Philosophical Issues
A Supplement to NOÛS

Philosophical Issues, 24, Extended Knowledge, 2014
doi: 10.1111/phis.12025

A PARTIAL DEFENSE OF EXTENDED KNOWLEDGE

Berit Brogaard
University of Miami

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 Andy Clark and Richard Menary. On the weak hypothesis, there are intrinsic representations in the brain but some types of cognition, knowledge or memory are constituted by particular types of external devices or environmental factors that extend beyond the skull and perhaps beyond the skin. This type of view was defended, for example, by Andy Clark and David Chalmers. After drawing this distinction and clarifying the notions of causal influence and constitution, I defend the second weaker hypothesis with respect to procedural knowledge and knowledge of action and show why this sort of view supports what we might call a 'situationist-friendly virtue epistemology'.

1. Introduction

Extended cognition is the idea that mental states, or representations, depend in an important sense not just on the brain but also on the body or the external environment. I shall here use 'mental state' and 'mental representation' interchangeably without assuming that a mental state or a mental representation is in the skull. The idea of extended knowledge can be seen as a special instance of extended cognition. If knowledge is extended, it depends in an interesting way on entities that extend beyond the skull and perhaps the skin. The main question that shall concern us here is that of whether there are any cases of extended knowledge.

Extended cognition should not be confused with content externalism. One could hold that no mental states reach out of the skull but still hold that the contents of mental states are objects and property instances in the

environment. So, content externalism is not a special version of extended cognition. That said, there is indeed an interesting connection between extended cognition and content externalism. We will explore that connection further in section 5.

When defenders of extended cognition say that a particular type of cognition is extended, they typically don't mean that a mental state in the brain causally depends on certain entities in the external environment (see e.g., Clark and Chalmers 1998; Adam and Aizawa 2001, 2010; Menary 2006, 2007; Clark 2007, 2008; Aizawa 2008; Theiner 2011). The hypothesis typically is a stronger one to the effect that entities in the external environment constitute particular types of mental state. 'Constitution' can be understood either in terms of essence or in terms of realization. On the first view, 'constitutes' should not be taken to be synonymous with 'is a part of'. There is a sense in which the bricks in my house partially constitute my house. They are part of my house right now in the actual world. But on a standard philosophical view of identity, I can replace a brick in my house without destroying the identity of my house. So, things can be parts of other things without this being a necessary fact. When defenders of extended cognition talk about a mental state being extended, they typically mean that it is *necessary* that certain external entities partially constitute the mental state (Noë, 2004: 221; Menary, 2006; Rowlands, 2009; Aizawa 2010). So, we can articulate the first notion of constitution as follows:

Essential Constitution

One entity x is (partially) constitutive of another entity y $=_{\text{def}}$ Necessarily, x is a part of y.

Representational and role-functionalist views of mental states construe the constituents of mental states in terms of essential constitution. Consider first representational views of mental states. Representational views take mental states to represent either concepts or objects and properties. For example, my belief that Socrates is human may be taken to be a state that represents Socrates and the property of being human. Even if that belief is a state of my brain, however, I can have this belief for many years while my brain undergoes various changes in structure and function. So, my belief is not identical to any particular set of neurons or any particular neural activity. My belief is a *state* that carries information, not a concrete particular in my brain. Its constituents are entities my belief has as a matter of necessity. For example, the property of representing Socrates is constitutive of my belief, and by that we mean that it's a necessary constituent of my belief, not that it happens to be part of my belief. The particular pattern of neuron activity or the neural network that instantiate the belief, on the other hand, is not constitutive of my belief.

Consider next role functionalism about mental states. Functionalism about mental states is the view that what makes something a mental state of a particular type depends on the way it functions in a cognitive or behavioral system. For example, pain may be taken to be any state which produces the belief that something is wrong with the body and a desire for the state to disappear and which tends to give rise to anxiety, moaning, whining or screaming. If this functional role is played by C-fiber stimulation, then a person is in pain when she undergoes C-fiber stimulation. Role functionalists, however, would identify the pain property with the functional role property itself, not the C-fiber stimulation. So, C-fiber stimulation is not a constituent of the pain state but the property of having a desire for the state to disappear is a constituent of the pain property in the envisaged scenario, and it's a constituent in the sense of being is a necessary part of the pain property.

Realizer functionalists are bound to treat constitution differently. Whereas role functionalists would identify pain with the functional role property itself, realizer functionalists would identify pain with C-fiber stimulation in the envisaged scenario. C-fiber stimulation, however, is not a necessary part of pain. If the pain role is played first by A-delta fiber stimulation and then by C-fiber stimulation, then the realizer functionalists would identify the pain state with both. So, this leads to a second definition of 'constitution'.

Realizer Constitution

One entity x is (partially) constitutive of another entity y that plays functional role F $=_{def}$ x is a part of y in all scenarios in which x (together with other entities) plays the F role.

We can define 'extended cognition' in the constitutional sense as follows, allowing for either interpretation of 'constitution':

Extended Cognition

A mental state M is extended $=_{def}$ M is partially constituted by entities that reach beyond the skull.

Extended knowledge is a special case of extended cognition:

Extended Knowledge

A knowledge state K is extended $=_{def}$ K is partially constituted by entities that reach beyond the skull.

Before we can address the question of whether there are any cases of extended knowledge, we need to draw some distinctions between different types of knowledge. Following an analogous distinction in psychology for memory (see e.g. Squire 2004), we can distinguish between declarative knowledge

and procedural knowledge. Declarative knowledge includes knowledge of personally experienced events, also known as 'episodic knowledge' (e.g., the knowledge that I talked to my friend Sonya last Thursday), and knowledge of impersonal facts, also known as 'semantic knowledge' (e.g. the knowledge that Harry S. Truman was the 33rd President of the United States). Procedural knowledge is knowledge of how to perform certain types of action, for example, the knowledge that enables you to swim. This distinction corresponds, roughly, to the standard philosophical distinction between knowledge-how and propositional knowledge (Stanley and Williamson 2001; Bengson and Moffett 2007; Brogaard 2009). However, the philosophical distinction is not as useful as the psychological distinction, as intellectualists about knowledge-how argue that knowledge-how is a form of propositional knowledge. Furthermore, there appear to be cases of knowledge-how that are not cases of procedural knowledge (e.g., knowing how to get to New York without having the ability to get there). I shall mostly employ the terminology from psychology here.

Another distinction that will be of relevance to the question of whether there are any cases of extended knowledge is that between belief-views and knowledge-first views. According to the standard belief views, only belief states can be knowledge states. Belief states are knowledge states when they satisfy certain state-external constraints (e.g., being justified or reliably formed). On the knowledge-first view, knowledge is a determinable, and a range of different cognitive states, including perception, memory and belief, are determinates of the determinable when they satisfy certain further state-external constraints (Williamson 2000: Intro, Brogaard 2011a).

Some thinkers have argued that all forms of cognition and hence all forms of knowledge are extended in the strong sense (Menary 2006, 2007; Clark 2005, 2007, 2008). Accordingly, there are no intrinsic knowledge representations in the brain, which is to say, there are no knowledge representations that are not at least partially constituted by external entities (including bodily abilities, external tools and conventions). A weaker hypothesis is that there are intrinsic knowledge representations in the brain but that some types of cognition are constituted by particular types of external devices or environmental factors that extend beyond the skull and perhaps beyond the skin (e.g., Clark and Chalmers 1998). Yet others have argued that all cognitive representations are intrinsic and hence that there are no forms of extended cognition (Adam and Aizawa 2001, 2010; Aizawa 2008).

Here I will provide arguments for a version of the weak hypothesis: the only plausible types of extended knowledge in the constitutional sense are knowledge of action and procedural knowledge. Along the way I will offer arguments against Noë's hypothesis that sensorimotor know-how is constitutive of perception and that informational states associated with external memory storage devices, such as notebooks and iPhones, can be constitutive of belief. I will conclude by showing that the hypothesis that many mental

states causally depend on entities beyond the skull or skin but are not constituted by these entities has interesting implications that suffice to make a weak extended-cognition hypothesis interestingly different from traditional models of cognition. Finally, I will offer a response to the argument that the view that many knowledge states causally depend on external devices is at odds with virtue epistemology.

2. Extended Perception

On the knowledge-first view, perceptual states that satisfy some further state-external constraints count as knowledge states (Brogaard 2011a). Given this view of knowledge, *if* perceptual states are extended (in the constitutional sense) and some of these states satisfy the required additional state-external constraints, then it follows that at least some knowledge states are extended. I will argue, however, that perceptual states are not in general extended.

In *Action in Perception* Alva Noë argues that perception is not a process in the brain, but a kind of skillful activity of the body as a whole. It's not entirely clear what this skillful activity consists in but it seems clear that whatever skillful activity amounts to, it is supposed to be constitutive of perception. As Noë lays out his position it appears to be a version of role functionalism, according to which the fundamental role of perceptual activity is to skillfully explore the environment, and that the property of perceiving is this higher-order property of skillfully exploring the environment. This is a common position among defenders of embodied, or enactive, perception.

Defenders of embodied perception tend to offer a range of empirical results showing how perception depends on skillful activity. Typical examples include (see e.g. Noë 2009):

> *One-Eyed Kittens:* Kittens who have one eye sewn shot until adulthood do not develop normal vision because they don't receive the proper kind of sensory input to the visual system.
>
> *Passive Kittens:* When kittens are moving around in their environments pulling carriages with other kittens in them, the kittens in the carriages see everything that the kittens pulling them see, but because the passive kittens don't actually use skillful action, their visual system doesn't develop properly.
>
> *Sensory Substitution:* A camera is mounted to a blind person or a blindfolded person's head and connected to vibrators against the skin to allow visual information presented to the camera to produce tactile stimulation on the skin. After a few hours people are able to make visual discriminations that have a visual character to them. So,

> the function of sensory neurons depends on how we interact with the world through skillful exploration.

The examples are supposed to show that entities outside the skull partially constitute perceptual states. However, at best they show that external entities causally influence perceptual states. Take the examples One-Eyed Kittens and Passive Kittens. In both cases an abnormal environment during the crucial window of opportunity within which experience shapes the brain's sensory and cognitive channels affects those sensory or cognitive mechanisms later in life. It is clear why these cases do not support a constitution thesis. Defenders of embodied perception would not want to say that *current* perceptual states literally extend into the past (or that current states can play the perception role only if some past states obtained). So, at best the kitten cases show that normal vision causally depends on certain active exploration and visual stimulation during early development, not that skillful action is partially constitutive of perceptual experience. There is, of course, room for arguing that the kittens cannot see properly because they did not develop the embodied skills required for proper vision and hence that the kittens' visual states aren't genuine perceptual states. I return to the question of whether embodied skills are constitutive of perception below.

The sensory substitution case *does* show that the function of the somatosensory cortex depends on how the perceiver explores the world. But this observation, too, is insufficient for demonstrating that a neural state can play the perception role only by being shaped by skillful action and hence it is insufficient for demonstrating that anything outside the skull is constitutive of perceptual experience. All that is needed to explain this case is the assumption that the function of the somatosensory cortex causally depends on how the perceiver explores the world. That is, we can grant that we can perceive objects in a visual-like way on the basis of tactile stimulation only by moving around in the environment. But this doesn't show that the motion is *constitutive* of perceptual experience as opposed to a condition that causally influences the functional activity of the underlying neural tissue.

There are reasons to doubt that skillful action could be partially constitutive of all perceptual states. One reason is that there are perceptual states that do not involve skillful action. For example, states of hallucination and illusion need not involve any action at all. The same goes for states of vivid dreaming and visual imagery, states that are very similar phenomenally and neurologically to visual experience (Ganis, et al. 2004). It may be argued that these are not genuine perceptual states. However, it is not clear how a non-question begging argument would go for restricting genuine perceptual states to veridical ones. Even if such an argument could be provided, it is unlikely to be sufficient to show that perceptual states always involve skillful activity. For there presumably could also be very simple veridical perceptions that involve no skillful activity on the part of the perceiver. Immobilizing a

perceiver is unlikely to prevent her from perceiving a uniformly colored wall or hearing a ringtone.

On a closer reading of Noë (2004), however, it doesn't seem that he ultimately wants to claim that it really is skillful activity that is constitutive of perception. What he seems to commit himself to is that perception is partially constituted by the exercise of sensorimotor know-how. Though this proposal is initially more promising than one that requires actual skillful activity, there are at least two considerations against thinking that sensorimotor capacities are constitutive of perception in general.

First, while it is true that sensorimotor capacities are not (normally) merely in the head, as they also involve certain muscle development and unsevered neural connections between the brain and the muscles, there is no doubt also a brain component to these capacities. So, if sensorimotor capacities were constitutive of perceptual experience, we should expect brain imaging studies and single neuron recordings to reveal the brain component of these capacities. But brain imagining studies and single neuron recordings do not reveal any such component for the general case (see e.g. Heeger 1999 and Ganis et al. 2004 for a range of neural correlates underlying visual perception). To be sure, there are imaging studies showing that when people read action-related sentences or retrieve memories of actions, there is activation in motor areas (Calvo-Merino et al. 2005; see also Rissman and Wagner 2012 for a review). But these findings do not support Noë's hypothesis. He is not simply saying that motor capacities are constitutive of perceptions of action. He is making a claim about all cases of perception, including perception of a colored background, perception of a ringtone, and so on.

Second, on Noë's (2004) enactive view, perceptual experience is the exercise of sensorimotor know-how. There are many places in the (2004) book where Noë appears to take sensorimotor know-how to be knowledge of how sensory inputs vary as you move (e.g., p. 78). But knowing how sensory inputs vary as you move is not true know-how. It is not true procedural knowledge. It's rather in the same category as knowing how warm the seminar room is when every seat is taken or how President Obama responded to the agreement to curtail greenhouse gas emissions at the climate meeting in Copenhagen as the meeting approached the end. Despite the deceptive linguistic appearance, this sort of language does not reflect knowledge-how but knowledge-that. You know how Obama responded if you know that he told the world leaders 'not to talk, but to act'. Knowledge of how sensory inputs vary as you move thus is more akin to knowledge of what a tilted coin looks like when it's not tilted. No one would disagree that perception is partially constituted by information about viewpoint-dependent properties and their relation to size, shape and color constancies. It may be objected that although we know what to expect when we see a coin from several angles, it would be nearly impossible to put these expectations down in words. Sensorimotor know-how thus seems to be implicit rather than explicit and,

it may be argued, is therefore not a candidate for propositional knowledge. However, this doesn't follow. There is a lot of knowledge that we cannot put into words but which nonetheless is a form of propositional knowledge. A good example is the information processed in the fusiform gyrus in the brain's ventral stream in response to faces. I doubt anyone could put the information that underlies face recognition into words but that doesn't make face recognition (or knowing a face) a form of procedural knowledge or know-how.

Elsewhere Noë says that 'the basis of perception, on our enactive, sensorimotor approach, is implicit practical knowledge of the ways movement gives rise to changes in stimulation' (2004: 8). This way of formulating the enactive view seems to implicate the parts of the dorsal stream potentiated for online-action representation as a constituent, or supervience basis, of perceptual experience. But as David Milner and Melvyn Goodale have argued, vision for action and vision for perception are two distinct cortical streams that are functionally and anatomically separated in the brain (Goodale and Milner 1992; Goodale, Milner, Jakobson, and Carey 1991; Milner and Goodale 1995, 2008; Brogaard 2012). The ventral stream processes information for object recognition, whereas the dorsal stream computes information for those aspects of action that take place rapidly and without conscious deliberation. If Milner and Goodale's dissociation hypothesis is correct, then the exercise of sensorimotor know-how is not a constituent of perceptual experience, unlike what Noë appears to say (Brogaard 2011c). Although there are important interactions between the two streams, the dorsal streams does not interact with the ventral stream in every case of vision for perception, as Noë's theory would require.

Although Noë's general claim that sensorimotor capacities or sensorimotor know-how are constitutive of perception appears unmotivated, it may well be that certain forms of perception and understanding of action described linguistically or presented sensorily are partially constituted by a form of sensorimotor know-how. For example, in one study it was shown that attempts to memorize actions in order to imitate them activate dorsal stream processes, whereas attempts to memorize actions in order to recognize them on a later occasion activate ventral stream processing. In the study, brain scans were performed on individuals presented with videos of meaningful and meaningless actions and requested to observe the stimuli either to imitate the action or to recognize it (Decety et al. 1997). The findings indicated that observation of action in order to imitate the action was specifically associated with bilateral activation of the dorsal pathways, reaching the premotor cortex. The ventral pathway was activated when the task was to observe in order to recognize the action after the scan. In another study, PET scans were performed in two sessions using meaningless and meaningful actions (Grezes, Costes, and Decety 1998). In the first session, subjects were asked to look at videos without any specific aim. In the second, they were asked to look at the videos with the aim of imitating the actions presented. In the

first session it was found that the pattern of brain activation was dependent on the nature of the movements presented. In the second session, the task representation initiated information processing in the dorsal "action" pathway.

However, despite the evidence that sensorimotor areas may become activated in perception of action tasks, it would be premature to draw any conclusions from these empirical studies about whether perception of action extends beyond the skull. As I will argue below, genuine procedural knowledge does indeed extend outside the skull. But the claim that knowledge of action is extended in the same way requires further argument. I believe an argument can be made for this claim. As we will see below, the observations that sensorimotor areas of the brain undergo radical change when the brain is forced to control a new and differently proportioned body make it plausible that some forms of (declarative) knowledge of action are partially realized in extracranial entities. I will return to that point below.

3. Extended Belief and Memory

On both the belief-view and the knowledge-first view, belief that satisfies certain state-external constraints (e.g., being justified or reliably acquired) counts as knowledge. So, another way to show that there is extended knowledge is to show that there are cases of extended belief, and that some of those cases are also cases of knowledge, and hence cases of extended knowledge. Andy Clark and David Chalmers have offered a well-known argument in favor of the claim that beliefs can be constituted partly by features of the external environment. The central case they rely on in their argument runs as follows:

> First, consider a normal case of belief embedded in memory. Inga hears from a friend that there is an exhibition at the Museum of Modern Art, and decides to go see it. She thinks for a moment and recalls that the museum is on 53rd Street, so she walks to 53rd Street and goes into the museum. It seems clear that Inga believes that the museum is on 53rd Street, and that she believed this even before she consulted her memory. It was not previously an *occurrent* belief, but then neither are most of our beliefs. The belief was sitting somewhere in memory, waiting to be accessed.
>
> Now consider Otto. Otto suffers from Alzheimer's disease, and like many Alzheimer's patients, he relies on information in the environment to help structure his life. Otto carries a notebook around with him everywhere he goes. When he learns new information, he writes it down. When he needs some old information, he looks it up. For Otto, his notebook plays the role usually played by a biological memory. Today, Otto hears about the exhibition at the Museum of Modern Art, and decides to go see it. He consults the notebook, which says that the museum is on 53rd Street, so he walks to 53rd Street and goes into the museum.

> Clearly, Otto walked to 53rd Street because he wanted to go to the museum and he believed the museum was on 53rd Street. And just as Inga had her belief even before she consulted her memory, it seems reasonable to say that Otto believed the museum was on 53rd Street even before consulting his notebook. (Clark and Chalmers 1998: 13)

Though Clark and Chalmers' position appears to be a form of realizer functionalism and they speak as if Otto's notebook is constitutive of Otto's belief, I shall assume that they mean that an informational state associated with Otto's notebook is constitutive of Otto's belief. If we were to tear off a small piece of one of the notebook's blank pages, this would not make a difference to Otto's belief. This is because it's the informational state of the notebook that is said to be constitutive of the belief, not the concrete particular notebook. This is not to say that we could just copy over all the information in the notebook without this giving rise to changes in Otto's beliefs, as it's plausible that, say, discoloring and smudge on the pages after long-term use is somehow assisting Otto in accessing the information. At least we cannot rule this out. But changes to the notebook itself that do not affect the information it contains should not have any effect on Otto's beliefs. So, the hypothesis being set forth here is that an informational state of the notebook is constitutive of Otto's belief, not that a concrete particular serves as a constituent.[1] Like Clark and Chalmers, I shall often avoid the somewhat tedious talk of an informational state of an object and simply talk about the object itself. But this should be taken as shorthand for the more tedious formulations.

Having made this clarification, we are now in a position to explore why Clark and Chalmers think that Otto has an extended belief about the location of the Museum of Modern Art. They say that the belief is extended because once we assimilate Otto's notebook to Inga's long-term memory, we will see that Inga and Otto's cases are explanatorily similar. Clark and Chalmers put it as follows:

> In relevant respects the cases are entirely analogous: the notebook plays for Otto the same role that memory plays for Inga. The information in the notebook functions just like the information constituting an ordinary non-occurrent belief; it just happens that this information lies beyond the skin . . . . Certainly, insofar as beliefs and desires are characterized by their explanatory roles, Otto's and Inga's cases seem to be on a par: the essential causal dynamics of the two cases mirror each other precisely. (Clark and Chalmers 1998: 13)

Otto's and Inga's cases do seem to be explanatorily on a par. The notebook and Inga's long-term memory appear to play the same roles in the causal explanations of Inga and Otto's behavior. Otto wants to see an exhibition at the Museum of Modern Art, intends to go to the museum and believes it's on 53rd Street. This explains Otto's walking to 53rd Street.

However, on further scrutiny, it's not Otto's and Inga's standing beliefs (assuming they both have them) that play a role in the causal explanation of their behavior, it's their occurrent beliefs about the location of the Museum of Modern Art. Unretrieved memories cannot play the required explanatory role in intentional action. But as the story is told, Inga's and Otto's occurrent beliefs are not constituted by notebooks or any other entities beyond the skin. So, they are not extended.

However, there is an easy route from the conclusion that a standing belief is extended to the conclusion that the corresponding occurrent belief is also extended. It certainly seems plausible that occurrent beliefs that consist in retrieving information from long-term memory are extended just when the long-term memory is extended. It may even be that continuing memory retrieval is needed to keep a belief in working memory for the duration of time it takes to walk to 53rd Street. While Inga keeps the information in mind by recruiting semantic areas of the brain via the hippocampus, Otto keeps the information in mind by continuing to consult the notebook. On this variation on the story, Otto's occurrent belief, which is an explanatory component of the causal explanation of his behavior, crucially involves the notebook and hence is extended.

However, I don't think Clark and Chalmers' argument works, even if it is modified in this way. To see why it doesn't, it will help to get clear on what Clark and Chalmers' argumentative strategy is. Their reference to the explanatory roles of belief and desire in action may make it seem that they draw their conclusion that Inga and Otto believe the same thing on the basis of a similarity in their overt behavior. However, we cannot make inferences about what people believe merely on the basis of their overt behavior. Consider the following case. Olga hears from a friend that there is an exhibition at the Museum of Modern Art, and decides to go see it. She thinks for a moment and recalls that the museum is either on 43rd Street or on 53rd Street. As she lives uptown, she has to pass 53rd Street to get to 43rd Street, so she walks to 53rd Street, spots the museum and enters. Here it seems clear that Olga believes that the museum is either on 43rd Street and on 53rd Street. If, however, we make a judgment about whether Olga and Inga believe the same thing on the basis of their overt behavior, we are going to wrongly infer that they both believe that the museum is located on 53rd Street, as there is no difference in their overt behavior. So, old-fashioned behavioral functionalism fails.

However, Clark and Chalmers' argumentative strategy is not simply to make conclusions about whether Inga and Otto believe the same thing on the basis of their overt behavior. Nor is it to draw conclusions on the basis of memory processes while bracketing overt behavior. Rather, they adhere to a form of psychofunctionalism, comparing both overt behavior and the memory process: Both Inga and Otto hear that there is an exhibition at the Museum for Modern Art and decide to go see it. Otto accesses information about the museum's location in his notebook, whereas Inga accesses that

information in a semantic memory area of the brain. They walk to 53rd Street and enter the museum. So, Clark and Chalmers conclude that both Inga and Otto had a standing belief about the location of the museum; the difference is simply that Inga's information is stored in the brain, whereas Otto's information is stored in the notebook.

However, comparing Inga and Otto's overt behavior and memory processes does not establish that they both have a standing belief about the location of the museum. Clark and Chalmers only focus on one of Inga's and Otto's behaviors, namely that of walking from home to 53rd Street. This, however, is a very coarse-grained description of their actions. Their actions while they are on their way to 53rd Street are very different. If we were to ask why Otto but not Inga constantly stops on the way to 53rd Street and looks in a notebook, the explanation would be that while both Inga and Otto desire and intend to go to the Museum of Modern Art, only Inga remembers where it is located. While Inga may need to consult her biological long-term memory, she doesn't need to stop what she is doing to consciously consult a notebook. Otto merely remembers where to look to find the information about the location. So, unlike Inga, Otto does need to stop what he is doing to consciously consult a notebook. The assumption that Otto doesn't remember the museum's location is the best explanation of the overall differences in their behavior. It explains why only Otto consciously consults a notebook before leaving the house and why he keeps stopping what he is doing to consciously consult the notebook. But this explanation of the differences in their behavior requires positing that only Inga has long-term memory about the location of the Museum of Modern Art, and hence that only Inga has a standing belief about the location of the Museum of Modern Art. If this is right, then no extended standing beliefs are involved in the story.

But now, if Otto doesn't have a standing belief about the location of the Museum of Modern Art because he doesn't have long-term memory about its location, then his occurrent belief is not extended either. Inga's occurrent belief does not involve holding *newly learned* information in mind. As we just established, she is retrieving information from long-term memory. She is not relearning the information. Otto's occurrent belief, on the other hand, does not involve retrieving information from long-term memory. It involves learning new information over and over again. So, Otto's case does not involve any kind of extended belief.

Something similar can be said in cases that do not involve Alzheimer's disease. If I remember where the Museum of Modern Art is located but you need to look it up on the iPhone you always carry around, then I have some long-term memory that you don't have, viz. long-term memory of where the museum is located. It's that difference that explains this particular difference in our otherwise very similar behavior. But since standing beliefs are stored in long-term memory, we have different standing beliefs. So, your iPhone does not partially constitute a standing belief about where the museum is

located. But occurrent belief about the location of the museum involves either retrieval of a belief stored in long-term memory or it's a new belief acquired on the basis of perception. Your occurrent belief about the museum's location does not consist in a retrieval of information from long-term memory, since you don't have long-term memory of the museum's location, as we just established. So, the belief that explains why you are walking to 53rd Street is a new belief acquired on the basis of consulting your iPhone. It is not a belief that results from *retrieving* information from the iPhone-serving-as-long-term-memory. The iPhone constitutes neither your standing belief nor your occurrent belief about the location of the museum. So, you don't have any extended beliefs in the constitutive sense.

These considerations do not establish that there are no cases of extended knowledge. They merely undermine a particular argumentative strategy for arguing that external storage devices and search engines are constitutive of belief and belief-based knowledge. Clark and Chalmers' argument is sometimes cashed out an argument from a particular form of integration of an external device into a knowledge-gathering system to the conclusion that the device (or strictly speaking, the informational state of the device) is constitutive of the cognitive system (see e.g., Aizawa 2010). I think Clark and Chalmers' argument is better seen as an argument from explanatory roles. But that is quite irrelevant, if in fact the integration argument is sound. But I don't think it is. A number of factors ensure that the notebook is sufficiently integrated into Otto's knowledge-gathering system. I think the three most important factors are: (i) He is aware of using it as a knowledge-gathering device, (ii) he uses it skillfully and (iii) he relies on it routinely as a knowledge-gathering device. But the inference from this level of integration to the conclusion that the notebook is constitutive of his belief is not valid. This is because the very notion of integration presupposes that the device is part of a knowledge-*gathering* system, that is, it presupposes that there is a system for acquiring internal knowledge that is distinct from the internal knowledge acquired. We will return to that distinction below.

4. Extended Procedural Knowledge

I turn now to procedural knowledge, which I think is a type of knowledge where we can find clear and unsurprising instances of extended knowledge. To see why procedural knowledge is a good candidate to be a type of knowledge that can extend not just beyond the skull but also beyond the skin, let us briefly consider the current debate about knowledge-how. There have traditionally been two main views of knowledge-how: The anti-intellectualist view and the intellectualist view. On the anti-intellectualist view of knowledge-how, originally due to Gilbert Ryle (1946, 1949: ch. 2), one knows how to A only if one has the practical ability to A. The anti-intellectualist account is motivated by examples of the following sort (Brogaard 2011b). Tim is going

on a skiing vacation. In preparation for the trip he carefully studies two renowned skiing 101 books. But upon his arrival at his destination he finds, to his surprise, that he doesn't know how to ski. Fortunately, the ski holiday resort is spawned with able skiing instructors who are more than willing to teach Tim how to ski. As Tim is an excellent scholar, Tim was, prior to his skiing vacation, in the possession of a vast amount of knowledge-that concerning skiing. Tim knew that to slow your speed as a beginner you should use the snow plow position, that to snow plow you must stand with the tips of the skis closer together than the tails, that to turn right your head should move toward the tip of your right ski, and so on. But he still didn't know how to ski. After ten days on the slope with his private skiing instructor Tim had acquired the ability to ski. Only then could Tim claim to know how to ski.

One of the main worries about the anti-intellectualist view is that it doesn't have a way of accounting for cases of procedural knowledge that do not require abilities. Here is an example from John Bengson and Marc Moffett (2007). The Olympic figure skater Irina Slutskaya cannot perform a quintuple salchow. Still, it makes good sense to say that Irina knows how to perform a quintuple salchow. She knows exactly what one ought to do to perform one; she just can't do it.

Perhaps it could be said that while this is a case of knowledge-how, it is not a true form of procedural knowledge. Cases like that of Irina Slutskaya, however, have motivated intellectualists to take knowledge-how to be reducible to some form of knowledge-that (Stanley and Williamson 2001; Bengson and Moffett 2007; Brogaard 2007; Brogaard 2008; Brogaard 2009). According to Jason Stanley and Timothy Williamson, for example, S knows how to F iff for some contextually relevant way w which is a way for S to F, S knows that w is a way for her to F. To account for cases like that of Tim learning how to ski through practice, Stanley and Williamson argue that knowledge-how sometimes requires having the knowledge in question under a certain practical guise. Prior to his skiing vacation Tim had knowledge of how to ski but he didn't have the knowledge under a practical guise. So, it was only once he acquired the ability to ski that he knew how to ski.

I have previously argued for a unified view of knowledge-how (Brogaard 2011b). It's a kind of knowledge-first view that takes certain kinds of ability states that satisfy certain state-external constraints to count as knowledge. So, some forms of knowledge-how are special kinds of belief states and other forms are special kinds of ability states. It may well be that only the cases involving ability states are true cases of procedural knowledge.

On the account I proposed, the relevant kinds of ability states involve both a bodily and a mental component. Consider Paul Snowdon's example of a case where ability and know-how seem to come apart:

> A man is in a room, which, because he has not explored it in the least, he does, as yet, not know how to get out of. In fact, there is an obvious exit which he can easily open. He is perfectly able to get out, he can get out, but does not know how to (as yet). (Snowdon 2003: 11)

It seems perfectly alright to say that the man has the ability to get out of the room (he just has to look around) and yet it seems highly plausible that he doesn't know how to get out. He doesn't know how to get out because there presently is no way w such that he knows that w is a way to get out.

The example trades on an ambiguity in the word 'ability'. In one sense of the word, S has the ability to A just in case s is in an ability state with a content that represents a certain procedure for how to A, and S has the bodily capacities for carrying out the procedure. In another sense, S has the ability to A just in case S has certain bodily capacities which, if combined with the right sort of procedural information, will put S in a position to A. The man in Snowdon's example is not in a state with a content that represents a procedure for getting out. There is a procedure (namely looking around) which, when internalized by the man, will put him in a position to get out. Only the first kind of ability is of the sort possessed by agents in ability states.

'Ability', of course, is frequently used in the latter sense in ordinary language. For example, we might say: 'Of course, you can swim, everyone can swim, you just have to learn it first' or 'Of course, she is perfectly able to walk, she just doesn't know how to yet, she is only eleven months'. Or consider a variation on the man-in-the-room example. To get out one must press a button behind the bookshelves, step on a particular floor plank, and yell 'out' three times. Even so, saying the following seems perfectly fine: 'Of course, the man is perfectly able to get out. He just has to press a button behind the bookshelves, step on a particular floor plank, and yell "out" three times'.

However, in neither the original case nor the variation can we attribute to the agent an ability state with a content that represents a procedure for achieving the intended result. Hence, the agents in these scenarios do not know how to get out (as yet). They have not yet internalized the relatively simple procedures which will lead to their escape. The sorts of abilities that are relevant to knowledge-how are abilities that correspond to procedures which have been internalized by the agent. The view that mental-state-involving abilities are knowledge-how states thus naturally leads to a disjunctive theory of abilities as either (a) states which are constituted by bodily capacities and procedures that have been internalized by the agent, and which are therefore essentially mind-involving, or (b) bodily capacities which, if combined with the right sort of procedural information, will put the agent in a position to achieve the relevant end. The former kind of ability is fundamentally a knowledge-state, viz. a knowledge-how state.

On an account of procedural knowledge that involves abilities states of this kind, procedural knowledge states clearly go beyond the skull. There may be a subclass of procedural knowledge states that do not. For example, if knowledge of how to perform certain arithmetic calculations is a form of procedural knowledge, then that need not reach out of the skull. But in most cases, procedural knowledge requires a certain development of muscles and neural connections between the brain and the muscles. In those cases, the bodily aspect of the ability states are constitutive of the procedural knowledge states.

But there is also a subclass of procedural knowledge that reaches not just beyond the skull but also beyond the skin. Miguel Nicolelis and his colleagues have shown that it is possible for monkeys to make a robot move, using only their brain, even when the robot is thousands of miles away (Nicolelis 2011). The team trained monkeys to walk upright on treadmills. While they were walking, brain activity from multiple cortical areas was recorded to locate the activity that determined the locomotion patterns of the monkeys. The signals from the monkeys' brains were then sent in real time to Japan to make a humanoid robot walk, a robot vastly bigger than the monkeys. The researchers in Japan then sent visual images back from Japan to the laboratory at Duke. These images were projected in front of the monkeys, so the monkeys could see the legs of the robot moving from behind. By rewarding the monkeys when they made the robot make a correct step, the team was able to turn off the treadmills, and the monkeys would then continue making the robot move with their brains as long as they were rewarded. What made this possible was a change in the firing patterns of the neurons of motor areas to assimilate the components of the robot. In this case, it is clear that the capacities of the robotic components partially constitute the monkeys' procedural knowledge states. This shows that procedural knowledge states can extend not just out of the skin but in fact thousands of miles across the globe.

The results also have implications for (declarative) knowledge of action. On a version of realizer functionalism, one of the functions of some forms of knowledge of action is to be able to imitate the action. The sensorimotor areas of the brain are quite plastic and can accommodate radical changes in the bodily components that are represented. The sensorimotor areas of the brain can control vastly larger and heavier bodily components than they normally do. This plasticity of sensorimotor areas suggests that the ability to imitate action must be partially instantiated by a bodily component that has shaped the function of the neural region. Of course, the bodily abilities would not need to be the abilities of any actual body or robotic component, as anything in principle could generate the feedback signals needed for the neural changes to occur. Rather, the ability to imitate an action seems to require the existence of bodily abilities that are either simulated or actually instantiated. If this is true, then we can grant that at least some forms of

knowledge of action extend beyond the skull and the skin. In most of us, the knowledge state that plays this imitation role is a state of the sensorimotor areas of the brain and the abilities of our attached bodies. But in other cases the knowledge state that plays this role is a state of the sensorimotor areas of the brain and the abilities of detached bodies, such as the "bodily" components of robots. So, a case can be made for the view that some forms of knowledge of action are partially constituted by representations of bodily ability that are external to the brain. For example, if the sensorimotor areas normally control a robot, then the knowledge state that plays this role is a state of the sensorimotor areas of the human or monkey brain and the "bodily" abilities of the robot.

5. The Importance of Causal Dependence

So far I have argued that the only forms of knowledge that are genuinely extended are cases of procedural knowledge and some forms of knowledge of action. This conclusion does not allow us to conclude that the mind extends beyond the skin, as it's not the mental representations involved in procedural knowledge that can extend beyond the skin but the bodily abilities. There is, however, a wide range of cases in which mental states causally depend on entities outside the skull and outside the skin, for example, belief states that result from the use of external devices, such as computers and search engines.

Defenders of embodied and extended cognition are unlikely to find this conclusion very exciting. They might say that if they give up on the idea that most mental states are constituted by external entities, then a principled distinction cannot be drawn between their approach and that of traditional cognitive scientists.[2] However, I think that conclusion is too quick. We should not underestimate implications of the weaker alternative defended in this paper.

For example, the weaker view of extended knowledge offered here can provide a novel explanation of how it is possible to gather knowledge while relying extensively on external devices. The very process of knowledge gathering relies on procedural knowledge. Psychologists' knowledge of how to perform experiments, astronomers' knowledge of how to use telescopes and philosophers' knowledge of how to do literature searches are all crucial elements of the knowledge-gathering project. In many of these cases, instruments and tools may well become extensions of the researchers' own bodies in much the same way that the robotic components became extensions of, or even replacements of, the monkeys' bodies in Miguel Nicolelis' experiments. To put an idea expressed by Duncan Pritchard and Orestis Palermos (Pritchard 2010; Palermos and Pritchard 2013) to our own use: the instruments and tools are plausible extensions of the agents' body when they are

appropriately integrated within the agent's cognitive character such that her cognitive success is to a significant degree creditable to her cognitive agency. If this is right, then it is to a large extent *extended* procedural knowledge that allows us to gather semantic knowledge about the world. This is not to say that the semantic knowledge that we gather through these processes is itself extended, but only that we would be unable to gather semantic knowledge without extended procedural knowledge. Why would we be unable to gather semantic knowledge without extended procedural knowledge? Following Pritchard and Palermos, we could say that without the proper integration the collected information does not count as knowledge. This is best understood within the larger framework of virtue epistemology, and I deal with that in the next section. The important point here is that we cannot gather knowledge using tools that are not sufficiently integrated into the knowledge-gathering system. For example, if a guardian angel ensures that you mostly form true beliefs in spite of the fact that your knowledge-gathering procedures are fundamentally flawed, we would want to deny that your true beliefs count as knowledge (Brogaard 2006).

While the weaker hypothesis defended here does have important consequences for our understanding of knowledge, there are also some virtues we lose by adopting only a weaker causal-dependence hypothesis for most knowledge states. One consequence of rejecting the stronger hypothesis is that we lose an easy route from the extended mind hypothesis to content externalism. If a mental state is extended, then an informational state of an external entity is part of the mental state. Since it's highly plausible that the informational content of the state determines the nature of the state, some of the content of extended mental states must be contents of the informational states of external entities and must therefore itself be external. So, given a widely accepted supervenience relation between state and content, extended cognition entails content externalism, but not vice versa. Of course, we can only infer that *some* of the content is external but this position is nonetheless at odds with narrow content internalism.

An interesting question here is whether extended knowledge entails not just content externalism but also externalism about justification. Here I am using the word 'justification' neutrally as denoting one of the normative components that turn belief into knowledge. There does not seem to be any simple argumentative routes from extended knowledge to externalism about justification. On most internalist views of justification, beliefs play an important role as justifications. On the standard belief-views of knowledge, if knowledge is extended, it is likely because the underlying belief state is extended. But it's a separate question whether one has access to the information carried by those belief states, and it certainly seems plausible that one could find reason to defend both extended knowledge and access internalism about justification.

So, by giving up the strong extended view for most forms of knowledge, we lose the easy route from extended knowledge to content externalism but it is unlikely to make a massive difference to arguments about externalism about justification. As I will now show, there is another consequence of rejecting a strong extended knowledge hypothesis: A strong, extended knowledge hypothesis provides us with a simple reply to an extended knowledge objection to virtue epistemology. But as I will argue, a natural reply to more problematic challenges for virtue epistemology also provides a response to the extended knowledge objection.

6. A Situationist-Friendly Virtue Epistemology

It has frequently been argued that when an agent relies extensively on external resources, then her cognitive access should be attributed to the external device rather than the agent herself (Brogaard 2006; Lackey 2007; 2009; Vaesen 2011). This sort of concern is particularly bothersome for virtue epistemologists who take knowledge to be cognitive success that results from the exercise of reliable cognitive mechanisms. Consider again Otto who relies on his notebook in order to gather information about the location of the Museum of Modern Art. As the notebook causally influences Otto's occurrent beliefs about the location of the museum, it could be argued that Otto's cognitive success mostly is due to the notebook rather than his exercise of reliable cognitive mechanisms. But this poses a prima facie problem for virtue epistemologists who grant that the notebook enables Otto to have knowledge of the museum's location and also hold that Otto has that piece of knowledge in virtue of having achieved cognitive success through the exercise of his own cognitive resources.

Defenders of extended semantic knowledge, such as Clark and Chalmers (1998), appear to have a straightforward way of dealing with these types of cases. If Otto's cognitive system is extended, then the cognitive success is a result of the exercise of that cognitive system. The problem arises only once we deny that those cases of knowledge are extended but grant that external resources causally influence those knowledge states. Above we argued that there is reason to think that notebooks are not constitutive of long-term memory and hence are not constitutive of belief. If this is correct, then we need to consider the status of those views that characterize knowledge in terms cognitive success.

A similar problem arises for a brand of virtue epistemology that takes knowledge to be cognitive success that results from the exercise of robust virtuous character traits (such as intellectual conscientiousness, intellectual openness, etc.) and not merely from the exercise of reliable (and sometimes unconscious) cognitive mechanisms. The robust character traits, which are acquired through upbringing, learning etc, enable the person who possesses

them to be cognitively successful in a wide variety of circumstances. So, she will be reliability disposed to form true beliefs in virtue of possessing those robust character traits.

This type of virtue epistemology is affected by the above considerations regardless of whether we take notebooks, iPhone, laptops, etc. to be constituents of knowledge states or merely causal influences. Suppose it's a virtuous trait to tend to reflect carefully on the evidence. An Alzheimer's patient like Otto will eventually lose many of these virtuous traits. He might jump to the conclusion that his wife is having an affair with the neighbor on the basis of seeing her talk to him. However, for a while Otto may be able to use external resources to help make up for deficits in his cognitive abilities. For example, he may make a habit out of using an app on his iPhone that can evaluate how well a given piece of evidence supports a given conclusion. As a result he may continue to be cognitively successful for a long time. But regardless of whether one takes beliefs to be extended or not, it could be argued that external resources are not constitutive of virtuous personality traits. So, Otto's successful occurrent belief that his wife is not having an affair with the neighbor is not the result of the exercise of virtuous personality traits.

If we want to maintain one of these versions of virtue epistemology, I don't think the solution is to argue that reliable cognitive mechanisms or virtuous personality traits are partially constituted by external resources. The reason is that a similar problem arises independently of considerations of which external resources are involved in achieving cognitive success. The problem is that virtue epistemology seems at odds with scientific evidence concerning our actual epistemic behavior. Evidence from experimental work apparently suggests that our intellectual behavior in a particular situation is so massively influenced by *epistemically irrelevant situational features* that there is no consistent way that we behave over time (see e.g., Olin and Doris 2013). Hence, neither cognitive abilities nor so-called character traits enable the agents who possess them to be cognitively successful in a wide range of situations. Thinkers who make this point on the basis of experimental findings point to the influence of epistemically irrelevant situational factors in almost all circumstances. Judges who have had lunch make more lenient judgments than hungry judges. A person's intellectual abilities based on that person's CV are judged to be more impressive if the name on the CV is male compared to when it's female. Black students' SAT scores will vary depending on whether they are asked about their ethnicity before or after the test. And I could go on and on (for a wide variety of interesting examples see e.g. Tavris and Aronson 2007). The empirical literature is spawned with examples of how irrelevant epistemic factors can influence our cognitive performance. These situational influences are hardly something we can easily rise above through *reflection*, as they tend to affect judgments through rapid subpersonal mechanisms. Certain forms of training and education may help

us become less affected by them. Yet if the situational influences are as widespread as situationists claim, then there does seem to be a challenge here, a challenge that is potentially more devastating than that from deliberate reliance on external resources. The external resources, unlike many of the situational factors situationists cite, are epistemically relevant and engaged by us on a personal level rather than a subpersonal level.

One response to the situationist challenge rests on the observation that even if cognitive agents have no robust character traits or cognitive skills that remain stable across a wide range of situations, agents nonetheless tend to have stable situation-dependent traits or skills. As Peter Railton (2011) argues, there are significant individual differences in people's personal dispositional profiles or signatures. These signatures are 'if, then' dispositions specifying how people behave cognitively in particular types of situations. Consider how people behave cognitively in homely settings versus stressful test-taking situations. Stress can impact cognition differently in different individuals. So, even if very few individuals will have stable cognitive dispositions across these situations, there are still significant differences in how individuals behave cognitively in the stressful situations, which may or may not mirror the individual differences in the homely situations. These personal dispositional signatures are too narrow to serve as a new type of traits on which to hang virtue epistemology. But we can take the signatures that lead to cognitively successful behavior in favorable or neutral situations to count as the virtuous traits. The unfavorable situations are those situations in which *epistemically irrelevant* factors significantly shape cognitive behavior. Whether we have eaten lunch or not will not typically affect our ability to accurately perceive our surroundings, though it clearly can affect other types of cognitive processes. An example of a favorable situation is one in which we rely on new technology to improve our reasoning abilities. When you rely on your calculator to multiply large numbers, the use of the calculator makes the situation favorable, other things being equal. Situations in which visibility is not optimal, in which we don't carefully reflect on the evidence, in which the tools we use to gather knowledge are not sufficiently integrated into our cognitive character, and so on, are unfavorable situations, regardless of whether they, by sheer coincidence, afford cognitive success. The cognitive abilities and character traits and combinations thereof that are reliable *in an epistemically relevant sense* are those that are reliable in favorable or neutral situations. We can then say that knowledge is cognitive success that results from exercising those of our cognitive abilities and character traits that are reliable in favorable and neutral situations *in those situations*.

The response to the objection from the causal influences on our cognitive successes, then, is to regard these influences as being either epistemically relevant or irrelevant. When they are epistemically relevant by being cognitively integrated and improving our cognitive abilities, the cognitive success that results from the knowledge-gathering process is knowledge. The credit

attributed to the agent in the case of cognitive successes that causally depend on external resources may turn out to vary widely across situations. That is, epistemic action that relies on external resources 'demands spread of *epistemic credit*' (Clark and Chalmers 1998: 8). For example, using other people's testimony may in some cases be so easy that the agent who is gathering information may deserve very little credit (Lackey 2007, 2009). But this spread of epistemic credit is only a problem for credit views of knowledge, not for virtue epistemology.

7. Conclusion

On the constitutional view of extended knowledge, a knowledge state is extended iff entities that extend beyond the skin are constituents of the knowledge state. On a weaker view, knowledge states are not extended in this way, though knowledge states causally depend on the knower's body and external environment. I have argued here that the only candidates to be extended knowledge states in the constitutive sense are knowledge of action and procedural knowledge. However, this limited extended knowledge view can provide us with an explanation of how tools and devices that appear to extend our bodies can be so crucial to the knowledge-gathering process.[3]

Notes

1. So, unlike what Chalmers sometimes says, he doesn't lose his mind if he loses his iPhone as long as he has proper backup of the information.
2. In the 1960s traditional cognitive scientists treated the human brain as a computational problem-solving device. On this approach, the body is an output system attached to a cognitive processing system the way that the engine and the rotors of a helicopter may be attached to a computer that turns the helicopter's parts on and off. Traditional cognitive scientists furthermore tended to ignore that the particular bodies we have may shape our cognitive resources. For example, if our bodies were unable to bend the way they actually are, our concept of a chair might have been different. But this was not seen as an important issue in traditional cognitive science.
3. I am grateful to Orestis Palermos for very helpful comments on a previous version of this paper.

References

Adams, F. and Aizawa, K. (2001). "The Bounds of Cognition." *Philosophical Psychology* 14: 43–64.

Adams, F. and Aizawa, K. (2008). *The Bounds of Cognition*. Malden, MA: Blackwell Publishers.

Aizawa, K. (2010). "The Value of Cognitivism in Thinking About Extended Cognition." *Phenomenology and the Cognitive Sciences* 9 (4): 579–603.

Bengson, J. and M. Moffett. (2007). "Know-How and Concept-Possession." *Philosophical Studies* 136: 31–57.

Brogaard, B. (2006). "Can Virtue Reliabilism Explain the Value of Knowledge?" *Canadian Journal of Philosophy* 36: 335–354.

Brogaard, B. (2007). "Attitude Reports: Do you Mind the Gap?" In B. Weatherson, ed., *Philosophy Compass*.

Brogaard, B. (2008). "Knowledge-The and Propositional Attitude Ascriptions." *Grazer Philosophische Studien* 77: 147–190.

Brogaard, B. (2009). "What Mary Did Yesterday: Reflections on Knowledge-wh." *Philosophy and Phenomenological Research* 78: 439–467.

Brogaard, B. (2011a). "Primitive Knowledge Disjunctivism." *Philosophical Issues* 21: 45–73.

Brogaard, B. (2011b). "Knowledge-How: A Unified Account." In *Knowing How: Essays on Knowledge, Mind, and Action*, J. Bengson and M. Moffett, eds., Oxford: Oxford University Press: 136–160.

Brogaard, B. (2011c). "Are There Unconscious Perceptual Processes?", *Consciousness and Cognition* 20 (2011): 449–63.

Brogaard, B. (2012). "Vision for Action and the Contents of Perception." *Journal of Philosophy* 109 (10): 569–587.

Calvo-Merino, B., Glaser D. E., Grezes J., Passingham R. E., and Haggard P. (2005). "Action Observation and Acquired Motor Skills: An fMRI Study with Expert Dancers." *Cerebral Cortex* 15: 1243–1249.

Clark, A. (2005). "Intrinsic Content, Active Memory, and the Extended Mind." *Analysis* 65: 1–11.

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

Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: Oxford University Press.

Clark, A. and Chalmers D. J. (1998). "The Extended Mind." *Analysis* 58 (1): 7–19.

Decety, J., Grèzes, J., Costes, N., Perani, D., Jeannerod, M., Procyk, E., et al. (1997). "Brain Activity During Observation of Actions. Influence of Action Content and Subject's Strategy." *Brain* 120: 1763–1777.

Ganis, G., Thompson W. L., Kosslyn S. M. (2004). "Brain Areas Underlying Visual Mental Imagery and Visual Perception: an fMRI Study." *Brain Res Cogn Brain Res.* 20 (2): 226–41.

Goodale, M. A. and Milner, A. D. (1992). "Separate Visual Pathways for Perception and Action." *Trends in Neurosciences* 15: 20–25.

Goodale, M. A., Milner, A. D., Jakobson, L. S., and Carey, D. P. (1991). "A Neurological Dissociation between Perceiving Objects and Grasping Them." *Nature* 349: 154–156.

Grezes, J., Costes, N. and Decety, J. (1998). "Top-Down Effect of Strategy on the Perception of Human Biological Motion: A PET Investigation." *Cognitive Neuropsychology* 15: 553–582.

Heeger, D. J. (1999). "Linking Visual Perception with Human Brain Activity." *Current Opinion in Neurobiology* 9: 474–479.

Lackey, J. (2007). "Why We Don't Deserve Credit for Everything We Know." *Synthese* 158 (3): 345–61.

Lackey, J. (2009). "Knowledge and Credit." Philosophical Studies 142: 27–42.

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

Menary, R. (2007). *Cognitive Integration: Mind and Cognition Unbound*. London: Palgrave Macmillan.

Milner, D. and Goodale, M. A. (1995). *The Visual Brain in Action*. New York: Oxford University Press.

Milner, A. D. and Goodale, M. A. (2008). "Two Visual Systems Re-Viewed." *Neuropsychologia* 46: 774–785.
Nicolelis, M. (2011). *Beyond Boundaries: The New Neuroscience of Connecting Brains with Machines—and How It Will Change Our Lives*. New York: Times Books.
Nöe, A. (2004). *Action in Perception*. Cambridge, MA: MIT.
Nöe, A. (2009). *Out of Our Heads: Why You Are Not Your Brain, and Other Lessons from the Biology of Consciousness*. New York: Hill and Wang.
Olin, L. and Doris, J. M. (2013). "Vicious Minds." *Philosophical Studies*. doi:10.1007/s1109801301533
Palermos, O. and Pritchard, D. (2013). "Extended Knowledge and Social Epistemology." *Social Epistemology Review and Reply Collective* 2 (8): 105–120.
Pritchard, D. H. (2010). "Cognitive Ability and the Extended Cognition Thesis." *Synthese* 175 (1): 131–51.
Railton, P. (2011). "Two Cheers for Virtue, or, Might Virtue be Habit Forming?" *Oxford Studies in Normative Ethics* 1: 295–330.
Rissman, J. and Wagner A. D. (2012). "Distributed Representations in Memory: Insights from Functional Brain Imaging." *Annual Review of Psychology* 63: 101–128.
Rowlands, M. (2009). "Extended Cognition and the Mark of the Cognitive." *Philosophical Psychology* 22 (1): 1–19.
Ryle, G. (1946). "Knowing How and Knowing That." In *G. Ryle*, *Collected Papers*, *vol. II: Collected Essays (1929–1968),* New York: Barnes and Noble, Inc.
Ryle, G. (1949). *The Concept of Mind*. New York: Barnes and Noble, Inc.
Snowdon, P. (2003). "Knowing How and Knowing That: A Distinction Reconsidered." *Proceedings of the Aristotelian Society* 104: 1–29.
Squire, L. R. (2004). "Memory Systems of the Brain: A Brief History and Current Perspective." *Neurobiology of Learning and Memory* 82 (3): 171–177.
Stanley, J. and Williamson, T. (2001). "Knowing How." *Journal of Philosophy* 98: 411–44.
Tavris, C. and Aronson, E. (2007). *Mistakes Were Made (But Not by Me): Why We Justify Foolish Beliefs, Bad Decisions, and Hurtful Acts*. New York: Harcourt.
Theiner, G. (2011). *Res Cogitans Extensa: A Philosophical Defense of the Extended Mind Thesis*. Bern, Switzerland: Peter Lang GmbH, Europaischer Verlag der Wissenschaften.
Vaesen, K. (2011). "Knowledge without Credit, Exhibit 4: Extended Cognition." *Synthese* 181: 515–529.
Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.