THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Cognitive ability and the extended cognition thesis

Citation for published version:

Pritchard, D 2010, 'Cognitive ability and the extended cognition thesis', *Synthese*, vol. 175, pp. 133-151.
<https://doi.org/10.1007/s11229-010-9738-y>

Digital Object Identifier (DOI):

[10.1007/s11229-010-9738-y](https://doi.org/10.1007/s11229-010-9738-y)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Synthese

Publisher Rights Statement:

Pritchard, D. (2010). Cognitive ability and the extended cognition thesis. *Synthese*, 175, 133-151doi: 10.1007/s11229-010-9738-y. The final publication is available at link.springer.com

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

COGNITIVE ABILITY AND THE EXTENDED COGNITION THESIS

DUNCAN PRITCHARD

University of Edinburgh

ABSTRACT. This paper explores the ramifications of the extended cognition thesis in the philosophy of mind for contemporary epistemology. In particular, it argues that all theories of knowledge need to accommodate the ability intuition that knowledge involves cognitive ability, but that once this requirement is understood correctly there is no reason why one could not have a conception of cognitive ability that was consistent with the extended cognition thesis. There is thus, surprisingly, a straightforward way of developing our current thinking about knowledge such that it incorporates the extended cognition thesis.

0. INTRODUCTORY REMARKS

One of the most interesting proposals in the recent literature in the philosophy of mind has been the suggestion that there is no in principle bar to cognitive processes extending beyond the skin of the agent. As one of the foremost exponents of this thesis has expressed the matter, “cognitive processes are not located exclusively inside the skin of cognizing organism” (Rowlands 1999, p22). This is the so-called *extended cognition* thesis. This thesis poses a radical challenge to our normal theorising about cognition which largely takes it as given that cognitive processes take place exclusively under the skin of the agent. Indeed, insofar as one treats cognitive processes as mental processes, then the extended cognition thesis generates a more radical conclusion still—*viz.*, that the mind can extend beyond the skin of the agent too (the *extended mind* thesis). Even so, there does seem to be some very strong support available for the extended cognition thesis, and so it is a proposal that must be taken seriously, not only by philosophers but also by anyone working in a research field

which is concerned with the nature of cognition, such as psychology, linguistics and informatics.¹

Given the wealth of literature that has been generated by the extended cognition thesis in the philosophy of mind (not to mention in other fields outwith philosophy, such as informatics), it is puzzling that so little has been done to explore the epistemological ramifications of this proposal, especially since this is clearly a thesis which has important implications for epistemology. The aim of this paper is to try to rectify this by examining how the theory of knowledge might accommodate this proposal. As we will see, far from being an uncomfortable fit with contemporary thinking about knowledge, there is in fact a key thread in current epistemology—regarding the relationship between cognitive ability and knowledge—which fits quite snugly with the extended cognition thesis. Indeed, reflecting on the relationship between knowledge and cognitive ability may arguably offer some additional support for the extended cognition thesis.

1. KNOWLEDGE AND COGNITIVE ABILITY

One of the master intuitions about knowledge which guides contemporary epistemology is the idea that knowledge is the product of cognitive ability. A true belief, no matter what else of epistemic relevance can be offered in its favour (e.g., that it is safe, sensitive, backed by reasons, epistemically blameless, and so on), will not count as a case of knowledge if it is not the product of cognitive ability. Call this the *ability intuition*.

Consider, for example, an agent who is, unbeknownst to her, a brain in a vat (BIV) who is being ‘fed’ beliefs about her environment by supercomputers. Such an agent might nonetheless form true beliefs about her environment. Moreover, those true beliefs so formed may well be modally stable in the relevant ways across a range of possible worlds such that they are both safe (i.e., they could not easily have been false) and sensitive (i.e., were the proposition believed to be false, then the agent would no longer have believe it).² The beliefs might well have additional epistemically relevant properties too, such as being backed by good reasons (and so justified, in an epistemically internalist sense at least) and epistemically blameless (in that no epistemic norm is being intentionally, or at least culpably, flouted).³ Even so, no-one in contemporary epistemology supposes that these true beliefs can amount to

knowledge, and the most natural explanation of why is that the fact that these beliefs are true in no way reflects the cognitive abilities of the agent. After all, despite the best cognitive efforts of the agent, the reason why her beliefs are true has nothing at all to do with these cognitive efforts, and reflects instead merely the work of the supercomputers.

In crediting an agent with knowledge we are thus, amongst other things, crediting her with having a relevant cognitive ability which played some key part in the production of the target true belief. An adequacy condition on any theory of knowledge is thus that it is able to accommodate the ability intuition. This entails that all theories of knowledge should include an ability condition of some sort (i.e., an epistemic condition which can accommodate the ability condition), or should at least include an epistemic condition which, amongst other things, does the work of an ability condition. Inevitably, there has been some debate about what constitutes a cognitive ability and thus regarding how strong the ability condition on knowledge should be.

One might think that a cognitive ability is simply a belief-forming process which is reliable (i.e., more likely, by a significant margin, to generate true beliefs rather than false beliefs). This account of cognitive ability is, however, too weak. This is because one can formulate cases where agents have formed true beliefs via reliable belief-forming processes but where the reliability in question does not reflect a cognitive ability on the part of the agent. Indeed, we can use the BIV case just noted to make this point, since we can stipulate that the supercomputers *only* ‘feed’ true beliefs to the agent concerned. To take a simple illustration, suppose that our hero is only fed beliefs about the current time (GMT) and the location of the British Prime Minister, but that these beliefs are always true. In the relevant sense such beliefs are reliable, in that they are indeed more likely to be true than false (indeed, they are pretty much guaranteed to be true). But this is not a case of knowledge, and the primary reason for this is that such reliability does not reflect a cognitive ability on the part of the agent.

Note that the problem here is not simply that we are dealing with a sceptical scenario (i.e., one designed to call the agent’s knowledge into question *en masse*), since we can run similar examples which lack this feature. One such case is that of ‘Temp’, who regularly forms his beliefs about the temperature in the room by consulting a thermometer on the wall. Unbeknownst to Temp, however, he is looking at a faulty thermometer which is randomly fluctuating within a given range. Nonetheless, his beliefs so formed are reliable because

(again unbeknownst to him) there is someone hidden in the room next to the thermostat who every time she sees Temp head towards the thermometer adjusts the thermostat so that the temperature of the room corresponds to the reading on the thermometer. Accordingly, Temp's way of forming his beliefs is highly reliable. Clearly, however, Temp does not know the temperature of the room, and the reason for this is that his reliability does not reflect his cognitive ability at all, but merely the helpful assistance of the hidden helper.⁴

In both the Temp and the BIV case the source of the reliability is external to the skin of the agent, and so one might think that this is crucial to these cases such that so long as the source of the reliability is under the agent's skin then that would suffice to ensure that the agent's true beliefs so formed are the result of her cognitive ability, and hence in the market (on this score at least) for knowledge. Interestingly, however, we can easily imagine a case which has the relevant features as the Temp and BIV examples but where the source of the reliability is under the skin of the agent. Consider, for example, 'Alvin' who has an unusual brain lesion, a guaranteed side-effect of which is that it prompts him to randomly, but reliably, form true beliefs about the product of fairly complicated arithmetical sums. So described, this is certainly a reliable belief-forming trait that Alvin has, and yet intuitively Alvin does not have knowledge of these mathematical propositions. As in the Temp and BIV cases, the primary reason for this is that it clearly has nothing to do with Alvin's cognitive abilities that his beliefs in this regard are true, but is instead a fortunate consequence of the otherwise unfortunate fact that he has a brain lesion. In short, Alvin's true beliefs in the target propositions do not satisfy the ability intuition.⁵ Nonetheless, the source of the reliability of Alvin's belief-forming process is clearly beneath his skin.

So if one is to capture the requirement that knowledge involves a true belief that is due to cognitive ability, a mere appeal to the reliability of the target belief-forming process—even where the source of that reliability lies under the skin of the agent—will not suffice. So what in addition is required for a genuine cognitive ability to be manifested? As a number of epistemologists have noted, the answer to this question lies in the extent to which the reliable belief-forming process is *integrated* within, and therefore a part of, the cognitive character of the agent, where an agent's cognitive character is her integrated web of stable and reliable belief-forming processes.⁶ In particular, what is required is that the reliable belief-forming process in question needs to be sufficiently integrated within the agent's cognitive character such that one would—to a significant degree anyway—credit the

cognitive success (i.e., the true belief) that results to the agent's cognitive agency rather than to some other factor. As we will now see, however, there are two main proposals put forward in this regard in the recent literature, and so in order to explore this matter further we first need to state these two views.

2. STRONG AND WEAK ACCOUNTS OF COGNITIVE AGENCY

We can formulate these two accounts of cognitive agency—specifically, of the relationship between cognitive ability and knowledge—that can be found in the recent literature as follows:

- (COGA_{WEAK}) If S knows that p , then S 's true belief that p is the product of a reliable belief-forming process which is appropriately integrated within S 's cognitive character such that her cognitive success is to a significant degree creditable to her cognitive agency.
- (COGA_{STRONG}) S knows that p iff S 's true belief that p is the product of a reliable belief-forming process which is appropriately integrated within S 's cognitive character such that her cognitive success is primarily creditable to her cognitive agency.

Note that both these theses in effect treat the question of the extent to which the cognitive success is creditable to the agent's cognitive agency as being the test by which one determines whether it should count as the product of cognitive ability. More generally, on both views one can use this test to determine whether a reliable belief-forming process is appropriately integrated within an agent's cognitive character such that it counts as a *bona fide* cognitive ability.⁷ Where the two accounts differ is on the extent to which the cognitive success in question is creditable to the agent's cognitive agency. Whereas the strong account, (COGA_{STRONG}), insists on a very demanding relationship between cognitive success and cognitive agency on this score, such that knowledge can be defined in terms of the satisfaction of this relationship, the weaker account, (COGA_{WEAK}), allows other factors to be substantively relevant to one's acquisition of knowledge.

It should be clear that both accounts of the relationship between knowledge and cognitive agency can explain why agents like Temp and Alvin lack knowledge. In the former case it is not Temp's cognitive agency which is producing his cognitive success, but rather a feature of the environment (i.e., the activities of his 'helper'). Accordingly, we do not credit

Temp's cognitive success to his cognitive agency to any significant degree, and hence he fails to satisfy both (COGA_{STRONG}) and (COGA_{WEAK}). The case of Alvin is different, in that it is not a feature of the environment which is the source of Alvin's cognitive success but rather a feature of his brain, but the essential point is the same. After all, it has nothing to do with Alvin's cognitive agency that he is forming true beliefs in this fashion, and hence he also fails to satisfy both (COGA_{STRONG}) and (COGA_{WEAK}). Indeed, we would naturally say that is *in spite of* Alvin's cognitive agency that he is cognitively successful in this regard. So even where the source of one's cognitive success lies under the skin of the agent, it can still be the case that the source of that cognitive success lies outwith the agent's cognitive agency.

Where these two accounts of the relationship between knowledge and cognitive agency will come apart will be in terms of what they would demand of the agents in the Temp and Alvin cases before they could count as having knowledge of the target propositions. That is, they will hold conflicting accounts of what Temp and Alvin would need to do to integrate their reliable belief-forming processes within their cognitive character. For now, let us focus on what the strong account would require in this regard, since what suffices for the strong account will *a fortiori* also suffice for the weak account as well (we will return to consider weaker accounts of cognitive integration below). What would clearly suffice in this regard would be if the agent concerned were to come to know both what the true source of the reliability of his belief-forming process was and that it was reliable.⁸

For example, suppose that Temp became appropriately aware that he had a helper in this way, perhaps initially through the testimony of a trustworthy informant and then subsequently by covertly observing the helper in action. Temp thus comes to know not just that the reading on the thermometer is a reliable indicator of the ambient temperature, but also gains knowledge of why this is the case even despite the fact that the thermometer is broken. Would this make a difference to the epistemic standing of Temp's beliefs? Intuitively, it would, in that Temp could now come to know what the temperature of the room is by observing the reading on the thermometer. But notice that it is also now true of Temp that his cognitive success is primarily creditable to his cognitive agency. After all, Temp is now able to take cognitive responsibility for this cognitive success. In becoming aware of the relevant facts, then, this hitherto merely reliable belief-forming process becomes integrated within Temp's cognitive character, such that his cognitive success is now primarily creditable to his cognitive agency. In this way, mere reliability in his belief-forming process is

converted into the exercise of a genuine cognitive ability, one that can in principle deliver knowledge.

The same goes for the case of Alvin. For now suppose that Alvin becomes aware that there exists brain lesions of this sort, and gains additional good grounds for supposing that he possesses just such a brain lesion, such that he now knows that he has one of these brain lesions. Perhaps, for example, he comes across an article about these brain lesions in a reliable newspaper and researches the matter in reliable medical journals and on this basis comes to know that he is the victim of the brain lesion in question. Intuitively Alvin's beliefs in the target mathematical propositions now qualify as knowledge. But notice that it is also true that Alvin has in this way integrated this belief-forming process within his cognitive character to a sufficient degree that his cognitive success is now primarily creditable to his cognitive agency, rather than being creditable to something external to his cognitive agency (albeit a factor which was under his skin, and hence in this sense internal). Like Temp, Alvin knows both that his beliefs in this regard are reliable and also what the source of this reliability is, and as a consequence he is able to take cognitive responsibility for his cognitive success in this regard.⁹

There is thus an account of cognitive integration available which is compatible with both (COGA_{WEAK}) and (COGA_{STRONG}). This raises the question of which of these two accounts of the relationship between knowledge and cognitive ability that we should opt for. Relatedly, it also raises the question of whether there is a kind of cognitive integration which suffices for knowledge but which is only accounted for by (COGA_{WEAK}). Interestingly, in the recent literature it has been (COGA_{STRONG}) that has had the most adherents, with a view of roughly this sort endorsed by such figures as Ernest Sosa (1988; 1991; 2007), Linda Zagzebski (1996; 1999), and John Greco (2002; 2007; 2008; 2009).¹⁰ The reason why these epistemologists are attracted to this more robust account of the relationship between knowledge and cognitive ability is that it has a number of theoretical advantages over the weaker proposal. In particular, (COGA_{STRONG}) offers an elegant and highly plausible account of knowledge while also casting light on some central problems in epistemology, such as concerning the nature of epistemic value.¹¹ Of course, (COGA_{STRONG}) only has these theoretical advantages if the proposal is true. Unfortunately, as we will now see, this proposal faces some pretty formidable obstacles, obstacles which are not faced by its competitor principle, (COGA_{WEAK}).

The difficulties come from two directions. First, there are worries about the right-to-left entailment in this biconditional, in that there seem to be cases where an agent's cognitive success is primarily creditable to her cognitive agency but where she lacks knowledge nonetheless. The clearest examples in this respect concern cases where the agent's belief in the target proposition is only luckily true on account of some unusual feature of the environment. This point gets missed by the proponents of (COGA_{STRONG}) because they tend to focus on cases involving knowledge-undermining epistemic luck where the luck in question is not environmental in this way.

Consider, for example, the following case involving 'Roddy', which follows the same basic model as the cases famously offered by Edmund Gettier (1963).¹² Roddy looks into a field and using his highly reliable perceptual abilities he spots a sheep-shaped object. Accordingly, he forms a belief that there is a sheep in the field. Moreover, this belief is true, in that there is indeed a sheep in the field. The twist in the tale, however, is that what Roddy is looking at is not in fact a sheep at all, but rather a big hairy dog which is obscuring from view the real sheep hidden behind. His belief is thus only true as a matter of luck, in the sense that his true belief is *unsafe*—viz., it could so very easily have been false. Intuitively, Roddy does not have knowledge of what he believes, even though his belief is true and the product of his reliable perceptual abilities. (COGA_{STRONG}) offers a straightforward explanation of why this is the case. After all, the luck involved in the formation of Roddy's true belief would surely suffice to ensure that his cognitive success would not be primarily creditable to his cognitive agency, but would be instead credited to the happy co-incidence that there happened to be a sheep in the field hidden from view behind the sheep-shaped object.

So far so good for (COGA_{STRONG}). The problem, however, is that not every form of knowledge-undermining epistemic luck is of this Gettier-style kind, with 'environmental' epistemic luck being a case in point. In Gettier-style cases of the sort just described something 'intervenes' between the agent's cognitive ability and her cognitive success, albeit in such a way that it does not undermine the agent's cognitive success. The Roddy case offers a vivid illustration of this phenomenon, in that Roddy thinks that he is looking at a sheep, and does truly believe as a result that there is a sheep in the field, but it is not a sheep that he is in fact looking at. Compare this case with an example where environmental epistemic luck is present.

Consider the case of 'Barney' who uses his reliable perceptual faculties to form a

belief that there is a barn before him. This belief is true, in that there is a barn before him. Like the Roddy case, this belief is unsafe, and thus is only true as a matter of luck. Unlike the Roddy case, however, Barney really is looking at a genuine barn, and hence nothing has ‘intervened’ between his cognitive ability and his cognitive success in the way that it does in the Roddy case. Instead, the source of the luck is entirely in the environment, in that Barney just happens to be in an environment where almost everything that looks like a barn is in fact a clever fake (a barn façade), and were Barney to have happened to look at any of these fakes then he would have continued to believe that what he is looking at is a barn regardless.¹³

The problem posed by the Barney case for defenders of (COGA_{STRONG}) is that unlike the Roddy case it does seem right to say that Barney’s cognitive success is primarily creditable to his cognitive agency, even though the luckiness of his true belief means that it is not in the market for knowledge. More generally, while environmental epistemic luck is just as knowledge-undermining as standard Gettier-style epistemic luck, it seems entirely compatible with the cognitive success in question being primarily creditable to the agent’s cognitive agency. After all, Barney really does employ his reliable perceptual faculties to see a genuine barn. The moral seems to be that there is sometimes more to having knowledge than having a cognitive success which is primarily creditable to one’s cognitive agency.

A straightforward way for the proponent of (COGA_{STRONG}) to respond to this problem would be to weaken the account slightly so that it is only offering a necessary condition for knowledge, rather than a necessary and sufficient condition. This would, in any case, make the strong account analogous to the weak account which itself only offers a necessary condition on knowledge. We thus get the following principle:

(COGA_{STRONG*}) If *S* knows that *p* then *S*’s true belief that *p* is the product of a reliable belief-forming process which is appropriately integrated within *S*’s cognitive character such that her cognitive success is primarily creditable to her cognitive agency.

It should be clear that (COGA_{STRONG*}) does not face the same problem as (COGA_{STRONG}) when it comes to the Barney case, since it is open to the defender of (COGA_{STRONG*}) to grant that Barney lacks knowledge on account of his failure to satisfy some further necessary condition which is required for knowledge (e.g., a safety condition).

Unfortunately, even this weaker construal of the strong account of cognitive agency is not immune to counterexamples, for it seems that the account on offer is still too austere. The most straightforward examples which illustrate this concern cases where agents appropriately

trust the word of others and thereby gain testimonial knowledge.

Consider the case of ‘Jenny’.¹⁴ Jenny gets off the train in an unfamiliar place and asks the first person she meets for directions. This person is indeed fully knowledgeable about the area and communicates this knowledge to our hero, who subsequently forms a true belief about the way to go. Does Jenny gain knowledge in this case? I think the natural answer is that she does, at least provided we read the details of the case in the right way. For I take it that we are assuming here that Jenny is in an epistemically friendly environment—it is not as if, for example, this town is renowned for its dishonest informants. Moreover, I take it that we are also reading into the case that Jenny is suitably responsive to epistemically relevant factors—it is not as if, for example, she would ask someone who would clearly not be a good informant (e.g., someone who was clearly a tourist), and it is not as if she would believe whatever she was told, even when it was obviously false. With these details of the case made explicit, then I think our intuitions are clear that Jenny does gain knowledge. Indeed, if Jenny does not gain knowledge in this case then, given that quite a lot of what we take ourselves to know we gained on a similar basis of trust, it follows that we have significantly less knowledge that we typically suppose ourselves to have.

Crucially, however, it does not seem plausible to suppose that Jenny’s cognitive success is primarily creditable to her cognitive agency. After all, given that the bulk of the cognitive work here was done by Jenny’s informant, then it would seem that if any particular person is primarily creditable for Jenny’s cognitive success then it is her rather than Jenny.¹⁵ Note that this is not to say that Jenny’s cognitive success is not to a significant degree creditable to her cognitive agency, since as we noted above we need to read the example such that Jenny’s cognitive success is substantively creditable to her cognitive agency before it counts as an intuitive case of knowledge in the first place (i.e., this is not a counterexample to the ability intuition). Still the point remains that Jenny’s cognitive success is not primarily creditable to her cognitive agency, and hence even (COGA_{STRONG*}) is under threat.

Indeed, once we diagnose what is going on in this case, we begin to see why the strong account of cognitive agency—and thus of the relationship between knowledge and cognitive ability—is so implausible (in both its forms). For the crux of the matter in the Jenny case is that in epistemically favourable environments knowledge is sometimes gained even though very little cognitive ability is exercised on the part of the subject. One can see this point by noting that had Jenny been in an epistemically unfavourable environment and

exhibited the same degree of cognitive ability then we would not have credited her with knowledge. There is thus an interplay between the extent of the cognitive ability required for knowledge and the epistemic favourableness of the environment for forming true beliefs in that case.

That there is this interplay between cognitive ability and favourable environments is also demonstrated by the Barney case. For remember that the key feature of this case is that even despite the cognitive ability exhibited by Barney, such that his cognitive success is primarily creditable to his cognitive agency, he still does not count as having knowledge because of how epistemically *unfriendly* his environment is. In contrast, had the environment been friendlier—if, in particular, his belief so formed had not been subject to environmental epistemic luck—then he would have counted as having knowledge.

The moral thus seems to be that while sometimes the exercise of very little cognitive ability can suffice for knowledge, equally sometimes the exercise of a great deal of cognitive ability can fail to suffice for knowledge, with in each case the crucial factor being the friendliness of the cognitive environment. But note that this point decisively favours a weak account of cognitive agency—and thus of the relationship between knowledge and cognitive ability—over its stronger counterparts. For what is not in doubt in the Jenny case is that (COGA_{WEAK}) is satisfied, as noted above. Moreover, given that (COGA_{WEAK}) only sets a necessary condition on knowledge, it is not called into question by the Barney case and cases like it which demonstrate that the exhibition of a high level of cognitive ability does not suffice for knowledge. Indeed, by the same token (COGA_{WEAK}) is not called into question by cases involving standard Gettier-style luck either (like the Roddy case), for the natural explanation of what is going on here is that while the agent is indeed exhibiting a significant degree of cognitive ability, and is in addition cognitively successful, the epistemic luck involved means that the agent's cognitive success is not creditable to her cognitive agency since it ensures that there is no appropriate connection between the agent's cognitive agency and her cognitive success (i.e., such cases reinforce the ability intuition).¹⁶

We are thus led to a relatively weak account of cognitive agency, and thus of cognitive ability, albeit one that demands far more of a belief-forming process if it is to be knowledge-conducive than that it be merely reliable. As we will see, such an account of cognitive agency in fact fits rather neatly with the extended cognition thesis.

3. THE EXTENDED COGNITION THESIS

We noted above that the mere fact that the source of the reliability of an agent's belief-forming process is beneath the skin does not suffice to ensure that this process represents a cognitive ability on the part of the agent—i.e., a cognitive process which can generate knowledge. What is important is rather that the reliable belief-forming process be integrated within the cognitive character of the agent such that the cognitive success in question is to a significant degree creditable to the agent's cognitive agency. Now one might naturally infer from this that 'cognitive ability' is a sub-class of 'reliable belief-forming process where the source of the reliability is completely under the skin of the agent'. Notice, however, that nothing that we have said so far supports this entailment since we have not even considered potential cases of cognitive ability where the latter characterisation does not hold. Nevertheless, it is common currency within epistemology to assume that something like the above entailment holds. In effect, what is being assumed here is that the extended cognition thesis is false.

Now many epistemologists will hold that this is an entirely reasonable assumption, on the grounds that factors relevant to one's cognition which are outwith one's skin can play at most a causal or instrumental role. For example, that one uses a pen to work out a mathematical problem doesn't make the pen part of the cognitive process—instead, it is just an instrument which helps the cognitive process to take place. We can cast this response in terms of the (COGA_{WEAK}) thesis outlined above by observing that were one to gain knowledge of the mathematical proposition in question then we would credit one's cognitive success to one's cognitive agency (e.g., one's skill in reasoning and so forth) *rather than* to the pen, which would be thought an incidental part of the cognitive process. That is, we would not consider the pen to be part of one's cognitive agency.

But of course intuitions like this about the pens used by mathematicians to gain mathematical knowledge need not be in conflict with the extended cognition thesis, for the claim by proponents of this thesis is only that in certain cases the cognitive process might genuinely extend beyond the skin. Indeed, the overarching dialectic of the seminal paper on extended cognition by Andy Clark & David Chalmers (1998) is in terms of a 'why not' question which is encapsulated in the following passage:

If, as we confront some task, a part of the world functions as a process which, were it to go on in the head, we would have no hesitation in accepting as part of the cognitive process, then that part of the world is [...] part of the cognitive process.” (Clark & Chalmers 1998, 8)

Call this the *parity principle*. What informs this principle is the laudable egalitarian aim of avoiding irrelevant bias. After all, the philosophical tendency to treat cognitive processes as being exclusively located in the head (and thereby under the skin of the agent) is meant to be something that one is inexorably led to by reflecting on the very nature of cognitive processes, and not merely a self-fulfilling prejudice against extended cognition (such that any potential case of extended cognition is ruled-out by default because it extends beyond the skin of the agent). According to Clark & Chalmers, however, once one endorses this principle then one does seem to be committed to the possibility of extended cognition.¹⁷

Let’s begin by focussing on the case that Clark & Chalmers (1998) offer regarding Otto. Otto suffers from Alzheimer’s disease, and as a consequence he is gradually becoming aware that his memory is fading. In order to counter this, he starts carrying with him a notebook in which he records the kind of information that he requires on a day-to-day basis. On the traditional way of casting the matter, the notebook does not itself form part of Otto’s cognitive processes, even though it may be utilised by those processes. The challenge posed by the parity principle, however, is to explain why not. After all, if the information stored in the notebook were instead stored within Otto’s head in some way, and if that information were readily accessible to the same degree as the information in the notebook, then we would not hesitate to treat this information resource as part of Otto’s cognitive processes. The parity principle thus enjoins us to regard Otto’s notebook as part of the cognitive process, and so we get support for the extended cognition thesis.

Of course, not every case in which an agent uses a notebook in this way counts as a potential case of extended cognition. In particular, the Otto case is meant to exhibit two key features which set it apart from other cases to make it a good illustration of extended cognition. The first is that it is built into the case that the notebook is readily accessible and consistently available. Indeed, notice that we can strengthen the example on this score by having the notebook attached to Otto in some way—fixed to his arm, say—and designed in such a way that it is not hostage to circumstances (e.g., it can’t easily be damaged or get lost, and so on). If this condition were not met, then we would be far more inclined to treat the notebook as a mere instrument which assisted Otto on a particular occasion, rather than being

a genuine component of his cognitive processes.

The second feature that is critical to the Otto case is that Otto has self-consciously decided to ‘extend’ his cognitive process in this way: aware that his (non-extended) memory is failing, this is the means by which he ensures that he can still get access to the information that he requires. Moreover, he has a standing endorsement of the information in the notebook. It is not as if, for example, Otto treats each piece of information he comes across with suspicion, and so seeks additional grounds for accepting it (which is not to say that he won’t sometimes treat the contents of the notebook with suspicion, but then that is true of one’s (non-extended) memories anyway). If this were the case, then this would again incline us to treat the notebook as a mere instrument employed in a cognitive process rather than as a *bona fide* part of the cognitive process.

The interesting question for our current purposes is what, if anything, the Otto case tells us about the relationship between knowledge and cognitive ability, and thus about cognitive agency. First off, note that Otto’s cognitive success does seem to be sufficiently creditable to his cognitive agency for him to satisfy (COGA_{WEAK}), and there are thus no grounds on this score for denying him knowledge—indeed, intuitively Otto has knowledge of what he believes. But if we regard Otto’s cognitive processes as extended then it follows that Otto’s cognitive agency extends beyond his skin as well. It is this claim that we need to explore further.

4. EXTENDED COGNITION AND COGNITIVE ABILITY

The first thing to note about the Otto case is how Otto’s acquisition of the notebook, and his systematic use of it, represents a great deal of epistemic virtue on his part. A lesser cognitive agent—i.e., one who was less interested in gaining and retaining true beliefs about his environment—would have acquiesced in the loss of his (non-extended) memory and so accepted the epistemic consequences. Moreover, notice that the way in which Otto employs the notebook also reflects his epistemic virtue. An agent less concerned with epistemic goods would not, for example, go to the lengths that Otto goes to in order to ensure that this information resource is readily available to him but really would just use this notebook as a mere incidental aid to his cognition.

More generally, it seems that what Otto is doing in employing the notebook in this way is—arguably, at least—integrating this information resource into his cognitive character in a relevant way for the extended cognitive process that results to genuinely count as one of his cognitive abilities. Recall that we noted above that Temp and Alvin could integrate their reliable belief-forming processes into their cognitive character, and thereby be in a position to acquire knowledge through these processes, by coming to know both that the target belief-forming process is reliable and what the source of this reliability was. This is just what Otto has done, however, for while his (non-extended) memory is failing and so cannot be trusted, he knows that he can generally trust what the notebook tells him and why.

In contrast, if Otto had no awareness at all of the source of the reliability of his belief-forming processes, nor that it was reliable, then it is hard to see why we would now regard the true beliefs that he forms as a consequence as knowledge. Imagine, for example, that Otto is simply fitted with a device which provides him with reliable information about his environment and he unquestioningly consults it when necessary while never questioning the source of this information or its epistemic pedigree.¹⁸ It remains true that this is a reliable belief-forming process, but what no longer seems to be the case is that this is a belief-forming process which can generate knowledge. In particular, any true belief so formed seems to conflict with the ability intuition, as encoded in (COGA_{WEAK}). This is because we would not deem Otto's cognitive success as being to any significant degree creditable to his cognitive agency, but rather treat it as creditable to some feature external to his cognitive agency (i.e., the source of the reliability of the device in question).

The moral seems to be that even cognitive processes that extend outside the skin of agent can count as part of one's cognitive agency just so long as they are appropriately integrated within one's cognitive character. This raises a number of interesting questions.

To begin with, this is a good juncture at which to re-evaluate an issue we raised earlier regarding what is at minimum required for a reliable belief-forming process to become integrated within one's cognitive character. As we noted above, coming to know both that the process is reliability and the source of this reliability will suffice in this regard, but this still leaves open the possibility that meeting a weaker requirement might sometimes suffice as well.

One factor that seems salient in this regard is whether the source of the extended cognition was always present or whether it was added at a later juncture. In the Otto case we

clearly have an example of the latter, but suppose we imagine a variant of this case where the subject—let's call him 'Tempo'—is fitted from birth with a highly reliable device which records the ambient temperature and Tempo grows up in a culture where it is taken for granted that one consults one's temperature-recording device in order to form beliefs about the ambient temperature. Interestingly, in a case like this it seems entirely unnecessary for Tempo to know that this is a reliable belief-forming process or what the source of the reliability is before he can gain knowledge via this process. For imagine that Tempo is a young child who has never even considered these questions. Wouldn't we nonetheless straightforwardly regard him as gaining knowledge via this belief-forming process? Moreover, wouldn't we regard Tempo's cognitive success as being to a significant degree creditable to his cognitive agency, such that he satisfies (COGA_{WEAK})? If that's right, however, then it seems that we do intuitively regard this reliable belief-forming process as being sufficiently integrated into Tempo's cognitive character to count as one of his cognitive abilities.

The intuition being elicited here is often cited in support of externalist theories of knowledge. According to such theories, one can have knowledge even while lacking good reflectively accessible grounds in favour of one's belief in the known proposition. Externalists often point out that our innate reliable belief-forming processes, so long as they are functioning in an environment for which they are suited at any rate, can generate knowledge even if the agent concerned lacks any good reflectively accessible grounds in support of the beliefs so formed—for example, even if the agent lacks any good reflectively accessible grounds for thinking that the faculties in question are reliable.¹⁹

Imagine, for example, a child employing her reliable perceptual faculties in such a way that she gains a true belief in an environment which is suited to the operation of these faculties. Clearly, we would not expect the child to have any meta-beliefs about the source of the reliability of her belief-forming processes, and hence she is not in a position to have any (meta-)knowledge in this regard. But wouldn't we regard her cognitive success when employing her reliable perceptual faculties as properly creditable to her cognitive agency nonetheless? Indeed, wouldn't we straightforwardly treat this cognitive success as an instance of knowledge? What the Tempo case seems to illustrate, and this is in keeping with Clark & Chalmers's parity principle, is that what is important in this respect is only that the cognitive abilities are present from the off—and are thus in this sense 'innate' (henceforth we will use

this term in this loose way)—and not whether they are natural rather than artificially added to the agent.²⁰

When it comes to one's innate reliable belief-forming processes, then, whether they are natural or artificially added, it seems that we intuitively relax what is needed for cognitive integration, for in these cases it is not required that the agent concerned know that these processes are reliable or why. It is an interesting question why this should be. If we allow innate reliable belief-forming processes to be knowledge-conducive even though the agent lacks knowledge of the source of this reliability, then why would we demand that Otto should have this knowledge before we treat him as a knower (or, for that matter, Alvin)? On reflection, however, it becomes clear that there is a natural explanation of this phenomenon.

To begin with, note that in the Otto case there is a *change* in the agent's cognitive processes, whereas in the Tempio case there is no change in this regard. The import of this is that such change cries out for the agent to take a reflective stance on the epistemic standing of this change, something that is not at all present in the Tempio case. For note that if we were to imagine Tempio being fitted with this device at a later stage, then we would require him to form a view as to the reliability of this process, and the source of this reliability, before we would regard the process as knowledge-conducive.

Moreover, notice that the Tempio case does involve genuine cognitive integration. In order to see this point, suppose we re-describe the case so that the beliefs produced by the attached device sometimes conflict with the beliefs produced by his other faculties. For instance, suppose that sometimes this device tells him that the temperature is quite cold and yet he feels quite warm (and, say, observes a fire in the hearth). In treating Tempio as gaining knowledge of the temperature in the room as a result of being fitted with this device, we are surely supposing him to be an agent who would respond to such discrepancies were they to emerge. If this is not the case—if Tempio is happy to live with such apparent inconsistencies in his beliefs—then it ceases to be far so obvious that he does gain knowledge from this device. And notice the reason for this. For what Tempio demonstrates when he does respond appropriately to these apparent inconsistencies in his beliefs is that there is a weak kind of cognitive integration in play as regards even the belief-forming process involving this appended device. In contrast, if even this weak form of cognitive integration were not present then I think we would be reluctant to ascribe the relevant knowledge to Tempio.

So when it comes to changes in one's cognitive processes we tend to demand a form

of cognitive integration which is quite austere, while when it comes to innate cognitive processes, even where this involves extended cognitive processes, we tend to be more liberal and allow a much weaker sense of cognitive integration to be compatible with knowledge. We should expect that there is a continuum in play here, such that between these two extremes there is a range of cases which elicit mixed, or at least uncertain, intuitions.

Suppose, for example, that there was a change in the agent's cognitive processes, but that the agent was somehow shielded from being made aware of this, in such a way that she was not epistemically culpable (i.e., it wasn't simply that she took no interest in what was happening to her). Consider, for example, a variant of the Tempo case where the device was added at a later juncture. Perhaps, for instance, Tempo* comes out of a coma with this device fitted and is somehow psychologically manipulated so that he comes to trust the information from this device while continuing to non-culpably be unaware that this device has been artificially implanted in him. Can Tempo* gain knowledge by using the reliable belief-forming process that makes use of this device? To begin with at least, I take it that intuition is not on Tempo*'s side. After all, wouldn't we simply credit his cognitive success to the machinations of those who fitted him with this device rather than (to any significant degree anyway) to his cognitive agency? If so, then he does not count as exhibiting a cognitive ability and so does not qualify for knowledge.

Interestingly, however, as time goes on this intuition lessens. Imagine that Tempo* has had this device fitted for ten years, say. Does he still not qualify as having knowledge in the relevant respects? Clearly, the matter is now less certain. (And note that what goes here applies with equal force to the analogous variant of the Alvin case where the lesion in question is long-standing). But notice that we can explain what is going on here; why our intuitions are becoming less definitive. For assuming that Tempo* is an otherwise epistemically good fellow—i.e., even despite what cognitive impairment was brought about in order to unknowingly set him up with this extended belief-forming process in the first place—we would expect that he is displaying the kind of weak cognitive integration that we noted above was primarily relevant when it comes to innate reliable belief-forming processes. That is, there is now a track-record of beliefs formed via this process which have generally cohered with the beliefs formed via Tempo*'s cognitive abilities (and if they hadn't cohered, we may suppose, then Tempo* would have spotted this and responded accordingly). With this point in mind it is not surprising that our intuitions are conflicted on this score. For while

we would ordinarily demand a stronger form of cognitive integration in this case, that there is this weak form of cognitive integration present surely has an influence on our intuitions in this regard. We can thus account for why it strikes us a penumbral case.

So once we understand the ability intuition on knowledge along the lines set out by (COGA_{WEAK}), then there seems a straightforward way of accommodating the extended cognition thesis within contemporary thinking about knowledge. Let me close by noting how (COGA_{WEAK}) is in fact more conducive to the extended cognition thesis than either (COGA_{STRONG}) or (COGA_{STRONG*}). The reason for this is that once we move towards interpreting the ability intuition such that it demands that the agent's cognitive success be primarily creditable to the agent's cognitive character, then this might lessen our willingness to allow certain putative cases of extended cognition.

Think again about the Otto case, for example. As we described the case above, Otto is solely responsible for setting-up and maintaining this new cognitive process, but clearly we can imagine variations of this case where this isn't so. Suppose, for instance, that Otto's wife helps him get set-up with this device, and plays an important role in ensuring that it functions effectively (in practice, of course, this would be quite likely). Adding this feature of the case does not prevent Otto from continuing to gain knowledge from this belief-forming process by the lights of (COGA_{WEAK}), since it will still be true that Otto's cognitive success is significantly creditable to his cognitive agency, even despite the involvement of a third-party in this regard. When it comes to either (COGA_{STRONG}) and (COGA_{STRONG*}), however, matters are not so clear. For given the involvement of a third-party is it still right to suppose that Otto's cognitive success is *primarily* creditable to his cognitive agency? This is hard to judge.

More generally, an advantage of (COGA_{WEAK}) is that it can accommodate the thought that sometimes—e.g., when the agent is in a very epistemically friendly environment—very little cognitive ability on the part of the subject can suffice for knowledge, and this is advantageous when it comes to accommodating cases of extended cognition. For we might suppose that part of what drives support for at least some cases of extended cognition is that the agent is in the relevant respects in an epistemically friendly environment such that although she is being assisted in various ways in attaining her cognitive goals—e.g., by receiving help from other agents within that environment—she is nonetheless exhibiting a sufficient degree of cognitive ability to count as a knower. Given that paradigm cases of extended cognition typically involve an agent adapting her current belief-forming processes,

we might expect that receiving help from a third-party is in practice quite a normal feature of cases of extended cognition (such ‘adaption’ is hard to do completely solo, we might suppose). But if that’s right, then it is inevitable that strong accounts of the relationship between cognitive success and cognitive agency are less suited to accommodating the extended cognition thesis.²¹ It is fortunate, then, that such accounts are independently problematic, as we demonstrated above.

5. CONCLUDING REMARKS

So far from being in conflict with contemporary thinking about knowledge, it seems that the extended cognition thesis may actually constitute a natural way of thinking about cognitive agency. In order to see this, however, we first needed to clarify what is involved in exhibiting a cognitive ability and thereby offer the right account of the relationship between knowledge and cognitive agency. As we have seen, there are independent reasons for favouring a relatively weak account in this regard. Interestingly, however, this weak account of cognitive agency is well-placed to accommodate cases of extended cognition. We thus have strong grounds for supposing that contemporary epistemology, through its commitment to the ability intuition, should incorporate the extended cognition thesis into the theory of knowledge.²²

REFERENCES

- Adams, F., & Aizawa, K. (2001). ‘The Bounds of Cognition’, *Philosophical Psychology* 14, 43-64.
- (2008). *The Bounds of Cognition*, Blackwell, Oxford.
- Chisholm, R. (1977). *Theory of Knowledge* (2nd Ed.), Prentice-Hall, Englewood Cliffs, NJ.
- Clark, A. (1998). *Being There: Putting Brain, Body and World Together Again*, MIT Press, Cambridge, MA.
- (2007). ‘Curing Cognitive Hiccups: A Defense of the Extended Mind’, *Journal of Philosophy* 104, 163-92.
- (2008). ‘Memento’s Revenge: The Extended Mind, Extended’, *The Extended Mind*, (ed.) R. Menary, Ashgate Press, Aldershot.
- Clark, A., & Chalmers, D. (1998). ‘The Extended Mind’, *Analysis* 58, 7-19.
- Cowart, M. (2006). ‘Embodied Cognition’, *Internet Encyclopaedia of Philosophy*, (eds.) B. Dowden & J. Fieser, <http://www.iep.utm.edu/e/embodcog.htm>.
- Gertler, B. (2007). ‘Overextending the Mind?’, *Arguing about Mind*, (eds.) B. Gertler & L.

- Shapiro, 192-206, Routledge, London.
- Gettier, E. (1963). 'Is Justified True Belief Knowledge?', *Analysis* 23, 121-3.
- Goldman, A. (1976). 'Discrimination and Perceptual Knowledge'. *The Journal of Philosophy* 73, 771-91.
- (1986). *Epistemology and Cognition*, Harvard University Press, Cambridge, MA.
- Greco, J. (1999). 'Agent Reliabilism', *Philosophical Perspectives* 13, 273-96.
- (2000). *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*, Cambridge University Press, Cambridge.
- (2002). 'Knowledge as Credit for True Belief', *Intellectual Virtue: Perspectives from Ethics and Epistemology*, (eds.) M. DePaul & L. Zagzebski, Oxford University Press, Oxford.
- (2007). 'The Nature of Ability and the Purpose of Knowledge', *Philosophical Issues* 17, 57-69.
- (2008). 'What's Wrong With Contextualism?', *The Philosophical Quarterly* 58, 416-36.
- (2009). 'The Value Problem', *Epistemic Value*, (eds.) A. Haddock, A. Millar & D. H. Pritchard, Oxford University Press, Oxford.
- Haddock, A., Millar, A., & Pritchard, D. H. (2009). *The Nature and Value of Knowledge: Three Investigations*, Oxford University Press, Oxford.
- Lackey, J. (2007). 'Why We Don't Deserve Credit for Everything We Know', *Synthese* 158, 345-61.
- Lehrer, K., & Cohen, S. (1983). 'Justification, Truth, and Coherence', *Synthese* 55, 191-207.
- Menary, R. (2006). 'Attacking the Bounds of Cognition', *Philosophical Psychology* 19, 329-44.
- (2007). *Cognitive Integration: Mind and Cognition Unbounded*, Palgrave Macmillan, London.
- (ed.) (2008). *The Extended Mind*, Ashgate Press, Aldershot.
- Plantinga, A. (1993a). *Warrant and Proper Function*, Oxford University Press, Oxford.
- (1993b). *Warrant: The Current Debate*, Oxford University Press, Oxford.
- Pritchard, D. H. (2005). *Epistemic Luck*, Oxford University Press, Oxford.
- (2008a). 'Greco on Knowledge: Virtues, Contexts, Achievements', *The Philosophical Quarterly* 58, 437-47.
- (2008b). 'Sensitivity, Safety, and Anti-Luck Epistemology', *The Oxford Handbook of Scepticism*, (ed.) J. Greco, 437-55, Oxford University Press, Oxford.
- (2008c). 'Virtue Epistemology and Epistemic Luck, Revisited', *Metaphilosophy* 39, 66-88.
- (2009). 'Anti-Luck Virtue Epistemology', *manuscript*.
- (Forthcoming). 'Apt Performance and Epistemic Value', *Philosophical Studies*.
- Rowlands, M. (1999). *The Body in Mind: Understanding Cognitive Processes*, Cambridge University Press, Cambridge.
- (2009). 'Extended Cognition and the Mark of the Cognitive', *Philosophical Psychology* 22, 1-19.
- Rupert, R. (2004). 'Challenges to the Hypothesis of Extended Cognition', *Journal of Philosophy* 101, 389-428.
- Sosa, E. (1988). 'Beyond Scepticism, to the Best of our Knowledge', *Mind* 97, 153-89.
- (1991). *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge University Press, Cambridge.

- (2007). *Apt Belief and Reflective Knowledge*, Oxford University Press, Oxford.
- Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge University Press, Cambridge.
- (1999). ‘What is Knowledge?’, *Epistemology*, (eds.) J. Greco & E. Sosa, 92-116, Blackwell, Oxford.
- Vaesen, K. (2009). ‘Knowledge Without Credit’, *manuscript*.

NOTES

- ¹ The *locus classicus* for the extended cognition thesis is Clark & Chalmers (1998). For some of the other key defences of extended cognition, see Clark (1998; 2006; 2007) and Rowlands (1999). See also Menary (2007). For two recent critical discussions of this thesis, see Adams & Aizawa (2008; cf. Adams & Aizawa 2001) and Rupert (2004; cf. Rupert 2009). See also Gertler (2007), though note that her focus is more on the related extended mind thesis than the extended cognition thesis. For a useful survey of the literature on extended cognition, see Cowart (2006). See Menary (2008) for a recent volume of articles on the extended mind thesis.
- ² For further discussion of safety and sensitivity, see Pritchard (2008b).
- ³ The thesis that the envatted agent’s beliefs could nonetheless enjoy a high degree of epistemic support, at least to the extent that they are internalistically justified, is known as the ‘new evil genius intuition’. See Lehrer & Cohen (1983).
- ⁴ For further discussion of the Temp case, see Haddock, Millar & Pritchard (2009, ch. 3).
- ⁵ The Alvin case is adapted from one offered by Plantinga (1993b).
- ⁶ Although there are various conceptions of cognitive character in the literature—some of them quite demanding, some of them very undemanding—we don’t need to take a stance on this issue here, as I explain why below (see endnote 7). For a very undemanding conception of cognitive character, one cast along broadly reliabilist lines, see Greco (e.g., 2000). For a very demanding conception of cognitive character, one cast along neo-Aristotelian lines, see Zagzebski (1996).
- ⁷ It is for this reason that we don’t need to engage with the detailed debate about the nature of cognitive character here (see endnote 6), since we can bypass this discussion for our purposes by simply focussing on the question of whether we would treat the agent’s cognitive success as appropriately creditable to her cognitive agency. Where the agent’s cognitive success passes this test, then we will treat it as being due to a reliable belief-forming process that is appropriately integrated within the agent’s cognitive character, and thus due to cognitive ability; where the agent’s cognitive success fails to pass this test, then we will not treat it as being due to cognitive ability. In this way we can gain a good enough grip on what constitutes cognitive character without having to engage with the wider epistemological discussion of this notion, and we can thereby simplify our discussion.
- ⁸ It might also be necessary to add that this new knowledge should cohere with the other relevant beliefs that the agent holds, though in most cases—and possibly all cases—this will be entailed by the acquisition of the new knowledge anyway, thereby making the addition of this further clause redundant. I comment further on this point below—see endnote 9.
- ⁹ I noted above in endnote 8 that one might be tempted to add a coherence clause to the strong account of cognitive integration on offer, such that the new knowledge that the agent acquires coheres with other relevant beliefs he holds. In most cases, however, if not all of them, this clause will be redundant, as the Temp and Alvin cases illustrate. After all, it is hard to see how Temp or Alvin could acquire the new knowledge in question if the new beliefs they form do not cohere with other relevant beliefs that they hold. In the Temp case, for example, if Temp has, say, independent reason to doubt the reliability of the thermostat that the helper uses, then won’t this undermine his knowledge that this belief-forming process is reliable? Still, I leave it an open question whether the addition of this clause may be necessary to deal with certain cases.
- ¹⁰ Arguably, one could add Plantinga (1993a) to this list. Note that in earlier work Greco (e.g., 1999; 2000) advanced a conception of the relationship between knowledge and cognitive ability which was much more in keeping with (COGA_{WEAK}).
- ¹¹ See Greco (2009) for a clear statement of the theoretical advantages of this account of knowledge, particularly with regard to its putative application to the topic of epistemic value.
- ¹² This example is adapted from one offered by Chisholm (1977, 105).
- ¹³ This case was originally offered by Goldman (1976), although he credits the example to Carl Ginet.

¹⁴ This example is due to Lackey (2007), though note that it was originally designed to illustrate a slightly different point.

¹⁵ Alternatively, one might regard this as a case of distributed cognition such that the cognitive system comprising of both Jenny and her informant should be regarded as primarily creditable for Jenny's cognitive success. Although I believe that there are cases of distributed cognition, I take no stand here on whether this case falls into this category since the important point for our purposes is just that Jenny's cognitive success is not primarily creditable to her.

¹⁶ For further discussion of weak versus strong accounts of cognitive agency, see Haddock, Millar & Pritchard (2008, ch. 2) and Pritchard (2008a; 2008c; 2009; *forthcoming*).

¹⁷ There are worries about the parity principle, not least of which that it is side-steps other apparently defining marks of the cognitive, such as that it involves some sort of non-derived, and not merely (temporally) immediate, access to information. For some recent fine-grained discussions of the notion of the cognitive on this score, see Adams & Aizawa (2001), Clark (2007; 2008), Menary (2006; 2007), Rowlands (2009), and Rupert (2004).

¹⁸ Note that I am simply talking here of a device which supplies Otto with reliable information, rather than a 'memorial' device in which Otto himself can record information for his later self. The problem with thinking of the device along memorial lines is that one then needs to tackle the vexed question of how Otto can both record information in this notebook and retrieve information from it without having any view of the epistemic pedigree of the information. (After all, if his memory is so bad that he doesn't remember entering this information, then in what sense is he a believer at all? But if he has any memory of entering the information, then that would surely have an effect on the epistemic pedigree he assigns to the information so retrieved).

¹⁹ The *locus classicus* for discussions of epistemic externalism is Goldman (1986). See also Plantinga (1993a; 1993b).

²⁰ Note that in talking of the reliable belief-forming processes as being present "from the off" I am glossing over an important complication in this regard, which is those belief-forming processes that we are primed for 'from the off', but which are only activated at a later point, if at all. The paradigm example in this regard is the (apocryphal) example of the so-called 'chicken-sexer' cognitive ability much beloved by epistemic externalists. For while agents are not born with chicken-sexing abilities, the idea runs that they are primed for such abilities, should they encounter suitable further stimuli (in this case, spending plenty of time around chicks). For our purposes I think we can safely set reliable belief-forming processes of this type to one side, independently interesting though they are. For further discussion of chicken-sexer cases, and their relevance to the externalism/internalism debate in epistemology, see Pritchard (2005, *passim*).

²¹ See Vaesen (2009) for a persuasive argument against a strong conception of the ability principle by appeal to the extended cognition thesis.

²² I am grateful to Evan Butts, Andy Clark, Luciano Floridi, Sandy Goldberg, Jesper Kallestrup, Orestis Palermos, Tom Roberts, Ernie Sosa and Krist Vaesen for helpful discussions on this topic. My research into this area has been supported by the award of a Philip Leverhulme Prize.