al 2014)).

They consider three possible accounts: an identity theory which claims that a group's knowledge-how is identical to the sum of knowledge-how possessed by its members, a distributive intellectualist theory which claims that a group's knowledge how to V is composed by each of its members knowing of some way w that it is a way to V, and their preferred anti-intellectualist account which identifies group knowledge-how with a disposition to engage in responsible performances which are grounded in dynamical properties of the group agent. This view which has interesting connections with David Lowenstein and Benjamin Elzinga's Rylean approaches to individual knowledge-how, which both stress the importance of responsible performance (Lowenstein 2017, Elzinga 2018).

Reading Blomberg's and Palermos and Tollefsen's pieces, I wondered whether they were considering all possible views. Philosophers typically discuss the distributive and collective readings of predicates applied to collectives, but linguists also discuss the *semi-distributive* or *cumulative* reading. Consider the following sentence (adapted from Lahiri 2002 p189):

1. The witnesses know which children were present at the fireworks display.

The distributive reading of (1) distributes the predicate across the members of the group, claiming that all of the witnesses know which children were present. The collective reading claims that the predicate is a property of the collective, but not necessarily any of its members (much like 'surrounded the house'). By contrast, the semi-distributive reading claims that some members of the group between them *contribute* to the fulfilment of the predicate without any individual completely fulfilling the predicate. This reading is salient in a case in which none of the witnesses saw every child who was present, each saw that some sub-set of the children were present, and combining these partial lists adds up to a complete list. The semi-distributive reading of an ascriptions of agent's knowledge claims that a group knows what it is up to in virtue of its members knowing what they are doing where this adds up to what the group is doing, and this reading of a knowledge-how ascription claims that a group knows how to do something when its members know how to do relevant sub-activities. It would be an interesting project to think through what accounts of collective agent's knowledge and knowledge-how that exploited this reading would look like.

Joelle Proust considers what kinds of collective attitudes group agents can possess. She argues against Gilbert's popular joint acceptance account, and the view that group agents can accept propositions. She worries about whether these attitudes are governed by suitable norms of rationality (also a theme in Lackey 2016), and proposes that groups can only have the attitude of acceptance under consensus.

Sabine Roeser considers the role that art works about technology can play in socially extended reasoning about risks. She stresses that deliberation about risks need not only be carried out in cold deliberative terms, and can involve emotional engagement with art objects like fluorescent rabbits, and work focusing on climate change (as I write, one hundred tonnes of Artic ice are slowly melting in central London in an installation conceived by Olafur Eliason and Mink Rosing to coincide with COP24).

Holly Arrow and Alexander Garinther offer an overview of the ways in which social psychologists think about socially shared cognition, which they apply to examples drawn from Jean Hatzfield's work with survivors of the 1994 Rwandan genocide. They distinguish a number of different levels of sociality involving different collective mechanisms, ranging from dyads engaged in joint attention, to macrobands of several hundred that construct a shared identity. This was an extremely helpful overview of the social psychology literature offering a range of tools for epistemologists thinking about social cognition. I did find it a little surprising that this contribution leant solely on Hatzfield's work. The Rwandan genocide is a complex and contested subject, and it would have been good to see a more diverse range of sources being used, including work by Rwandans.

Alessandra Tanesini flips around the virtue reliabilist's approach to knowledge, considering the ways in which *ignorance* can be a cognitive achievement. Drawing on evidence from psychology, Tanesini proposes a social account of collective memories, according to which the function of memories is partly to solidify social bonds, and the formation and maintenance of memories is socially filtered. She argues that memories are mediated by a process of *socially-shared retrieval induced forgetting*, whereby the retrieval of certain features of an event leads to other features fading into obscurity. For example, public discourse focusing on the positive features of the British Empire which might make it harder for people to recall its numerous genocides. This focus on ignorance helps us to see the dark side of the social scaffolding of cognition, including the ways in which epistemic technologies and informational niches can lead to epistemic failings. Tanesini distinguishes three kinds of memory ignorance: self-deceptive memories that serve an individual or group's self-conception, retrieval-induced forgetting via social processes, and the construction of warped informational niches. She argues that these processes can lead to epistemic injustice, undermining the self-trust of people groups targeted by this miscognition.

Georg Thiener and Nikolaus Fogle make the case for thinking of Bourdieu as an externalist about knowledge. They argue that his notion of *habitus* picks out a kind of environmentally dependent and socially tuned knowledge which we can think of as a kind of extended knowledge-how. They develop the social nature of the habits by considering how to respond to Turner's challenge to give an account of how practice gets transmitted, developing three mechanisms of transmission: diachronic co-dependence, instrumental dependence, and co-constitutive dependence.

Paul Smart considers cases where individual and collective epistemic virtues come apart, in particular cases where individual epistemic vice might contribute to collective virtue. This phenomenon is an example of what Mayo-Wilson, Zollman, and Danks (2011) call *the Independence Thesis*: the view that methodological prescriptions for individuals and groups might come apart. He considers four examples of this phenomenon: Zollman's claim that dogmatism contributes to useful network structures (2010), Solomon's work on the virtues of heuristics and biases (1992), Nickerson's work on confirmation bias (1998), and work on impure motivations of scientists (Zollman 2018). These phenomena are fascinating, and deserve further study.

Harry Halpin considers how the exploitation of social processes by algorithmic processing can help to solve the frame problem, giving a nice history of the transition in computer science from seeing computers as independent symbol-processing machines to extensions of human capabilities, and a focus on human-computer interactions.

Throughout the volume, I was puzzled about how to think about the relation between social extension and collective mental states. The contributions seemed to be concerned with two different and connected senses of the social extension of knowledge.

The first sense is the idea that one person's mental life can be broadened or extended through her relations with other people, meaning that her mental life depends on her social position. This sense involves an extension of knowledge, but not of the knowing subject: individuals know more in virtue of their social position.

The second sense is the idea that social relations can influence our mental lives by creating supra-human entities that have distinct mental properties. This is a phenomenon of *collective mentality*, involving an extension not just of what is known, but of the subject which is the knower.

To fix the two senses of social extension, consider an analogy: trees often depend on one another for nutrients via subterranean transfer. There are two ways we might think about this dependence: that trees stretch their nutritive capacities by depending on one another (individual dependence on social factors), or that the dependence involves a new being constituted by all of the trees together (collective states).

Throughout the volume, different contributions are concerned with different senses of extension. O'Madragain is concerned with social dependence in individual knowledge of meaning, and Theiner and Fogal with the social dependence of habitus. By contract de Ridder, Tollefsen and Palermos are concerned with collective knowledge, and Proust with collective acceptance. Others – such as Wray – contrast individuals' social dependence with collective knowledge. We see this split in the wider literature. Tollefsen and Palermos argue that social extension and collective mental states are just different names for the same phenomenon (Tollefsen 2005, Palermos 2015), whereas Sandford Goldberg uses 'social extension' to refer to ways in which individuals' epistemic states depend on their social position (Goldberg 2012, forthcoming).

How do individual social dependence and collective knowledge relate to one another? In his early discussion of scientific collaboration, Hardwig suggests that extended justification and collective knowledge are competing explanations of the same phenomenon (Hardwig 1985: 348-9). Extreme versions of the two views might be in competition: if everyone in a group already knows what everyone else does due to a social dependence, this weakens the case for ascribing knowledge to the collective. Another possibility is that social extension and collective knowledge are distinct but related phenomena. This raises the question of how the two phenomena relate. Perhaps individuals can extend their mental states by relying on the outputs of collective agents (although see Boyd forthcoming on groupstrapping), and perhaps *collective* agents can extend their cognition by relying on individual mental states (as de Ridder argues in this collection).

References

- Bird, Alexander (2010). Social knowing: The social sense of 'scientific knowledge'. *Philosophical Perspectives* 24 (1):23-56.
- Boyd, Kenneth (forthcoming). Epistemically Pernicious Groups and the Groupstrapping Problem. *Social Epistemology*:1-13.
- Clark, Andy (ed.) (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press.
- Clark, Andy & Chalmers, David J. (1998). The extended mind. *Analysis* 58 (1):7-19.
- Elzinga, B. (2018). Self-Regulation and Knowledge-How. *Episteme*, 15(1), 119-140. doi:10.1017/epi.2016.45
- Ford, Anton ; Hornsby, Jennifer & Stoutland, Frederick (eds.) (2011). *Essays on Anscombe's Intention*. Harvard University Press.

- Goldberg, Sanford C. (2012). Epistemic extendedness, testimony, and the epistemology of instrument-based belief. *Philosophical Explorations* 15 (2):181 - 197.
- Goldberg, Sanford C. (forthcoming). Epistemically engineered environments. *Synthese*:1-20.
- Hardwig, John (1985). Epistemic dependence. *Journal of Philosophy* 82 (7):335-349.
- Jenkins, Anna, Dodell-Feder, David, Saxe, Rebecca & Knobe, Joshua (2014). The Neural Bases of Directed and Spontaneous Mental State Attributions to Group Agents. *PLoS ONE* 9.
- Lackey, Jennifer (2016). What Is Justified Group Belief? *Philosophical Review Recent Issues* 125 (3):341-396.
- Lahiri, Utpal (2002). *Questions and Answers in Embedded Contexts*. Oxford: OUP.
- Löwenstein, David (2017). *Know-how as Competence. A Rylean Responsibilist Account*. Frankfurt am Main: Vittorio Klostermann.
- Mayo-Wilson, Conor ; Zollman, Kevin J. S. & Danks, David (2011). The Independence Thesis: When Individual and Social Epistemology Diverge. *Philosophy of Science* 78 (4):653-677.
- Nickerson, R (1998) Confirmation Bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2): 175-220
- Palermos, S. O., (2015). Active externalism, virtue reliabilism and scientific knowledge. *Synthese* 192 (9):2955-2986.
- Pritchard, Duncan (2010). Cognitive ability and the extended cognition thesis. *Synthese* 175 (1):133 - 151.
- Solomon, Miriam (1992). Scientific rationality and human reasoning. *Philosophy of Science* 59 (3):439-455.
- Thagard, Paul (2010). Explaining Economic Crises: Are There Collective Representations? *Episteme* 7 (3):266-283.
- Tollefsen, Deborah (2006). From extended mind to collective mind. *Cognitive Systems Research* 7 (2):140-150.
- Zollman, Kevin J. S. (2010). The Epistemic Benefit of Transient Diversity. *Erkenntnis* 72 (1):17-35.
- Zollman, Kevin J. S. (2018). The Credit Economy and the Economic Rationality of Science. *Journal of Philosophy* 115 (1):5-33.