Extended Knowledge and Social Epistemology

Orestis Palermos and Duncan Pritchard, University of Edinburgh

Abstract

The place of social epistemology within contemporary philosophy, as well as its relation to other academic disciplines, is the topic of an ongoing debate. One camp within that debate holds that social epistemology should be pursued strictly from within the perspective of individualistic analytic epistemology. In contrast, a second camp holds that social epistemology is an interdisciplinary field that should be given priority over traditional analytic epistemology, with the specific aim of radically transforming the latter to fit the results and methodology of the former. We are rather suspicious of this apparent tension, which we believe can be significantly mitigated by paying attention to certain recent advances within philosophy of mind and cognitive science. Accordingly, we attempt to explain how *extended knowledge,* the result of combining active externalism from contemporary philosophy of mind with contemporary epistemology, can offer an alternative conception of the future of social epistemology.

Introduction

Social epistemology is in many ways still an emerging field. Its aim, broadly construed, is to study the social dimensions of knowledge acquisition and information processing. As is the case with any nascent paradigm, however, its methodology, and thereby its place within contemporary philosophy (as well as its relation to other academic disciplines), is the topic of an ongoing debate. One camp (Goldman 1999; 2010) within that debate holds that social epistemology should be pursued strictly from within the perspective of individualistic analytic epistemology. In contrast, a second camp (Fuller 2007; 2012) holds that social epistemology is an interdisciplinary field that should be given priority over traditional analytic epistemology, with the specific aim of radically transforming the latter to fit the results and methodology of the former.

We are rather suspicious of this apparent tension, which we believe can be significantly mitigated by paying attention to certain recent advances within philosophy of mind and cognitive science. Accordingly, in §1, we will go through what we consider to be the most important aspects of the debate regarding the status of social epistemology. This will make apparent the relevance of *active externalism* from within contemporary philosophy of mind that we will introduce in §2. Finally, in §3, we will explain how *extended knowledge,* the result of combining active externalism with contemporary epistemology, can offer an alternative conception of the future of social epistemology.

Social Epistemology

William Alston (2005) opens *Beyond "Justification"* by asking what counts as epistemology in general. The reason Alston gives priority to this broader question is not to provide a positive response but rather to make clear that the answer is particularly elusive. Within the history of philosophy, as he notes, "thinkers were engaged in what *we* tend to call 'epistemology' long before anyone applied that label to what they were doing or, indeed, distinguished these efforts from other intellectual inquiries by any designation whatever." (Alston 2005, 1) Accordingly, we shouldn't be surprised that we are at a loss about how to pick out the 'purely epistemological' issues from this "heterogeneous grab bag of disparate materials." A "heterogeneous grab bag" that, Alston notes, is not only intimately connected with cognitive psychology but is best classified as such. (*ibid*., 2)

Within contemporary epistemology, however, where some of the topics seem to depart from the vicinity of cognitive psychology, things can get even trickier still. A case in point is the recent trend of social epistemology. On one hand, by conceiving of knowledge — typically the primary focus of epistemology — as a cognitive (i.e., mental) phenomenon, traditional epistemology has traditionally focused on the individual cognitive agent. Cognition, after all — it is largely held — rests within the individual's head. Accordingly, to account for knowledge, one should focus on the cognitive/epistemic properties of the individual agent. On the other hand, social epistemology conceives of knowledge primarily (and, sometimes, even entirely) as a social phenomenon that cannot be understood in the absence of the individual's socio-epistemic interactions. As such, it appears to be in stark contrast with the aforementioned traditional approach.

This is why Alston further makes the passing remark that certain parts of Alvin Goldman's *Knowledge in a Social World,* where some of its chapters explore the effects of social interactions and organizations on the quest for knowledge and well-formed belief, would probably be rejected by many contemporary epistemologists as "not real epistemology." (Alston 2005, 5) This claim, of course, could not have gone unnoticed by Goldman. Goldman, after all, is a leading proponent of reliabilism, which most contemporary epistemologists would readily classify as part of analytic epistemology.[1] Accordingly, in an attempt to reply to Alston's remark, Goldman (2010) has recently distinguished between two construals of social epistemology, only one of which he is happy to accept as "real epistemology."

Goldman's method for testing whether some approach counts as real epistemology comes in two steps. He begins by putting forward what he considers to be the "central features of epistemology 'as we know it' (in the analytic fashion)." (Goldman 2010, 3) The primary and probably the most important tenet that he provides is that (a) "the

---

[1] Briefly, according to reliabilism, knowledge is the product of a reliable belief-forming process. We focus on this account of knowledge in the section to follow.

epistemic agents of traditional epistemology are exclusively individuals." (*ibid*., 2) Even though we will here focus on (a), it should be helpful to state the rest of the features as well: (b) epistemology is a normative domain focusing on evaluative concepts such as justifiedness, rationality, and knowledge; (c) those evaluative concepts are not conventional or relativistic; (d) knowledge and justification either entail or are closely related to truth; (e) truth is objective and mind independent; and finally (f) the central task of traditional epistemology is the critical examination of doxastic decision making. Then, having so defined analytic epistemology, Goldman goes on to judge whether some perspective counts as real social epistemology by considering its degree of departure from the above core tenets.

Accordingly, 'Revisionism'— the first general approach that Goldman considers, and which takes its name from its attempt to revise epistemology as we know it — fails the test. The reason is that this approach studies doxastic attitudes embedded in their social contexts by relying on movements in postmodernism, social studies of science, and cultural studies, which not only fail to illuminate the nature and conditions of conventional epistemic concepts, but which also "generally seek to debunk or reconfigure" them. (*ibid*., 1)

In contrast, the second general approach to social epistemology, which Goldman further subdivides into 'preservationism' and 'expansionism', seems to fare better. Preservationist social epistemology studies themes such as testimonial knowledge and peer disagreement by focusing primarily on the individual epistemic agent, thereby *preserving* the core assumptions of analytic epistemology. Expansionist social epistemology, however, "ventures a bit further" than preservationism, "into unfamiliar terrain." (*ibid*., 3) It *expands* its interests to topics such as the epistemic properties of group doxastic agents and the influence of social systems and their policies on epistemic outcomes, possibly distancing itself from some of the core individualistic assumptions of mainstream epistemology. Nevertheless, according to Goldman, both preservationism and expansionism pass the 'real epistemology' test — though, notice, possibly not with the same credentials. Parts of expansionist social epistemology could be particularly problematic, as they seem to go against Goldman's first tenet of analytic epistemology (a worry that Goldman passes in silence).

Nevertheless, leaving the above worry for later on (we will see how it can be addressed in the sections to follow), Goldman's approach to demarcating real social epistemology has already raised other concerns. Specifically, Steve Fuller (2012) takes Goldman's 'real social epistemology' to be fundamentally mistaken. Of course, this could hardly come as a surprise. Fuller's general approach rejects most, if not all, of Goldman's core tenets of analytic epistemology, and would thereby be categorized (on Goldman's view) as offering a revisionist form of social epistemology.[2] We will

---

[2] Fuller appears to reject (a), (c), (d) and possibly (e) of Goldman's core tenets of analytic epistemology as listed above.

bracket Fuller's remarks on the concept of truth and its relation to knowledge, as well as his views about the role of epistemology. Instead, we will here concentrate on three particular criticisms that Fuller offers of Goldman's conception of 'real' social epistemology (i.e., analytic social epistemology).

To start with, Fuller first rejects Goldman's reliabilist approach to epistemology on the grounds that it is a mechanistic view that misses our sentient nature:

> Anything calling itself "epistemology" — including "social epistemology — not concerned with the formation of beliefs and only examining reliable processes for arriving at the truth provides no more than an account of knowledge fit for androids not humans — that is to say, an epistemology where all the action occurs without the mediation of consciousness (Fuller 2012, 269).

Second, he rejects Goldman's core individualism (which, remember, may turn out to be in any case problematic with respect to certain topics of expansionist social epistemology, such as the study of group doxastic agents):

> I have spoken of an 'outside in' vs. an 'inside out' orientation towards the problem of knowledge [Fuller 2007]. In keeping with the original premise of the book *Social Epistemology,* I start with the existence of knowledge as a social phenomenon — defined primarily in textual terms — as something that is already given in the world. I then proceed to determine how it is possible that diversely and imperfectly informed individuals could have organized themselves to produce such an authoritative body of work that exerts normative force in precincts far beyond the sites of knowledge production itself. In short, my focus has been on the social construction of epistemic standards, assuming that they arise from and are largely maintained by processes that are relatively indirect to the desires and capacities of the relevant knowers (Fuller 2012, 276).

Finally, his third criticism of analytic social epistemology is that it is an isolated academic discipline: "My version of social epistemology has been unique in conceiving of the field as inherently interdisciplinary, with the specific aim of transforming epistemology." (*ibid*., 271) And, interestingly, one of the relevant disciplines that Fuller singles out is the field of psychology, by noting the following:

> My interest in this field has always been, I believe, similar to that of Karl Popper (whose PhD was in educational psychology); namely, to arrive at an account of knowledge that acknowledges at once the depth of our natural liabilities and aspirations to transcend them artificially (Fuller 2012, 271).

Fuller's interest in psychology is reminiscent of Alston's point with regards to the intimate connection between cognitive psychology and what we may now call 'epistemology'. To those familiar with contemporary philosophy of mind and cognitive science, however, his last remark may further bring to mind the emerging current of *active externalism,* according to which human cognition is not always limited to our biological capacities, but often exceeds them by extending to the epistemic artifacts (or even agents) that we interact with.

Crucially, we think that this is more than mere coincidence. To explain why, for the remainder of the paper, we will focus on active externalism and its potential relation to contemporary epistemology. By doing so we will attempt to demonstrate that the choice (for any socially inclined epistemologist) need not be between Fuller's revisionism and Goldman's core individualism.

Extended Knowledge

As noted above, in conceiving of knowledge as a cognitive phenomenon mainstream epistemologists have tended to focus on the individual as the proper epistemic subject. Given that cognition is widely thought to rest within the individual's head, it follows that mainstream epistemology thus leads to a specific form of epistemic individualism, whereby the focus is the individual knower and the cognitive processes which take place under her skin.

This last claim, however, has lately been called into question by recent advances within philosophy of mind and cognitive science, and especially the currents of embodied cognition and active externalism (to be distinguished, as we will soon explain, from the meaning or 'passive' externalism associated with the work of philosophers like Hilary Putnam and Tyler Burge). To get a grip on what all these theories of mind amount to and how they differ from each other it would be helpful to consider their motivations in a chronological order of appearance. Before proceeding further, however, we should note that their common denominator, and hence the reason why we here group them together, is that they all deny the claim that cognition resides entirely within the individual's head.

Now, the first blow to the approach of internalism — the idea that a complete understanding of our minds can be achieved by an exclusive focus on our brains— came from meaning (or passive) externalism, which shows that some mental contents fail to supervene on intrinsic facts (i.e., facts that pertain solely to our brains); consequently, the opposite of internalism, *viz*., externalism about our minds, must be true. Specifically, according to this form of externalism, in order to have certain types of intentional mental states, such as beliefs and desires, it is necessary to be related to the environment in the right way. [3] Accordingly, studying our brains in isolation to

---

[3] The classical argument for meaning externalism involves the following thought experiment: Imagine a remote planet, Twin Earth, which is exactly like Earth, except that instead of water (H2O) it has a different substance, twin- water. Even though twin-water is a different chemical compound, say XYZ,

their natural (Putnam 1975) and social (Burge 1986) environments would necessarily be incomplete.

Not much later, however, several cognitive scientists and philosophers of mind (e.g., Varela, Thomson & Rosch 1991; Clark 1997), as well as roboticists (e.g., Brooks 1991*a*; Brooks 1991*b*), noted that not even the study of our brains as embedded in their environments is enough. This is because cognition is not just embedded but also embodied in the sense that aspects of the agent's body beyond the brain play a physically constitutive role in cognitive processing (i.e., literally speaking, with respect to several cognitive operations, our bodies are parts of our minds). In particular, for those aspects of an agent's mind that his/her brain is heavily *interdependent* with his/her body, we should think of the latter as a constitutive element of the agent's overall cognitive system. [4] According to embodied cognition, then, considerations pertaining to the agent's body as well as its interaction with his/her brain (and central nervous system) are essential for a complete understanding of the human mind.

Now, active externalism, as represented by the extended and distributed cognition hypotheses, is the extreme consequent of the approach of embodied cognition. We should note, however, that we here say 'extreme' only because of its radical conclusions. Indeed, for some it may be counterintuitive to accept that cognitive systems extend beyond our organisms to the artifacts we mutually interact with or

---

its macro properties are just like those of water: it looks and tastes like water, it can be found in the rivers and oceans on Twin Earth, and so on. Furthermore, imagine two intrinsically identical individuals: S who lives on Earth and twin-S who lives on Twin Earth, neither of whom knows anything about chemistry. Now, when S utters "water quenches thirst" he is expressing his belief that water quenches thirst, a belief that is true if and only if H2O quenches thirst. To the contrary, having always encountered twin-water and never having encountered or heard of water, when twin-S utters "water quenches thirst" our intuition dictates that he does not believe that water quenches thirst. Instead, twin-S expresses the belief that twin- water quenches thirst, a belief with different truth-conditions. Accordingly, we seem to have two intrinsically identical individuals who nevertheless have different beliefs, which means that some beliefs do not supervene on intrinsic facts. Therefore meaning externalism must be true. Similar arguments can be construed with respect to one's social environment.

[4] This interdependence is usually cast out in terms of "sensorimotor dependencies." Consider the following passages:

> "The basic claim of the enactive approach is that the perceiver's ability to perceive is constituted (in part) by sensorimotor knowledge (i.e., by practical grasp of the way sensory stimulation varies as the perceiver moves)" (Noë 2004, 12).
> "[*Perception*] is not a process in the brain, but a kind of skillful activity on the part of the animal as a whole" (Noë 2004, 2).
> "Perception is not something that happens to us or in us, it is something we do" (Noë 2004, 1).

Sensorimotor dependencies are relations between movements or change and sensory stimulation. It is the practical knowledge of loops relating external objects and their properties with recurring patterns of change in sensory stimulation. These patterns of change may be caused by the moving subject, the moving object, the ambient environment (e.g., changes in illumination), and so on. For more recent elaborations of the idea, see Hurley & Noë (2003), Noë (2003; 2004), Gallagher (2005), and Chemero (2009).

that cognitive processing may be distributed amongst several individuals and their artifacts. The spirit of the approach, however, is very similar to, if not the same as, that of embodied cognition. If we are willing to accept that our minds are embodied when our brains are heavily interdependent with our bodies, then there is no principled reason to deny that cognitive processes may be extended or even distributed in those situations that our brains are heavily interdependent with the artifacts we employ or the other agents we interact with. In fact, active externalism has been developed, refined, and defended by many philosophers (Clark & Chalmers 1998; Clark 2007, 2008; Hutchins 1995; Wilson 2000, 2004; Wheeler 2005; Menary 2006, 2007; Theiner 2011).[5] Accordingly, active externalism, in both of its forms, is a viable hypothesis that we believe can generate several ramifications within analytic epistemology, some of which might be particularly interesting within the context of the present discussion.

So how can we introduce active externalism within epistemology? The most obvious starting point is the approach of process reliabilism, due to its central focus on the reliability of epistemic *mechanisms, methods, and processes:*

> *Process Reliabilism*
> S knows that *p,* if and only if *S's* true belief is the product of a reliable belief-forming process.

Clearly, a formulation of knowledge along the lines suggested above is well suited to accommodate considerations originating from philosophy of mind and cognitive science whose aim is to provide a mechanistic understanding of the *cognitive processes* that constitute the machinery of the human mind.

Remember, however, that Fuller rejects the general approach of reliabilism on the grounds that it misses our sentient nature. By merely focusing on reliable processes we can go only so far as to provide an account of knowledge fit for automata. In other words, although process reliabilism is a good staring point for naturalizing epistemology, it is not targeted enough; knowledge is supposed to be a cognitive phenomenon, whereas process reliabilism allows *any* reliable belief-forming process to count as knowledge-conducive.

Process reliabilism, however, is not the whole story within the reliabilist approach. In order to flag the shortcomings of process reliabilism, contemporary epistemologists

---

[5] Which is not to say, of course, that the view is without its critics. The discussion of the objections facing active externalism is well beyond the scope of the present paper, however. Nevertheless, in brief, many of them point either to the dissimilarity between the inner cognitive processes and the external elements that are supposed to be parts of one's cognitive system (e.g., Rupert 2004; Adams & Aizawa 2008), or to the perceptive rather than introspective manipulation of those external elements. Others deny the mark of the cognitive to the alleged extended cognitive processes (e.g., Adams & Aizawa 2008), or claim that there cannot be a science of active externalism (e.g., Rupert 2004; Adams & Aizawa 2008). For a short discussion and reply to most of these objections, see Menary (2006).

have accentuated the importance of what has come to be known as the *ability intuition* on knowledge — *viz*., the idea that knowledge is belief that is true in virtue of the manifestation of cognitive ability (e.g., Sosa 1988, 1991, 2007, 2009; Plantinga 1993; Greco 1999, 2004, 2007; Pritchard 2009, 2010*a*, 2010*b*, 2012; Haddock, Millar & Pritchard 2010, chs. 1-4).[6] Upgrading process reliabilism with the ability intuition on knowledge gives rise to virtue reliabilism, which is supposed to encompass all of its predecessor's advantages while limiting the knowledge-conducive processes to only the genuinely cognitive ones. Specifically, virtue reliabilists hold that a belief-forming process counts as knowledge-conducive only if it is one of the reliable processes that make up or have been integrated into the agent's *cognitive character*.

According to this view, one's cognitive character usually consists of one's organismic cognitive faculties, one's memories and, in general, one's overall doxastic system. In addition, however, it can also consist of acquired methods of thought, "acquired skills of perception and acquired methods of inquiry, including those involving highly specialized training or even advanced technology." (Greco 1999, 287)

Now, within the literature, there are several variations of virtue reliabilism that compete with each other on how well they fare with respect to several thought experiments that test our intuitions on what may count as knowledge. To keep the dialectics simple, however, we will here present only a necessary virtue reliabilistic condition on knowledge that, arguably, is free from any known counterexamples.[7]

> *COGA_{weak}*
> If *S* knows that *p,* then *S*'s true belief that *p* is the product of a reliable belief-forming process, which is appropriately integrated within *S*'s cognitive character such that her cognitive success is to a significant degree creditable to her cognitive agency (Pritchard 2010*b*, 136-7).

---

[6] The ability intuition on knowledge was, initially at least, often introduced in order to do away with knowledge undermining epistemic luck:
> "To say that someone knows is to say that his believing the truth can be credited to him. It is to say that this person got things right due to his own abilities, efforts and actions, rather than due to dumb luck, or blind chance, or something else" (Greco 2004, 111).

[7] For a full virtue reliabilistic account of knowledge see Greco (1999; 2010). COGA_{weak} is formulated as only a necessary condition on knowledge because, arguably, it cannot accommodate all cases of knowledge-undermining luck. Accordingly, Pritchard argues elsewhere that it must be supplemented with an anti-luck condition on knowledge such as the safety principle. Consider for example Anti-Luck Virtue Epistemology: S knows that p if and only if S's safe belief that p is the product of her relevant cognitive abilities (such that her safe cognitive success is to a significant degree creditable to her cognitive agency). (Pritchard 2012, 20) Again, in Pritchard (2010a, 76) we can read:
> "Knowledge is safe belief that arises out of the reliable cognitive traits that make up one's cognitive character, such that one's cognitive success is to a significant degree creditable to one's cognitive character."

For a defense of the claim that virtue reliabilism can provide a full account of knowledge see Palermos (forthcoming).

Obviously, COGA$_{weak}$ accommodates the ability intuition on knowledge by claiming that a reliable belief-forming process is knowledge-conducive only if it has been *appropriately integrated* within the agent's cognitive character. What is of further interest to our present purposes, however, is that there is nothing in the formulation of COGA$_{weak}$ or in the concepts involved thereof that restricts knowledge-conducive cognitive abilities to those cognitive processes within the agent's head. To the contrary, the idea of a cognitive character that may consist of "acquired methods of inquiry including those involving highly specialized training or even advanced technology" seems to be compatible with, or even prefigure, the hypothesis of extended cognition.

In fact, we have both argued in the past that there seems to be a nice fit between the two views (Pritchard 2010*b*) or that, stronger, both views put forward the same conditions in order for a process to count as a cognitive ability (and thereby, according to virtue reliabilism, as knowledge- conducive) (Palermos 2011). Briefly, both views state that in order for a process to count as a cognitive ability it must be (a) reliable (i.e., not subject to critical scrutiny), (b) one of the agent's dispositions (such that it will be typically invoked), and (c) integrated within the rest of the agent's cognitive system/character (such that it will be easily accessible as if it were part of the agent's organismic cognitive apparatus). In addition, both theories understand the central idea of cognitive integration in terms of mutual interactions with other aspects of the agent's cognitive system (Clark 2010; Greco 2010; Palermos 2011).

Now, before moving on, two important remarks are in order here. First, the requirement for a belief-forming process to be integrated into the agent's cognitive character is a straightforward answer to Fuller's worry that mere process reliabilism pays no attention to our sentient natures. In contrast to process reliabilism, virtue reliabilism is not interested in isolated automatic processes that reliably generate true beliefs, but in the interwound totality of them that constitutes the agent's cognitive character. Possessing such a complex cognitive character may still be insufficient for enjoying consciousness, but it is hard to see how a being with a cognitive system that is functionally equivalent to ours would count as a mere android that falls short of generating knowledge. Second, since both virtue reliabilism and the extended cognition hypothesis put forward the same or at least very similar conditions for a process to count as a cognitive ability, we get the following two very interesting possibilities with respect to the ways in which knowledge may be produced.

First, we can claim that our knowledge-conducive cognitive characters may, and indeed many times do, extend beyond our organismic cognitive faculties. For instance, interpreting COGA$_{weak}$ along the lines suggested by the hypothesis of extended cognition, we can explain how a subject might came to know the position of a satellite on the basis of a telescope while remaining fast to the ability intuition on knowledge. Even though the belief-forming process in virtue of which the subject formed his true belief is for the most part external to his organismic cognitive agency, it still counts as one of his cognitive abilities as it has been appropriately integrated within his cognitive character. Moreover, the subject's belief satisfies COGA$_{weak}$,

since his believing the truth is significantly creditable to his cognitive agency (i.e., his organismic cognitive apparatus). It is the subject's organismic faculties that are first and foremost responsible for the recruitment of the extended belief-forming process (i.e., telescopic observation) on the basis of which he eventually formed a true belief with respect to the satellite's position.[8]

Now, having the resources to analyze such and similar cases in this way is important for three reasons. (1) We can explain how it is possible to acquire knowledge on the basis of artifacts while remaining fast to the guiding intuition that knowledge is belief that is true in virtue of cognitive ability (*viz.*, the ability intuition on knowledge). (2) Given that artifacts may be both hardware and software, we finally have the means and the incentive to provide a single account of knowledge that can be universally applied, despite the fact that knowledge can be attained via disparate processes whose (physical) implementation may be entirely unrelated;[9] it is by no means obvious why we should group vision, reasoning, memory, telescopic observation and so on together in the absence of a unified conception of all of them as (software or hardware) cognitive artifacts for reliably generating true beliefs. (3) Such an approach can reveal the partly social nature of many instances of individually produced knowledge. Previously, we noted that when an agent gains knowledge on the basis of an artifact, her cognitive success will be significantly creditable to her cognitive agency on account of her appropriately integrating the artifact within her cognitive character. The rest of the credit, however, should be attributed to the individuals that brought the relevant extended belief-forming processes about. Notice, however, that frequently we will not be able to attribute the rest of the credit to only one single individual, because, in most cases, in order to come up with such reliable belief-

---

[8] One may attempt to provide an alternative account of knowledge that is the product of the operation of epistemic artifacts, while also remaining fast to the ability intuition on knowledge, by claiming the following: in such cases, it is merely the agent's training and skill of using the artifact, as mirrored in the agent's neural/bodily architecture, that is the most salient factor in the causal explanation of the agent's cognitive success. This alternative account, however, wouldn't work. To see why, remember first that according to the ability intuition on knowledge and virtue reliabilism, in cases of knowledge, one's *belief must be true in virtue of cognitive ability.* Clearly, however, in cases of telescopic observations for example, the most important factor that explains the truth status of the agent's belief is the external component. To illustrate this, consider, on one hand, an untrained agent in possession of a properly working artifact. In that case, it is obvious that even though the agent will initially be unable to form any (true or false) beliefs, eventually — provided that he gains sufficient experience—not only will he form beliefs, but he will also reliably enjoy cognitive success. On the other hand, think about a well-trained agent, but in possession of a faulty artifact. In this case, despite the agent's excellent internal skills, it is evident that he would be unable to reach any (non-lucky) true beliefs, no matter how much he tried. It therefore seems that in such cases the most (and maybe the only) significant factor that explains the truth-status of the agent's belief is the epistemic artifact. In other words, since the agent's belief is true in virtue of the artifact, the virtue reliabilist must account for it being part of his cognitive system. For a more detailed treatment of this possible objection as well as why the extended cognition hypothesis is necessary for virtue reliabilism in order to account for advanced cases of knowledge, whereby the operation of epistemic artifacts is involved, see Palermos (2011).

[9] In other words, given that knowledge can be attained in such fundamentally different ways, it has so far been unclear why or whether we should assume that a single account of knowledge could apply to all of them.

forming processes the individual employs similar belief- forming processes or relies on knowledge that has been delivered by other individuals on the basis of further reliable belief-forming processes, and so on. Accordingly, in such cases the rest of the credit for the agent's cognitive success will have to be attributed to the individuals — and, in general, to the social structure — that brought the relevant reliable belief-forming process about. In other words, by combining COGA$_{weak}$ to the extended cognition hypothesis, we gain a view of knowledge whereby the individual agent can be an advanced epistemic agent only within a given social structure necessary for supplying him with the reliable-belief forming processes that he will later integrate within his cognitive character so as to come to know the truth of some proposition.

The second interesting ramification we get by combining virtue reliabilism with active externalism — and, in particular, with the hypothesis of distributed cognition — is that we can account for *epistemic group agents*. These are groups of individuals who exist and gain knowledge in virtue of a shared common cognitive character that primarily consists of a distributed cognitive ability. Such a collective cognitive ability emerges out of the members' mutual (socio-epistemic) interactions and is not reducible to the cognitive abilities possessed by the individual members, thereby allowing us to speak of a group agent in itself. This is important, because by being able to so conceptualize a group of people as a self-standing agent, we can use an *individualistic* condition on knowledge to account for knowledge that is collectively produced and which is, thereby, distinctively social.

For example, we can use COGA$_{weak}$ to explain how a research team gains knowledge on the basis of an experiment. Even though the knowledge-conducive belief-forming process consists of several experts *and* their experimental devices engaging in mutual (socio-epistemic) interactions, the *collective cognitive success* of believing the truth of some (scientific) proposition will still be significantly creditable to the group's cognitive agency (i.e., the set of the organismic cognitive faculties of its individual members). It is the set of these organismic cognitive faculties that is first and foremost responsible for the emergence and appropriate employment of the collective's belief-forming process. Crucially, however, given that this kind of cognitive success will only be attributed to the *set* of the member's cognitive agencies as a whole and to none of the individual members alone, such collectively produced beliefs won't be known by any individual alone, but by the group agent as a whole. In other words, by combining an individualistic condition on knowledge, such as COGA$_{weak}$, with the hypothesis of distributed cognition, we can make sense of the claim that a proposition is known by a group agent even though it is not known by any individual alone.

The upshot of the foregoing is a conception of knowledge that we call *extended knowledge.* Rather than understanding knowledge as an individualistic phenomenon (where the cognitive individual is confined within the cranium of the subject), we instead get a conception of knowledge which can be *extended* in terms of the subject's interactions with epistemic artifacts and which can also be *distributed* in terms of her social epistemic interactions. In short, the combination of active

externalism with a virtue reliabilistic epistemology as embodied in COGA$_{weak}$ allows us to claim that we can have extended and distributed cognitive systems that generate knowledge in the same way that our individual/organismic cognitive apparatus does.

Extended Knowledge and Social Epistemology

Contemporary epistemology, therefore, has the resources to address a variety of disparate intuitions about knowledge without departing radically from the analytic tradition. At the same time, however, there is no need to suppose that all social epistemology must remain dogmatically attached to individualism in order to qualify as 'real epistemology.' To elaborate a bit further on our suggested methodology, we would finally like to compare the extended knowledge approach with some of the central aspects of both Goldman's and Fuller's understanding of social epistemology, starting with the possible connections of the field with other academic disciplines.

As noted before, Fuller suggests that epistemology should be radically transformed on the basis of an interdisciplinary approach to the question of knowledge. Such an approach should be sensitive to any subject area that is or may become relevant, and, in particular, to the field of psychology, broadly construed. It should by now be obvious that we are sympathetic to such calls for interdisciplinary research. The extended knowledge approach to socializing epistemology is an inherently multidimensional research programme that relies for the most part on the combination of epistemology with the general field of cognitive science. Accordingly, the way we see it, any adequate theory of knowledge should be informed by and, in turn, inform back such disciplines as philosophy of mind and cognitive science, neuroscience, biology, sociology, social psychology, anthropology, or even robotics and artificial intelligence.

The machinery of the human mind (which may be extended as well as distributed) has only recently started being explored and understood. We therefore suggest that contemporary epistemology should grasp any available opportunity to tackle its subject matters from all these perspectives that are now on offer for the first time, as well as by bringing any relevant epistemological considerations to the forefront of such emerging disciplines. Such an opportunity, however, would hardly count as a transformation or revision of the field. If Alston's point about the intimate connection between epistemology and cognitive psychology is correct, then the future impact of cognitive science (and of related disciplines) on epistemology should amount to no more than the normal growth of the field as pursued within the tradition of naturalist philosophy.

It would therefore be a mistake to resist such a multidimensional approach to epistemology that promises a rounded understanding of our complex intellectual nature, a nature that even though it is shaped by our organismic abilities is not limited to them. Relatedly, think about Fuller's (and possibly Popper's) objective of formulating an account of knowledge that simultaneously reveals our natural limitations *and* our attempts to *artificially* rise above them. We believe this is an

insightful suggestion. For this suggestion is not merely in line with the extended knowledge research programme but in tune with its very hard core.

Such an approach to understanding knowledge, we believe, can generate several new and exciting avenues for research, some of which have for a long time been inaccessible. Consider, for example, the intersection between epistemology and the field of history and philosophy of science. Notably, these two intimately related fields have so far been at odds — an awkward situation owing to the fact that the former discipline has traditionally being individualistic whereas the latter has for the most part been socially oriented (hardly anyone could deny the social nature of the scientific process, especially after the publication of Thomas Kuhn's *The Structure of the Scientific Revolutions,* in 1962). The present account, however, could now provide a useful link between the two fields. Science is primarily performed by individual scientists employing their hardware and software epistemic artifacts or by research *teams* operating within scientific labs that are uniquely tailored to fit their purposes. Accordingly, the concepts of extended epistemic agents and epistemic group agents could become very handy for a mainstream epistemological analysis of the scientific progress. Indicatively, discussing the scientific revolution of the 16th century, Ronald Giere and Barton Moffat write:

> "No 'new man' suddenly emerged sometime in the sixteenth century [...] The idea that a more rational mind [...] emerged from darkness and chaos is too complicated a hypothesis" [Latour 1986,1]. We agree completely. Appeals to cognitive architecture and capacities now studied in cognitive sciences are meant to explain how humans with normal human cognitive capacities manage to do modern science. One way, we suggest, is by constructing distributed cognitive systems that can be operated by humans possessing only the limited cognitive capacities they in fact possess. (Giere & Moffat 2003, 308)

The extended knowledge project, therefore, clearly has the resources to provide an account of knowledge that places at its centre our historical inclination to artificially transcend our organismic epistemic boundaries, by establishing appropriate patterns of causal connections with bio-external structures (either in the form of artifacts or social institutions).

So, finally, let us move on to the more specific methodological considerations concerning the field of social epistemology, by focusing on Goldman first. The present approach seems to fall under both preservationist and expansionist social epistemology. It is expansionist as it is able to accommodate even the most progressive topics, such as the study of doxastic or epistemic group agents. But it is also preservationist as it retains most of analytic epistemology's core assumptions. Nevertheless, allowing for the possibility of epistemic group agents, the extended knowledge approach is not committed to Goldman's core individualism. We should also note, however, that this does not mean that our suggested approach runs counter to the spirit of traditional individualistic epistemology. On the contrary, it is

continuous with it in the sense that it draws from the arsenal of traditional (individualistic) epistemology in order to provide an understanding of knowledge that will apply equally well to both individual and collective epistemic agents. Consequently, we may say that even though it is not committed to ontological individualism, extended knowledge is a *methodologically* individualistic approach to the question of knowledge.

Does this mean that extended knowledge runs against Fuller's methodological remarks? Think about the possibility of epistemic group agents who generate knowledge in a way that cannot be reduced to the sum of the cognitive abilities possessed by their individual members. Obviously, such a possibility will be welcomed by any socially inclined theorist who wishes to give priority to the 'outside in'. Gaining individual knowledge (through, say, testimony or by reading a scientific journal) of a collectively produced true belief is a good case of downwards causation within epistemology, whereby some non-reducible social entity affects the epistemic status of an individual agent. Nevertheless, by remaining fast to the spirit and resources of individualistic epistemology in the way explained just above, and recognizing that group entities depend on — but are not entirely determined by — the cognitive abilities of their individual members, this approach allows for both upwards and downwards epistemic causation, in such a way that it can accommodate both the 'inside out' and 'outside in' approaches to social epistemology. Accordingly, even though we do not need to stick with Goldman's core individualism, there is no need to succumb to Fuller's call for an exclusive focus on the 'outside in', either. The middle grounds of an individualistic social epistemology are open to us.[10]

Contact details: S.O.Palermos@ed.ac.uk; duncan.pritchard@ed.ac.uk

References

Adams, Frederick and Kenneth Aizawa. 2008. *The bounds of cognition.* Oxford: Blackwell.

Alston, Walter. 2005. *Beyond "justification": Dimensions of epistemic evaluation.* Ithaca, NJ: Cornell University Press.

Brooks, Rodney. 1991a. "Intelligence without representation." *Artificial Intelligence* 47: 139-59.

Brooks, Rodney. 1991b. "Intelligence without reason." *Proceedings of 12th International Joint Conference on Artificial Intelligence,* 569–95.

Burge, Tyler. 1986. "Individualism and psychology." *Philosophical Review* 95 (1): 3-45.

Chemero, Anthony. 2009. *Radical embodied cognitive science*. Cambridge, MA: MIT.

Clark, Andy. 1997. *Being there: Putting mind, body, and world together again.* Cambridge, MA: MIT Press.

Clark, Andy. 2007. "Curing Cognitive Hiccups: A Defense of the Extended Mind." *Journal of Philosophy* 104: 163-92.

Clark, Andy. 2008. *Supersizing the mind: Embodiment, action, and cognitive extension.* Oxford: Oxford University Press.

Clark, Andy. 2010. "Memento's revenge: The extended mind, extended." In *Extended mind,* edited by Richard Menary, 43-66. Cambridge, MA: MIT.

Clark, Andy and David Chalmers. 1998. "The extended mind." *Analysis* 58 (1): 7-19.

Fuller, Steve. 2007. *The knowledge book: Key concepts in philosophy, science and culture.* Durham: Acumen.

Fuller, Steve. 2012. "Social epistemology: A quarter-century itinerary." *Social Epistemology* 26 (3-4): 267-83.

Gallagher, Shaun 2005. *How the body shapes the mind.* Oxford: Oxford University Press.

Giere, Ron and Barton Moffat. 2003. "Distributed cognition: Where the cognitive and the social merge." *Social Studies of Science* 33 (2): 301-10.

Goldman, Alvin. 1999. *Knowledge in a social world.* Oxford: Oxford University Press.

Goldman, Alvin. 2010. "Why social epistemology is real epistemology." In *Social Epistemology,* edited by Adrian Haddock, Alan Millar and Duncan Pritchard, 1-28. Oxford: Oxford University Press.

Greco, John. 1999. "Agent reliabilism." *Philosophical Perspectives* 13: 273-96.

Greco, John. 2004. "Knowledge as credit for true belief." In *Intellectual virtue: Perspectives from ethics and epistemology,* Michael DePaul and Linda Zagzebski, 111-34. Oxford: Oxford University Press.

Greco, John. 2007. "The nature of ability and the purpose of knowledge." *Philosophical Issues* 17 (1): 57-69.

Greco, John. 2010. *Achieving knowledge: A virtue-theoretic account of epistemic normativity.* Cambridge, UK: Cambridge University Press.

Hurley, Susan, and Alva Nöe. 2003. "Neural plasticity and consciousness: Reply to Block." *Trends in Cognitive Sciences* 7 (1): 342.

Hutchins, Edwin. 1995. *Cognition in the wild.* Cambridge, MA: MIT.

Kuhn, Thomas. 1962. *The Structure of Scientific Revolutions.* Chicago: University of Chicago Press.

Latour, Bruno. 1986. "Visualization and cognition: Thinking with eyes and hands." *Knowledge and Society: Studies in the Sociology of Culture Past and Present* 6: 1-40.

Menary, Richard. 2006. "Attacking the Bounds of Cognition." *Philosophical Psychology* 19 (3): 329-44.

Menary, Richard. 2007. *Cognitive integration: Mind and cognition unbound.* London: Palgrave Macmillan.

Nöe, Alva. 2003. "Causation and perception: The puzzle unravelled." *Analysis* 63 (2): 93-100.

Nöe, Alva. 2004. *Action in perception.* Cambridge, MA: MIT.

Palermos, Orestis. forthcoming. "Could reliability naturally imply safety?" *European Journal of Philosophy.*

Palermos, Orestis. 2011. "Belief-forming processes, extended." *Review of Philosophy and Psychology* 2 (4): 741-65.

Plantinga, Alvin. 1993. *Warrant and proper function.* New York: Oxford University Press.

Pritchard, Duncan. H. 2009. *What is this thing called knowledge?* London: Palgrave Macmillan.

Pritchard, Duncan. H. 2010a. "Knowledge and understanding." In *The nature and value of knowledge: Three investigations,* Duncan Pritchard, Alan Millar and Adrian Haddock, Part I, 5-89. Oxford: Oxford University Press.

Pritchard, Duncan. H. 2010b. "Cognitive ability and the extended cognition thesis." *Synthese* 175 (1): 131-51.

Pritchard, Duncan. H. 2012. "Anti-luck virtue epistemology." *Journal of Philosophy* 109 (3): 247-79.

Pritchard, Duncan. H, Alan Millar and Adrian Haddock. 2010. *The nature and value of knowledge: Three investigations.* Oxford: Oxford University Press.

Putnam, Hilary. 1975. "The meaning of 'meaning'." In *Language, mind and knowledge,* edited Keith Gunderson, 131-93. Minneapolis: University of Minnesota Press.

Rupert, Robert D. 2004. "Challenges to the hypothesis of extended cognition." *Journal of Philosophy* 101 (8): 389-428.

Rupert, Robert D. 2009. *Cognitive systems and the extended mind.* Oxford: Oxford University Press.

Sosa, Ernest. 1988. "Beyond skepticism, to the best of our knowledge." *Mind* 97: 153-88.

Sosa, Ernest. 1991. *Knowledge in perspective: Selected Essays in epistemology.* Cambridge: Cambridge University Press.

Sosa, Ernest. 2007. *A virtue epistemology: Apt belief and reflective knowledge, Volume 1.* Oxford: Clarendon Press.

Sosa, Ernest. 2009. *Reflective knowledge: Apt belief and reflective knowledge, Volume II.* Oxford: Clarendon Press.

Theiner, Georg. 2011. *Res cogitans extensa: A philosophical defense of the extended mind thesis.* Bern, Switzerland: Peter Lang GmbH, Europaischer Verlag der Wissenschaften.

Varela, Francisco, Evan T. Thompson and Eleanor Rosch. 1991. *The embodied mind: Cognitive science and human experience.* Cambridge, MA: MIT Press.

Wheeler, Michael. 2005. *Reconstructing the cognitive world: The next step.* Cambridge, MA: MIT.

Wilson, Robert A. 2000. "The mind beyond itself." *Metarepresentations: A multidisciplinary perspective,* edited by Dan Sperber, 31-52. New York University Press.

Wilson, Robert A. 2004. *Boundaries of the mind: The individual in the fragile sciences - Cognition.* New York: Cambridge University Press.