

*Philosophical Issues*  
 Extended Knowledge  
 Edited by Jesper Kallestrup and Duncan Pritchard

J. Adam Carter, Jesper Kallestrup, Orestis Palermos & Duncan Pritchard (University of Edinburgh)  
 ‘Varieties of Externalism’

ABSTRACT. Our aim is to provide a topography of the relevant philosophical terrain with regard to the possible ways in which knowledge can be conceived as extended. We begin by charting the different types of internalist and externalist proposals within epistemology, and we critically examine the different formulations of the epistemic internalism/externalism debate they lead to. Next, we turn to philosophy of mind and the effects of the different varieties of content externalism on our understanding of the epistemic internalism/externalism distinction. In light of the above dividing lines, we then focus on extended cognition. First, we examine the extent to which the view entails a commitment to either or both content and epistemic externalism. Second, we explore whether there are varieties of epistemic externalism that are better suited to accommodate the hypothesis of extended cognition. Finally, we attempt to delineate the field of collective and distributed cognition; we investigate how the topic intersects with the extended cognition proposal as well as with both content and epistemic externalism, and we comment on its potential input to social epistemology.

Sven Bernecker (UC, Irvine)  
 ‘Epistemic Agency, Instruments, and Testimony’

ABSTRACT. What distinguishes an internal epistemic process from the external instruments and information-presenting devices on which the process relies? According to the received view, an instrument is part of the cognitive process if the instrument meets certain conditions, such as unfettered accessibility, standing endorsement, cognitive integration, etc. The goal of the paper is threefold. The first goal is to show that none of the attempts to draw a sharp distinction between internal epistemic processes and external instruments is successful. The second goal is to cast doubt on attempts to distinguish between epistemic processes that rely on the testimony of others and those that rely on instruments and mechanisms. The third goal is to sketch an account of epistemic agency that is consistent with the epistemic extendness thesis.

Brit Brogaard (St. Louis)  
 ‘Extended Knowledge and Intrinsic Representation’

ABSTRACT. The paper starts out by distinguishing two closely related hypotheses about extended cognition. According to the strong hypothesis, there are no intrinsic representations in the brain. This is a version of the extended-mind view defended by Menary and Clark. On the weak hypothesis, there are intrinsic knowledge representations but some types of cognition, knowledge and memory are extrinsic and require external devices and particular environments for their subsistence. After drawing this distinction and clarifying the notion of an intrinsic representation, I defend the second weaker hypothesis and show why this sort of view supports what we might call a ‘situationist-friendly virtue epistemology’.

Shaun Gallagher (Memphis)  
 ‘Pragmatic Interventions into Extended and Enactive Conceptions of Cognition’

ABSTRACT. Clear statements of both extended and enactive conceptions of cognition can be found in John Dewey and other pragmatists. For example, Dewey argues against the idea of in-the-head or “armchair” cognition as “an event going on exclusively within the cortex [...] Hands and feet, apparatus and appliances of all kinds are as much a part of it as changes within the brain” (1916). He also argues for an understanding of perception starting “not with a sensory stimulus, but with a sensori-motor coordination [...] it is the movement which is primary, and the sensation which is secondary, the movement of body, head and eye muscles determining the quality of what is experienced” (1896). In this paper I’ll argue that we can find resources in the pragmatists, especially Dewey and Peirce, to address two ongoing debates: (1) in contrast to recent disagreements between proponents of extended vs enactive cognition, pragmatism support a more integrative view—an enactive conception of extended cognition, and (2) pragmatist views suggest ways to answer the main objections raised against extended and enactive conceptions—specifically objections focused on constitution versus causal factors, and the mark of the mental.

Axel Gelfert (Singapore)  
 ‘Scientific Evidence as Socially Extended Cognition’

ABSTRACT. In recent years, the idea has been gaining ground that our traditional conceptions of knowledge and cognition are unduly limiting, in that they privilege what goes on inside the single mind (or brain) of an individual reasoner. Instead, we are told, knowledge and cognition need to be understood as *embodied* (involving mind and body), *situated* (being dependent on the complex interplay between the individual and its environment), and *extended* (that is, continuous with, rather than separate from, the world ‘outside’). Whether these various interrelations and dependencies are ‘merely’ causal, or are in a more fundamental sense constitutive of knowledge and cognition, is as much a matter of controversy as the degree to which they pose a challenge to ‘traditional’ conceptions of cognition, knowledge and the mind. So far, the debate has mainly focused on questions in the philosophy of mind and knowledge. In the present paper, I plan to extend the idea of ‘extendedness’ to the philosophy of science, in particular to the concept of scientific evidence. Previous attempts to apply the ideas of ‘extendedness’ and ‘distributed cognition’ (Giere 2002, Nersessian et al. 2004, Magnus 2007, Brown 2009) have either focused on science at large (as a socially distributed activity with an obvious cognitive function) or on the cognitive role of scientific models, instruments, and measurement devices. (Regarding the latter class of objects and artefacts, historians and sociologists of science have amassed overwhelming evidence of their indispensability to the creation of scientific knowledge.) However, arguing that science is *in fact* too complex an activity to be pursued by any one individual is one thing, showing that scientific knowledge—as the *output* of this complex activity—is ‘socially extended’ is quite another. The former could be simply the result of contingent limitations, which we need to work together to overcome, whereas the latter cuts to the core of our understanding of scientific knowledge. A natural way that one might hope to resolve the tension between the collective nature of science and the individual’s claim to scientific knowledge would be by regarding the former as merely furnishing scientific evidence, to be interpreted by individual reasoners (and then passed on as scientific knowledge to others). Indeed, on most philosophical interpretations, evidence is just this: it is simply that which justifies an individual’s belief or boosts his or her credence. Yet, whatever the merits of such an individualistic conception of evidence in other areas of philosophical inquiry, it does not suffice to explain how it is that we can credit ourselves with scientific knowledge, or so I shall argue. This is because, from an ‘individualistic’ perspective, science (along with similarly ordered epistemic domains) would appear to be in danger of becoming a victim of its own success, since it produces far more empirical evidence than any one individual could ever hope to process. As John Hardwig (1991: 699) puts it, it would then seem ‘that there can no longer be knowledge in many scientific disciplines

because there is now too much available evidence'. (Note also that it is a moot point whether this is a 'practical' limitation on our part: If individual processing of the relevant data for a particular scientific question would require multiple—perhaps thousands of—human lifetimes, a 'merely' practical limitation quickly becomes a de facto *in-principle* limitation.) If we are to make sense of scientific knowledge—and of how it is that we, as individuals, can legitimately credit ourselves with at least some such knowledge—we need to acknowledge that scientific evidence itself depends on the existence of reliable (external) networks of human reasoners and non-human actors (e.g., technical artefacts, measurement instruments, computational devices, information filters etc.). Thus understood, science is complex and distributed not merely in its technical infrastructure and social organisation, but depends on mechanisms of extendedness also for its evidential and cognitive import.

Mikkel Gerken (Copenhagen)  
'Outsourced Cognition'

ABSTRACT. The dramatic developments in technologically enabled social cognition over the last decades call for a rethinking of many aspects of human cognition. According the hypothesis of *extended cognition*, we must revise our psychological categories by eliminating allegedly superficial distinctions between (internal) cognition and (external) processes. As an alternative to this proposal, I outline a hypothesis of *outsourced cognition* which seeks to respect distinctions that are operative in both folk psychology and empirical psychology. According to this hypothesis, outsourcing cognitive tasks to resources distinct from an individual's own cognitive resources is *itself* a cognitive process that is ascribable to the individual. I argue that this proposal is not merely conservative in preserving explanatory categories in folk psychology and the cognitive sciences but that it also contributes to a more fine-grained taxonomy of cognitive categories.

I will consider the epistemology of testimony as a case study. I will argue that important explanatory categories may be preserved by adopting the outsourced cognition hypothesis over the extended cognition hypothesis. Moreover, I will argue that the outsourced cognition hypothesis is also of *direct* explanatory value in informing the epistemology of testimony.

Sandy Goldberg (Northwestern)  
'De Re Trust'

ABSTRACT. When we take another person's word, we typically rely on them to have been reliable and responsible in their giving their word. In this paper I will argue that this phenomenon should be understood in terms of a kind of *de re* trust or reliance. After characterizing this phenomenon and arguing that it is seen in most ordinary cases of testimonial belief-formation, I will go on to argue that the phenomenon of *de re* trust has heretofore unrecognized implications for the epistemology of testimony. In particular, it is incompatible with any individualistic approach to the epistemology of testimony; and it offers strong support to the idea that testimonial chains of communication constitute epistemically extended chains of information-processing.

Frank Hofmann (Luxembourg)  
'The Great Unification: Semantics and Epistemology—All Grounded in One Structure?'

ABSTRACT. Recent developments in semantics and epistemology point toward a new theoretical possibility of unification. The common theme is externalism: to consider an individual or subject in its environment and to take the relations between the individual and the external objects as relevant to or even constitutive of certain semantic and epistemic features. *Prima facie*, the externalist perspective provides a number of reasons

for thinking that semantics and epistemology rest on a shared basic structure. What makes for meaning or content seems to make for epistemic status (epistemic justification, broadly understood, and knowledge) as well. Basically, the shared structure consists of a certain kind of correlation: one thing being correlated reliably with some other thing, in certain conditions. Meaning and content are grounded in certain correlations, and arguably epistemic justification and knowledge are so as well. If this is not merely a misleading superficial commonality, it should be possible to spell out the common structure of reliable coincidence, and to point to the *differentia specifica*. If successful, this would provide for a uniquely elegant unification of two central fields in philosophy, semantics and epistemology. And it would provide an indirect vindication of the externalist perspective. The possibility of such a ‘great unification’ is what I shall investigate into. If it does not turn out to be viable, this should also be of great importance since it should help us to understand better why we deal with two structurally distinct phenomena. I shall mostly look at teleo-semantics (Millikan, Dretske, Neander) and several versions of reliabilism, including teleo-reliabilism (Peter Graham), and virtue epistemology (Sosa, Greco, Pritchard) in order to examine the possibility of the great unification from the externalist perspective.

Christoph Kelp (Leuven)

‘Extended Cognition and Epistemology: Radically Extended Cognition or Knowledge-First Virtue Epistemology?’

ABSTRACT. The timeseeker looks at a stopped clock and forms a true belief about the time. The timekeeper, who knows that the clock is stopped, happens to observe the episode and, using his two radio clocks, confirms that the reading of the stopped clock is accurate. Had the stopped clock reading been inaccurate, the timekeeper would have alerted the timeseeker of this fact and enlightened her about the real time. This paper argues (i) that the timeseeker knows the time and (ii) that standard accounts of knowledge struggle to accommodate this intuition. Reason is provided to believe that allowing the timeseeker’s cognitive processes to extend beyond her skin (to include the timekeeper’s confirmation of the reading) serves to solve the problem this case poses for standard accounts of knowledge. However, the kind of cognitive extension needed is much more radical than the ones countenanced by champions of the extended cognition thesis in the philosophy of mind. It is shown that, even if this radical kind of cognitive extension is acceptable, a special problem remains for standard virtue theoretic accounts of knowledge. A knowledge first version of virtue epistemology not only avoids the special problem for standard virtue theoretic accounts of knowledge but also addresses the problem posed by the timekeeper case without resorting to radical cognitive extensions.

Jennifer Lackey (Northwestern)

‘Distributed Cognition and Group Knowledge’

ABSTRACT. It is often noted in the collective epistemology literature that knowledge is attributed to groups when information is distributed across its members, even if no individual member of the group knows the proposition in question. This is taken to show that groups can possess knowledge, where this is understood as over and above, or otherwise distinct from, any individual member possessing knowledge. In this sense, then, groups are said to have minds of their own. In this paper, I argue against this view, defending a reductionist account of group knowledge.

Ram Neta (UNC, Chapel Hill)

‘An Extended Defense of Access Internalism’

ABSTRACT. Access Internalism was once a popular thesis about epistemic justification, but it has now been almost completely abandoned. Externalists like Goldman, Alston, and Plantinga have levelled prominent objections to it, and even influential internalists like Feldman and Conee are eager to distinguish their own version of internalism (which they call “mentalism”) from access internalism. This paper draws on the resources provided by a particular version of the extended mind thesis to argue that the influential objections to access internalism all fail. If we accept a particular version of the extended mind thesis, we can then be access internalists with a clean conscience.

Kirk Michaelian (Bilkent)

‘External Memory: Implications for the Compatibility of Extended Cognition and Virtue/Credit Epistemologies’

ABSTRACT. As Vaesen (2011) has shown, an agent’s reliance on an external resource in the sorts of cases used to motivate the extended cognition hypothesis appears to pose a problem for virtue/credit epistemologies, since in such cases the cognitive success seems to be attributable to the external resource rather than to the agent. Drawing on empirical and theoretical work in a number of fields, this article argues that the apparent incompatibility between extended cognition and virtue/credit theories may rely on an overly simple view of the sort of agent-resource interaction at work in the relevant cases. The argument focuses on external memory, for two reasons: first, forms of external memory have figured centrally in the extended cognition literature; second, we are arguably currently experiencing an external memory explosion—a rapid development of new external memory technologies and a consequent rapid increase in our reliance on such technologies—which lends added urgency to understanding the epistemology of external memory, in particular. The main thrust of the argument is that considerations from a number of fields—including human-computer interaction, communication research, conversational remembering, memory and ageing, and metacognition—make clear that, in typical cases of reliance on external memory, the agent’s internal cognitive resources play a crucial role in achieving cognitive success (roughly, formation of a true memory belief): agents make use of sophisticated strategies to solve a number of problems related to the use of external memory, for example, selecting from available external memory resources, ensuring their availability for future use, storing information in appropriate resources, interpreting and evaluating retrieved information, integrating externally retrieved information with internally retrieved information, and resolving conflicts between information retrieved from different sources. The central role of the agent’s internal cognitive capacities here suggests that, at least when it comes to external memory, there need be no conflict between extended cognition and virtue/credit theories.

Joëlle Proust (Institut Jean-Nicod, Paris)

‘Epistemic Feelings and Epistemological Externalism’

ABSTRACT. Recent evidence indicates that a class of epistemic feelings, such as the feeling of knowing, can reliably predict the correction of (some of) one’s own cognitive performances (e.g., the probable validity of a memory retrieval, or of a perceptual discrimination). Another class, such as the feeling of being right, seems to allow agents to reliably evaluate in retrospect the probable correction of (some of) their cognitive responses. At least in some domains, feelings might thus be seen as providing thinkers with an internal, individualist, experience-based entitlement to having confidence in making, or having made, a valid cognitive decision. Against this view, it will be claimed that a local culture as well as a physical, temporally and materially structured environment are part of the very process through which epistemic feelings are generated, calibrated and used. First, each culture contextually gives priority to specific epistemic norms (e.g., consensus rather than truth). Second, it may or not offer appropriate

conversational and educational recurrent scaffolding for children and adults to form reliable epistemic feelings. Reliable self-appraisal is thus constituted in large part (rather than merely caused) by socially distributed features of cognition and by dynamic-environmental features (traditional grading vs. progress reports, SATs, etc.). Emotional forms of norm sensitivity qualify no less as an “extended mental” phenomenon than their analytic counterparts.

Robert Rupert (Colorado)  
‘Situated Self-Knowledge’

ABSTRACT. Semantic externalism is often taken to create a special problem of self-knowledge: if the contents of a subject’s thoughts are partly determined by environmental factors to which she has no immediate access, she cannot know, directly or *a priori*, the content of her own thoughts, which strikes many as implausible. Active externalism might appear to give rise to parallel problems: if the subject’s cognitive processing takes place beyond the boundary of her body, she cannot know *how* she is thinking; additionally, if the external processes in question involve external vehicles, one might worry that she does not even know *what* she is thinking, regardless of whether an internalist or externalist semantics for these vehicles is presupposed. A straightforward solution to the problems created by active externalism appeals to the distinction between personal and subpersonal processing: we shouldn’t expect human subjects to know how their subpersonal processing proceeds; thus, such external processes can belong to human subjects when not accessible to them. In the first half of the paper, I lay out these problems for active externalism and argue against the proposed fix, on two grounds: (1) it threatens incoherence to attribute the external processes in question to a subject if that subject is physically bounded, and (2) the view rests on an unsupportable version of the personal-subpersonal distinction. In the second half of the paper, I develop a positive account of what it is to be aware of one’s own cognitive processing and explore the extent to which we have such awareness in the age of the distributed self, keeping clearly in mind the wealth of empirical work cited in connection with situated views of cognition.

Sarah Sawyer (Sussex)  
‘Morals, Judgement and Reason’

ABSTRACT. This paper will concern two important debates in moral philosophy. The first (that between judgement internalists and judgement externalists) concerns the question of whether there is a conceptual connection between moral judgement and motivation. I will argue that an externalist theory of the mind demonstrates that the debate rests on a false presupposition, and hence that the disagreement can be overcome if an externalist theory of the mind is adopted. The second debate (that between reasons internalism and reasons externalism) concerns the relationship between moral reasons and an agent’s subjective motivational set. I will argue that an externalist theory of the mind can be used not only to ground an externalist theory of epistemic warrant, but also to ground an externalist theory of reasons.

Georg Theiner (Villanova)  
‘Between Linguaging and Languages: Rethinking Knowledge of Language’

ABSTRACT. The “distributed language movement” (DLM) has emerged as a new kid on the block in contemporary debates over the nature of language, and how it ought to be studied (Cowley 2007, 2009a; 2009a; 2011a). DLM subsumes a variety of closely related conceptions of language that are known as “linguaging” (Maturana 1988; Kravchenko 2006), “utterance-activity” (Cowley 2009b; Thibault 2011), “first-order language” (Love 2004), “dialogue” (Linell 2009), “colloquy” (Jennings and Thompson

2012), or “embodied, embedded language use” (Fowler 2010). Combining the framework of distributed cognition (Hutchins 1995; Hollan et al. 2000) with ideas from integrational linguistics (Harris 1981), dialogism (Linell 2009), systemic functional linguistics (Halliday 1975), ecological psychology (Gibson 1979), and embodied neuroscience (Damasio 1994), DLM conceives of language as an ecological, dialogical, sense-saturated, and non-local activity centered on achieving the inter-personal coordination of wordings (Cowley 2009a; 2011a, 2011b).

The underlying theoretical orientation of DLM is radically opposed to the classical cognitive-scientific view that language is an internalized computational system (“I-language”) in the mind/brain of individual speakers that should be studied independently of the contexts in which it is used (for classical statements of this view, cf. Chomsky 1986, 1995). However, it seems fair to say that orthodox fundamentalist versions of this internalist view no longer dominate the mainstream of linguistics and cognitive science, which has seen an influx of embodied, embedded, extended, and enactive (“4E”) approaches (for an overview, see Menary 2010). Comparing and potentially integrating DLM with various strands of 4E-approaches, and determining how much of a departure from internalism each of them really implies, is an important desideratum (see Dale 2012 and the papers discussed in that special issue).

But such a project is complicated by a number of factors. First, the vocabulary and methods that are used by DLM to propound its views are in many cases not germane to the order of business in cognitive science. Second, in order to accentuate their distinctive stance in opposition to internalism, proponents of DLM sometimes employ an excessively strong rhetoric that tends to obscure rather than to clarify the relationship to alternative criticisms of classical internalist views. Third, since internalism itself comes in a large variety of degrees and shades, it is far from clear that all the main tenets of DLM are indeed incompatible with certain non-classical versions of internalism. Consequently, it can be hard to assess the exact nature of their disagreement, the chances for a theoretical rapprochement, and what kind of evidence can be marshaled to support one or the other.

My goal in this paper is to help establish a shared conceptual framework that could be used to build bridges between what appear to be divergent conceptions of language. Rather than implausibly pitting “internalist” against “distributed” accounts of language, and thus to perpetuate what may turn out to be a false (or at least an insufficiently articulated) dilemma, I distinguish six central theses that are frequently run together in the literature on DLM (for a similar strategy in the context of embodied cognition, see Wilson 2002). Those six theses, each of which warrants an independent consideration, concern:

- (1) the psychological reality of mental representations (“representation thesis”);
- (2) the purely symbolic, amodal nature of mental representations (“modality thesis”);
- (3) the autonomy of linguistic vis-à-vis non-linguistic activities (“autonomy thesis”);
- (4) the ways in which linguistic meanings are partly constituted by the natural, social, or temporal environment in which linguistic interactions take place (“semantic externalism”);
- (5) the ways in which linguistic processes are bodily, socially, or culturally distributed (“process externalism”);
- (6) the emergent collective reality of linguistic structures (“linguistic collectivism”).

With my critical examination of these six theses, I aim to show that DLM should not be understood as a “package deal” utterly incompatible with internalism. Hopefully, this will lead theorists with an internalist bent to pay greater attention to the distributed and dynamic aspects of language use that have been the main focus of DLM. Conversely, proponents of DLM may eventually come to the realization that at least some of their

claims are quite naturally combined with certain internalist conceptions of linguistic structure and psychological constraints on language use. Perhaps a more eclectic mix between internalist and distributed approaches is exactly what is needed to understand the complex webs of linguistic scaffolding that we spin around us, but that also spin us.

Adam Toon (Bielefeld)

‘Virtuous Realism?’

ABSTRACT. Defenders of scientific realism often appeal to reliabilist theories of knowledge. One reason they do so is to argue for the possibility of gaining knowledge through the use of instruments in science. Put simply, according to reliabilism, the scientist can gain knowledge through an instrument just in case that instrument is, as a matter of fact, reliable. The scientist need not be able to offer reasons for thinking that this is the case. In particular, she need not be able to refute the arguments put forward by empiricists for doubting that instruments provide us with knowledge of unobservable entities.

Within contemporary epistemology, however, simple reliabilist theories are often thought to face serious difficulties. One such difficulty is that they fail to accommodate the intuition that knowledge is the product of a cognitive ability. To capture this intuition, virtue epistemologists argue that is not enough for knowledge that a belief-forming process is reliable; instead, it must also be appropriately integrated within the cognitive agent, such that her epistemic success may be credited to her cognitive agency. Rather than lending support to scientific realism, such accounts threaten to undermine it, by rendering the use of instruments highly problematic, since in such cases it appears that knowledge is primarily creditable to the instrument, not the scientist.

Despite its prominence within contemporary epistemology, the consequences of virtue epistemology for the philosophy of science remain little explored. Virtue epistemologists differ both in the way that they understand what it is for a process to be integrated into an agent’s cognitive character, and in the degree to which they require success to be creditable to the agent. Recently, some have further explicated the relevant notions of cognitive character and agency by looking to work on extended cognition in the philosophy of mind. By drawing on studies in the history and sociology of science, I will ask whether the various conditions proposed by virtue epistemologists are met in typical cases in which instruments are used in science. By doing so, I will assess whether virtue-theoretic approaches to knowledge support or undermine the case for scientific realism.

Krist Vaesen (Eindhoven)

‘Objectual Understanding without Knowledge, Exhibit 3: Extended Cognition’

ABSTRACT. Is objectual understanding a species of knowledge? Greco (*forthcoming*) has recently claimed that it is. According to Greco, to have understanding of a thing is just to have systematic knowledge of the thing’s dependence relations. Rather than rehearsing two earlier counterarguments—one due to Elgin (2004, 2007), one due to Kvanvig (2003), both of which Greco dismisses—I offer a new argument here. In particular, I show that cases of extended cognition present Greco with a dilemma. In such cases, an agent offloads her understanding of a complex thing onto the environment, which affords her to forget (i.e., to stop knowing) most of the facts about the thing’s complex dependence relations. In order to save his knowledge account of understanding, on one horn of the dilemma, Greco needs to deny the agent understanding, thereby losing one of the conspicuous features of having an understanding of a thing, namely that it picks out a persistent feature of agents, rather than a state which is turned on just when in the process of reasoning about the thing. On the other horn of the dilemma, Greco needs to accept that agents may know things just in virtue of information being extracranially



available to them. More generally, my paper illustrates the usefulness of discussions about extended cognition for our understanding of objectual understanding.

Hamid Vahid (Tehran)

‘Cognitive Penetration, the Downgrade Principle and Extended Cognition’

ABSTRACT. It has recently been argued that just as, say, prejudice or wishful thinking can generate ill-founded beliefs, the same is true of experiences. The idea is that the etiology of cognitively penetrated experiences can downgrade their justificatory force. This view, known as the Downgrade Principle, seems to be compatible with both internalist and externalist conceptions of epistemic justification. An assessment of the credentials of the Downgrade Principle is particularly important in view of the fact that not all cases of cognitive penetration are epistemically malignant. There are good and bad cases of cognitive penetration. While granting the intuitive plausibility of the Downgrade Principle, I consider and reject two defenses of this principle (due to Siegel and McGrath). I will also explain why both reliabilism and mentalism fail to sustain the distinction between good and bad cases of cognitive penetration. The Downgrade Principle is then discussed in the context of the extended mind thesis. Having first delineated an appropriate notion of a mental state for this context, I will try to show why it is still difficult to maintain the distinction between good and bad cases in such a context. Finally, it will be argued that an explanationist conception of epistemic justification might have a better chance of accounting for the distinction in question.

Åsa Wikforss (Stockholm)

‘Extended Cognition and Internalist Justification’

ABSTRACT. The paper discusses the relation between extended minds and extended knowers. It is argued that extended beliefs are (in principle) possible and that this implies that extended knowers are (in principle) possible, since an extended state qualifies as a belief if and only if it is a candidate for knowledge. Moreover, it is argued that this is so even if one endorses an internalist conception of justification—at least on the most plausible construal of internalism. Consequently, extended knowledge does not presuppose epistemic externalism, nor does it presuppose virtue epistemology, as is commonly assumed. This provides a further illustration of the need to keep the various ‘externalisms’ distinct: Just as vehicle externalism is compatible with content internalism, so is it also compatible with justification internalism.