

EXTENDED BELIEF AND EXTENDED KNOWLEDGE

Åsa Wikforss
Stockholm University

The hypothesis of extended cognition comprises two distinct claims: First, there is the claim that in many cases the external world aids our cognitive processing in such a way that it would be artificial, and scientifically unfruitful, to insist that only that part of the process which is within the skull is properly speaking cognitive. Let us call this the thesis of extended *processing*. Second, there is the thesis of extended *states*: the claim that ordinary folk psychological states, such as belief and desire, extend beyond the skin and skull. As Clark and Chalmers note, the second thesis is the more radical one (1998: 12). Much of what qualifies as cognitive processing is subpersonal and as such it is not obvious that it belongs to the mind in the first place. Why, then, should the process become cognitive only at that step in the causal chain that it enters the skull? Folk psychological states, by contrast, are personal level states that stand at the center of our ordinary conception of ourselves as minded creatures, capable of reasoning and action. If these states extend beyond the skull and skin, we need to reconsider the traditional conception of these as seated in the human body—an idea that has informed not only Cartesian conceptions of the mind but also (nearly) all 20th century theories of the mind.

The extended states thesis also has implications for related fields within philosophy, such as epistemology and moral philosophy, given the central role played by beliefs and desires in these fields. In particular, if ordinary personal level beliefs extend beyond the body, the question naturally arises if ordinary, personal-level *knowledge* similarly extends. Indeed, the step from extending belief to extending knowledge would seem very short: Belief was traditionally conceived of as the 'inner', subjective component of knowledge, whereas the additional components were conceived of as objective and external—truth, certainly, but also (on many accounts) justification. If the

subjective component of knowledge extends there should be little obstacle to extending the more objective components, and hence little obstacle to extending knowledge.

However, the prospect of extending knowledge raises a well-known worry, the worry that knowledge will *overextend*, or, as it is sometimes said, that there will be a ‘cognitive bloat’ (Clark 2008a).¹ In their 1998 paper, Clark and Chalmers did not merely argue for the *principled* possibility of external vehicles of belief, they argued for the very strong claim that belief very easily extends, and suggested that ordinary external information carrying devices (such as notebooks and computers), under conditions that are relatively easily met, can function as the vehicle of belief. But if this is so, and if the extension of belief brings with it an extension of knowledge, then it would seem as if there is a possible explosion in ordinary personal level knowledge: If you carry your iPhone with you in the right way, what is there to prevent the conclusion that you know all the information on the iPhone? What is there to prevent the conclusion that the ignorant student who has downloaded the textbook on ancient philosophy suddenly becomes an expert on the topic? Of course, it is part and parcel of the extended mind thesis that belief and knowledge does extend in rather surprising ways. However, no one would take it to be a virtue of the theory that the extension is unchecked, clashing with widely shared intuitions about what does and what does not count as a subject knowing that *p*.²

In response to this problem various strategies are possible. One strategy is to reject the extended mind thesis in its entirety. For example, it might be insisted that all states that properly belongs with the mind are occurrent, or conscious, and that consciousness does not extend.³ Or one might simply insist that all mental states need be neurologically realized.⁴ A second strategy is to accept the extended mind thesis but to argue that even if *belief* extends along the lines suggested by Clark and Chalmers, *knowledge* does not since the conditions necessary for knowledge are not met by extended beliefs—at least not in the problem cases, where we have a strong intuition that the subject does not know that *p*. For instance, it has been argued that knowledge requires a form of cognitive responsibility that is not present in the problem cases.⁵ On this proposal, there is a very close connection between extended cognition and epistemology; indeed, the theory of knowledge is to be judged in part by its capacity to solve the problem of knowledge bloat.

In this paper I shall attempt a third strategy. I shall argue that the fundamental problem arises at the level of belief, not at the level of knowledge: We prevent bloating of knowledge by preventing the bloating of belief. To do so we need not deny the *principled* possibility of extended beliefs, we just need to take seriously Clark and Chalmers’ suggestion that what is distinctive of belief is its *role in folk psychology*; the idea that in order for a given piece of external information to be an instance of an extended belief, it need play the role of belief in folk psychological explanations. This requirement, I argue,

is much more demanding than what is captured by Clark and Chalmers four conditions on belief, the so called ‘trust and glue’ conditions. Taking the folk psychological role seriously, therefore, allows us to draw a distinction of the required sort, between what does and what does not belong to the subject’s system of beliefs. Moreover, I suggest, the conditions under which belief extends are also conditions under which knowledge extends (assuming the beliefs are true)—or, rather, the features that make the extended state a belief are features that make it a good *candidate* for knowledge, even if they are not sufficient for knowledge. As a consequence, there is no need to adjust the theory of knowledge to prevent knowledge from overextending.

The paper is divided into four parts. Part 1 sets the stage by spelling out the details of the extended belief thesis. Part 2 presents the problem of ‘cognitive bloat’ and discusses the attempt to solve it by appealing to the notion of cognitive responsibility, a notion central to ability and virtue theories of knowledge. Part 3 discusses the role of belief in folk psychology, arguing that this prevents ‘belief bloating’, and part 4 spells out how this account of belief allows us to solve the problem of knowledge bloat even within the confines of a traditional, evidentialist conception of justification and knowledge.

1. The Extended Belief Thesis

Clark and Chalmers argue for the twin theses of extended processing and extended mind in a parallel fashion: In both cases the argument turns on the central role played by a given external feature in the subject’s cognitive life, in her reasoning and actions. And in both cases they appeal to the so called ‘Parity Principle’, the idea that the mere fact that a feature is external to a subject’s skin and skull, does not in itself disqualify the feature from being a genuine part of the subject’s cognitive processing or her mind: “If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world *is* (so we claim) part of the cognitive process. Cognitive processes ain’t (all) in the head!” (Clark and Chalmers 1998: 8).

Nevertheless, given that the one thesis concerns extended *processes*, the other extended *states* there are reasons to treat the theses separately. In the case of cognitive processing, the argument depends on the idea that there is no natural cut off point between the contributions of the internal and the external aspects. Consequently, the claim is not that the process as such is located outside the subject, but that the internal and external parts of the process are inextricably intertwined. As Clark has later put it, in these cases cognitive processing “bleeds into the world” (2008a: 70). In the case of belief, by contrast, the thesis is best understood as the claim that standing

beliefs may, indeed, be located outside the subject. One way to spell out the difference between the two theses is in terms of *realization*. Whereas the extended processing thesis implies that the cognitive process is realized in internal as well as external features, it is ‘spread out’ across the boundary of the skin and skull, the extended mind thesis is precisely the thesis that the psychological state is realized in the external feature, for instance in the notebook information. No doubt, the argument supporting the latter claim depends on the idea that the external feature interacts with the subject in a special way; it is only in virtue of the fact that the information plays the right causal role in the subject’s life, her reasoning and actions, that it can be said to realize (some of) her standing beliefs. However, the upshot is not that there is no natural cut-off point between the internal and external contributions but, rather, that the psychological state is realized in a feature which is wholly external to the subject’s skin and skull.⁶

This is quite clear in Clark and Chalmers’ well-known discussion of Otto, the Alzheimer patient who writes down important information in a notebook that he carries with him at all times. Thus, Otto notes down that MoMA is on 53rd Street, and when he desires to go to MoMA he simply looks this information up and proceeds to 53rd Street. Since this externally stored information plays the same causal role in Otto’s life as internally stored information in a normal subject, Clark and Chalmers argue, we should conclude that Otto believed that MoMA was on 53rd Street even before consulting the notebook (1998: 13). Clark and Chalmers recognize that there are many fine-grained differences between the role played by the notebook information and the role played by internally stored information. For instance, there is a difference in the phenomenology of retrieval: It is immediate in the case of Inga, but mediated by a visual experience in the case of Otto. However, they argue that these differences are all *shallow* and that to focus on them would be to miss that the notebook information plays just the role that belief does in guiding most people’s lives. That is to say, the differences are all non-essential to belief and do not prevent the conclusion that the information in the notebook realizes (some of) Otto’s standing beliefs. Consequently, Otto’s belief about the location of MoMA “is simply not in the head” (1998: 14).

This difference between the thesis of extended processing and the thesis of extended states also means that the two theses rest on rather different assumptions. In the case of cognitive processing the argument does not seem to presuppose any specific assumptions about the nature of mind, and the Parity Principle is relatively neutral: If we are concerned with a process it is obviously true that the causal role is decisive (this just seems to be what a process *is*, any process) and the claim that it is unmotivated to cut the process off at the limit of the skull seems *prima facie* plausible. This is a reflection of the fact, noted above, that it is not obvious that all cognitive processing belongs to the mind in the ordinary sense. However, in the case of mental states such as belief and desire more controversial assumptions

are needed. To begin, all versions of materialism according to which the neurological base of a psychological state is essential must be rejected, such as the identity theory. In relation to such theories, the Parity Principle is not neutral but obviously false. Indeed, applied to psychological states the principle would seem to be true only on the assumption that what is essential to these states is the *role* played by them. That is, it depends on a functionalist conception of the mind, or (at least) a functionalist conception of belief.⁷ In fact, although Clark and Chalmers do not explicitly mention functionalism, their arguments for the extended belief thesis implicitly invoke it. Thus, in an oft-cited passage, they write that “what makes some information count as a belief is the role it plays, and there is no reason why the relevant role can be played only from inside the body” (1998: 14).⁸

The extended state thesis therefore depends on a functionalist conception of psychological states. A central motivation behind functionalist theories of the mind is the idea that mental states are multiply realizable, allowing for the possibility that creatures with a different physiology could have minds just like ours. Originally, when functionalism was first introduced, people were talking about the possibility of non-neural realizers, such as silicon chips, etc. but it is of course equally possible that the realizers are outside the skull of the agent. Given a functionalist starting point, therefore, the principled possibility of extended states follows immediately; in this sense, the thesis is fairly trivial.⁹ However, the thesis that is of relevance here, is not that extended beliefs are *metaphysically possible*, but the very strong thesis that extended beliefs are possible here and now, that our standing beliefs may be realized in ordinary notebooks and iPhones. This is the claim that raises the worry that belief, and knowledge, overextends. And to support it, is not sufficient to appeal to the metaphysical possibility of external realizers, it must also be argued, as Clark and Chalmers do, that the role distinctive of belief is in fact played by the notebook information.

2. Overextending Knowledge

Although Clark and Chalmers do not claim to give an exhaustive account of the role of belief, they do suggest that the notebook information shares four central characteristics with Inga’s internally stored information and that this warrants the conclusion that (some of) Otto’s standing beliefs are realized in his notebook (1998: 17). These are the so called ‘trust and glue’ conditions:

- (i) the notebook is a constant in Otto’s life
- (ii) the information in the notebook is directly available without difficulty

- (iii) upon retrieving information from the notebook Otto automatically endorses it
- (iv) the information in the notebook has been consciously endorsed in the past

Let us grant, for the moment, that these conditions are all necessary for a piece of information to qualify as a belief. The question is whether they are sufficient: Do these conditions allow us to separate external information that can be properly said to be part of the subject's set of beliefs, and external information which is not part of this set? Although Clark and Chalmers's argument about Otto presupposes that the conditions are sufficient they have later expressed some uncertainty on this issue.¹⁰ It is easy to see why. If I carry my iPhone around the way Otto does his notebook, and I have downloaded (and accepted) everything there is to know about the Icelandic fishing industry, does it really follow that this information constitutes my standing beliefs? For instance, does it really follow that I have the standing belief that in 2011 the export production of marine products amounted to ISK 252 billion and increased in value by 14.4% from previous year?

Once we take the step from extended belief to extended knowledge the worry about 'bloat' becomes more acute. If the downloaded information belongs with my standing beliefs does it also follow that I now *know* everything there is to know about the Icelandic fishing industry? Have I suddenly become an expert in the topic? The 'bloating' worry arises with respect to cognitive processing as well, since we do want to maintain some distinction between that which properly belongs to my cognitive system and what belongs with causal background conditions. However, when it comes to personal-level, propositional knowledge overextending seems even more problematic given the central role of this kind of knowledge in our lives and in the community. Katalin Farkas mentions the case of Lotte who has downloaded a 37 volume of the history of Europe with a quick search function from a source she trusts completely (Farkas 2012). Does Lotte thereby acquire all these beliefs and become an expert on the history of England? Surely, Farkas argues, we do not want to obliterate the distinction between Lotte and an impressively erudite scholar of English history.

It might be thought that condition (iv) prevents this, since Lotte has not consciously endorsed everything in the 37 volumes. However, as Clark & Chalmers themselves suggest, condition (iv) is dubious, since there is the possibility that one acquires beliefs through subliminal perception or in other ways that do not involve conscious endorsement. Indeed, it would seem as if most of our beliefs are such that we have not consciously endorsed them, since most of our beliefs are simple inferences from other beliefs that we have consciously endorsed.¹¹ Still, Clark hesitates giving up condition (iv) since, he suggests, "to drop this requirement opens the floodgates to what many would regard as an unwelcome explosion of potential dispositional beliefs" (2008a:

80). This may well be true, but the condition cannot be maintained just to save the extended belief thesis.¹² If condition (iv) is not a necessary condition on standing beliefs, then suggesting that it is a necessary condition on the extended 'states' is just to say that the latter are not beliefs. Consequently, another strategy is required to prevent knowledge from over-extending.

An obvious strategy would be to question the move from belief to knowledge. After all, many of our (true) beliefs do not qualify as knowledge since they lack further conditions necessary for knowledge, such as justification. Perhaps, then, we can prevent the conclusion that Lotte knows everything in the 37-volume by appealing to one of these further conditions. Of course, there are theories of justification that do not seem well suited for the job. For instance, a purely reliabilist theory is insufficient since it is built into these examples of extended beliefs that all reliability conditions are met: Not only is the information in question true, it is reliably available to the subject and immediately endorsed. I will be a reliable source when it comes to the Icelandic fishing industry, as will Lotte when it comes to the history of England. If the theory of knowledge is going to help with the problem of over-extended knowledge, we need to go beyond reliabilism. One important alternative to purely reliabilist theories of knowledge is theories that stress the subject's cognitive *abilities* or virtues. Perhaps these theories are better suited to the task?

Duncan Pritchard (2010) has argued that there is a natural fit between ability theories of knowledge and extended cognition. There is a strong intuition, Pritchard argues, that a necessary condition on knowledge is that it is a product of a cognitive ability. For instance, a brain in a vat may reliably form true beliefs about her environment (as a result of the work of a super-computer) but since the reliability has nothing to do with the subject's cognitive abilities, Pritchard argues, her beliefs do not qualify as knowledge. To manifest a cognitive ability the reliable belief-forming process must be properly integrated within the cognitive character of the subject in such a way that her cognitive success can (to a significant degree) be credited to her *cognitive agency*—the subject must be able to take *cognitive responsibility* for her success. The idea that knowledge requires a cognitive ability, Pritchard argues, supports the possibility of extended knowledge. Thus, Otto takes cognitive responsibility for the notebook information in the required way: He has self-consciously decided to extend his cognitive process and his systematic use of the notebook represents a great deal of epistemic virtue on his part. He has a standing endorsement of the information in the notebook and makes serious efforts to ensure that the information is readily available to him. However, Pritchard argues, if Otto were merely fitted with a device that fed him reliable information about the environment, and had no awareness of the source of this reliability and made no efforts to confirm its epistemic qualities, then the true beliefs that he would form as a result would not qualify as knowledge. Consequently, on Pritchard's view,

knowledge extends but only when the beliefs can properly be seen as the result of the subject's cognitive agency.

A similar idea can be found in Roberts (2012). Roberts argues that the glue and trust conditions fail to distinguish the case where a subject uses an external device to be a reliable producer of representational outputs, from the case where the extended resource is properly integrated into the mental life of the subject. In order to secure this distinction, according to Roberts, we need to appeal to the notion of cognitive responsibility: There is knowledge in those cases where the subject is *responsible* for the cognitive activity, which means that the subject grasps and is sensitive to the relevant epistemic norms. This, Roberts argues, allows us to prevent cognitive bloat. The subject's cognitive activity must be the result of an implicit knowledge of the relevant norms governing the problem-solving behavior of the whole system. It is therefore not enough that her beliefs are reliable, the subject need also have internalized the relevant epistemic norms: "Although the subject's familiarity with the device may permit her to closely engage with the apparatus under appropriate circumstances, she does not bear responsibility for it when she does not grasp the normative significance of the output, or of the intermediate steps that were required to reach this product" (2012: 140).

However, there are problems with the appeal to cognitive responsibility and norm guidance. To begin, there is the traditional worry concerning doxastic voluntarism, the idea that we can choose what to believe. The talk of epistemic responsibility presupposes that we have a certain amount of control over what we believe, but there are powerful objections to this idea.¹³ For the most part belief formation is automatic, reflecting how we are built to respond to certain input. For instance, perceptual beliefs cannot be said to result from my cognitive abilities but are formed automatically in response to sensory input. When it comes to perceptual knowledge, it would seem, we are in the same position as the brain in the vat: We are fed information through a reliable mechanism that we have not chosen to endorse and cannot choose to disregard.¹⁴

Moreover, the claim that belief formation need be guided by norms implies a conception of belief formation that is much too cognitively demanding, since it requires both that the subject knows the relevant epistemic norms and that she has a meta-level access to her own first-order reasoning. In this respect, virtue epistemology runs into the same problems as access internalism, the idea that justification requires that the subject is able to access what justifies her first-order beliefs. Roberts' is aware of this problem, and grants that 'the intellectualist model' of cognitive rule-following is too demanding since it is not clear that ordinary speakers know much about any rules of reasoning. For this reason he suggests that the norm guidance is merely *implicit*, appealing to John Greco's idea that all that is required is that the subject 'countenances' the relevant norms (2012: 138). However, as Quine pointed out long ago, it is not clear how the talk of implicit

rule-following will allow us to preserve the distinction between being guided by a norm and merely acting in accordance with it. And unless this distinction can be upheld, the norms are nothing but an idle wheel: What matters is not whether there are norms involved, but merely whether the beliefs are reliable and supported by evidence.¹⁵

There is no doubt more to be said about this, but the difficulties noted here should indicate that it would be unfortunate if preventing cognitive bloat required falling back on the notion of cognitive responsibility. The question also arises whether the notion does the work required and allows us to exclude the cases we would like to exclude. Consider Lotte again. We can imagine that she is extremely conscientious. She goes to great lengths making sure that the encyclopedia she downloads is the best and most reliable source on the history of England (perhaps she goes around interviewing experts on the topic). And she makes sure that the information is with her at all times, and she has a standing endorsement of it. Arguably, this shows that Lotte is epistemically responsible, and that her cognitive success with respect to English history does depend on her cognitive agency. She is certainly not in the situation of Pritchard's subject who is fitted with an external device that he unquestioningly consults. What, then, is there to prevent the conclusion that Lotte does know everything that is in the encyclopedia?

I therefore do not think that the best response to the problem of cognitive bloat goes via the ability/virtue theories of knowledge. Instead, I shall suggest another strategy. The best way to prevent cognitive bloat, I want to suggest, is to prevent a bloating of belief. The problem with cases such as Lotte's is that the subject in these cases cannot be said to *believe* the external information.

3. Folk Psychology and the Role of Belief

It is important, at the outset, to keep in mind that there is an important difference between *belief* and *information*.¹⁶ This tends to get obscured in the debate, since it is common to speak of belief in terms of information. Clark and Chalmers, for instance, write that 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" (1998: 13). Belief is of course a state that carries information, in the sense that it has a content—but so do any number of other states (intentions, desires, imagination and hypotheses). What distinguishes belief from these other attitudes is not the content, but the *attitude*. What we need to focus on, therefore, is what sets this *attitude* apart from other information carrying states, such as assuming that *p*, imagining that *p* or merely entertaining the thought that *p*.

The starting point for Clark and Chalmers, again, is that belief is a folk psychological kind, and when they speak of 'the role' of belief what

they have in mind is its role in folk psychological explanations. This means that differences between extended and internal states that do not matter for the purposes of folk psychology, are to be considered ‘shallow’, as lacking relevance for the classification of a state as a belief. In later texts, Clark has been quite explicit on this. Commenting on the case of Otto, for instance, he writes:

[A]ll we meant is that *for most ordinary folk psychological purposes*, we look onto many of the very same patterns in Otto’s actual and counterfactual behavior by treating the notebook entries as part of the mechanical supervenience base for his standing beliefs. In that restricted sense, and only in that restricted sense, are the two resources said to govern behaviors in similar enough ways. Importantly, this is something they can thus do despite a multitude of other more subtler differences . . . (Clark 2011: 451)

I take this to be a very promising starting point. Arguably, belief plays a distinctive folk psychological role, both in practical and in theoretical reasoning: Belief provides *reasons* for actions and further beliefs, in a way that merely imagining that *p* or assuming that *p* does not. The difference between imagining going to the Oscar ball tomorrow, and believing that I will go to the ball tomorrow turns on this reason giving role: If I actually believe it then (given normal background conditions) I will act a certain way (fix my hair and buy a ball gown) and draw certain inferences (I infer that I won’t be home tomorrow and that I will see Tom Cruise tomorrow). Moreover, my belief will have some type of grounds (I believe I was sent an invitation). Arguably, if *none* of this is in place (no evidence, no inferences and no action) the state is not a belief but an imagination or a pretense. Put in terms of functionalism, belief is that state which is evidence sensitive (on the input side) and which provides reasons for further beliefs and actions (on the output side).¹⁷ I therefore accept Clark and Chalmers’ starting point. However, taking it seriously requires going beyond the ‘trust and glue’ conditions, and to address the question whether, in fact, the external information plays the right role in the subject’s practical and theoretical reasoning.

Clark and Chalmers focus on the role of belief in *practical* reasoning, in the production of action. The notebook information realizes Otto’s standing beliefs, they argue, since we can explain why Otto goes to 53rd Street precisely by appealing to the information in the notebook in just the way that we ordinarily explain actions. They compare Otto with Inga, an ordinary subject who has the standing belief (internally stored) that MoMA is on 53rd Street, and they argue that the notebook information guides Otto’s actions in the same way that Inga’s internal states guide her actions: Just as we explain Inga’s actions by appealing to her desire to go to the museum and her standing belief that the museum is on 53rd street, we can explain Otto’s action the same way. To insist that in Otto’s case we must go via the belief

that the museum is on the location written in the notebook, they argue, complicates the explanation unnecessarily: “We submit that to explain things this way is to take one step too many. It is pointlessly complex, in the same way that it would be pointlessly complex to explain Inga’s actions in terms of beliefs about her memory. The notebook is a constant for Otto, in the same way that memory is a constant for Inga; to point to it in every belief/desire explanation would be redundant. In an explanation, simplicity is power” (1998: 13–14).

The first thing to note, here, is that the added complexity in Otto’s case would hardly be *pointless*. After all, there is an *additional action* to be explained here: Why does Otto keep looking in his notebook? To explain that, we have to appeal to beliefs and desires that Inga does not have, such as the belief that the notebook contains important information about the location of places that Otto tends to forget, and the desire to find out the location of the museum of modern art. Simplicity may be power, but oversimplification falsifies the actual course of events—if there is an additional complexity at the level of actions, there should be one at the level of explanation.

In his foreword to Clark (2008a), Chalmers notes this difficulty and suggests that it is one principled place where the opponent of extended mind can resist. Otto must write in the notebook, and read what he has written, and this requires both perception and action, whereas in Inga’s case there is no such requirement. It might therefore be proposed that real perception and real action mark a boundary for the mental. Chalmers suggests that one can question this proposal, since one can question the idea that there is a deep difference between real perception and inner perception, between real action and mental action. However, he notes, this only works if the involvement of perception and action makes no important difference to the explanatory role of Otto’s extended state, and it is not so clear that this is the case, since one can ask the crucial question why Otto reached for his notebook: “This seems to be a perfectly good psychological question about the explanation of action. And the natural answer is: he wanted to get to the museum, he did not know its location, and he believed that the notebook contained the information. In this explanatory structure, we speak naturally as if Otto lacked the extended belief” (Chalmers 2008: xii).

I think that this does point to a very important difference between the case of Inga and that of Otto. The problem is not so much that Otto, in order to find out where the museum is, has to rely on perception. The problem is that Otto has to *find out* where the museum is, that the information in the notebook cannot interact *directly* with Otto’s internally stored beliefs and desires. In the case of Inga there is a very direct interaction between her beliefs and desires. Thus, after reading about the exhibit at MoMA, a simple piece of practical reasoning leads to her action:

Desire 1: *I go to MoMA.*
 Belief 1: *MoMA is on 53rd Street.*
 Action: *I go to 53rd Street.*

In Otto's case, by contrast, the connection between the information in the notebook and his internally stored beliefs and desires is mediated by a complex set of other beliefs and desires, including second-order beliefs. After reading about the exhibit at MoMA Otto reasons along the following lines:

Desire 1: *I go to MoMA.*
 Belief 1: *I do not remember where MoMA is.*
 Desire 2: *I find out the location of MoMA.*
 Belief 2: *This kind of information is written down in the notebook.*
 Desire 3: *I look in the notebook.*
 Belief 4: *It says in the notebook that MoMA is on 53rd Street.*
 Belief 5: *I believe everything written in the notebook.*
 Belief 6: *I believe that MoMA is on 53rd Street.*
 Action: *I go to 53rd Street.*

It might be objected that in Inga's case the reasoning is actually more complex as well, since she may have to *recall* where MoMA is. That is, she desires to go to MoMA and she makes an effort to recall where the museum is located just like Otto has to look up the location in the notebook.¹⁸ Clark (2008a) defends a response along these lines. He considers an objection that he labels the 'Otto 2-step'. According to this objection, all that Otto believes in advance is that the address is in the notebook (step 1) and this leads him to look in the notebook (step 2) which, in turn, leads him to have a new belief about the address of MoMA (2008a: 80). Clark responds that we could say the same in the case of Inga, since we could simply say that Inga's only antecedent belief was that the information was stored in her memory and that her retrieval of the information is a 2-step. We don't do it this way when explaining Inga's actions since it adds extra complexity, Clark argues, and for the same reason we should not do it in the case of Otto's actions.

However, this cannot be right. We cannot construe Inga's memory retrieval as a step-2 process since it cannot be said that all that Inga believes in advance is that the information is stored in her memory—that, after all, would entail that Inga does *not* have the standing belief that MoMA is on 53rd Street. To have a standing belief, again, is not simply to have some information stored, to have retained some true contents, it is to have a standing *attitude* towards this information, a belief. Instead, memory recall is a simple step 1 process: What happens is not that Inga has to go through second order beliefs about her own beliefs, but simply that a standing belief (as a result of the effort to recall) becomes *occurrent*.

Indeed, even if one were to insist that memory recall should be construed as a mental *action* (which does not seem *prima facie* plausible) Inga's reasoning will be distinct from Otto's. Inga would reason:

Desire 1: *I go to MoMA.*

Belief 1: *The location of the museum is stored in my memory.*

Desire 2: *I do memory recall.*

Action: *Memory recall*

Belief 2 (occurrent): *MoMA is on 53rd Street.*

Action: *I go to 53rd Street.*

There simply is nothing corresponding to Otto's Belief 4 here, to his belief about what it says in the notebook, or to his second-order beliefs 5 and 6. Instead, the memory recall immediately leads to Inga's occurrent Belief 2. This brings out why in Inga's case the longer explanation does indeed seem to add unnecessary complication—it does not actually change the structure of the explanation. Adding that she believes that the information is stored in her memory may serve to explain why she goes through memory recall (if, again, we think of this as an action in need of an explanation), but it can be eliminated without changing the explanation of her going to 53rd Street. In Otto's case, by contrast, the additional beliefs cannot be eliminated.¹⁹

There is therefore a striking difference between ordinary beliefs and Otto's extended 'notebook states': What is distinctive of normal first-order beliefs is precisely that they can interact *directly* with other first-order beliefs and desires, without having to be mediated by further beliefs and desires. In this very important respect Otto's notebook states do not function the way belief does in ordinary folk psychological explanations. The information as such is simply *inert*, incapable of directly interacting with Otto's further beliefs and desires. That is to say, the information is just *information*—not a belief.

This becomes clear, also, if we consider the role of belief in *theoretical* reasoning: The externally stored information does not interact with the rest of Otto's beliefs the way beliefs ordinarily do. For instance, there was a period when MoMA was temporarily moved to Queens. When Inga finds this out she automatically updates the belief that MoMA is on 53rd Street as well as the rest of her related beliefs (she no longer believes that MoMA is in Manhattan, etc.). The information in the notebook, by contrast, functions very differently: It is not, in itself, sensitive to the new evidence. Upon reading about the move Otto may of course decide that it is time to update the information in his notebook, but the new (internally stored) belief does not have any direct impact on the information in the notebook. Instead, it must go via Otto's beliefs about the notebook as well as via his desire to update the notebook in order for it to function as a support of his failing memory.

Just like in the case of Otto's practical reasoning, therefore, his theoretical reasoning will differ from Inga's in some crucial respects.

Indeed, until Otto changes the information in the notebook, the claim that his standing beliefs are realized in the notebook implies that he holds simple contradictory beliefs: He believes both that MoMA is on 53rd Street and that MoMA is not on 53rd Street. This irrationality, of course, spreads to Otto's further beliefs. For instance, upon hearing that MoMA has moved to Queens, if he were to infer (like Inga) that MoMA is not in Manhattan, he would believe both that MoMA is on 53rd Street (external belief) and that it is not in Manhattan (internal belief).²⁰ Similarly in the case of Lotte. For instance, the claim that she believes everything in the dictionary may imply that she believes (internally) that the battle of Hastings was in 1096 (being a historical ignoramus) and believes (externally) that it was in 1066. Indeed, in Lotte's case the extended mind thesis is likely to imply that there are countless inconsistencies and contradictions in the belief system—a 37 volume encyclopedia will contain much that is inconsistent with Lotte's 'internal' beliefs.

The failure of extended beliefs to interact properly with other beliefs is stressed by Daniel Weiskopf (2008). A central feature of ordinary beliefs, Weiskopf argues, is that they are "informationally integrated, and updated in concert with, other beliefs (and further mental states of the subject, such as desires)" (268). For example, beliefs about people's marital status and beliefs about their home addresses interact. However, Weiskopf argues, in the case of external 'beliefs' this type of informational integration fails. He considers a mixed situation where belief A is stored internally and belief B is stored externally, and a pure external situation where both beliefs are stored externally. Thus, there is Waldo who suffers less extreme memory loss than Otto, and stores some beliefs internally, some externally. He notes down that the museum is on 53rd street. A little later he learns that the museum has been torn down for a bypass and this he remembers biologically. In this situation, a normal subject will revise her belief that the museum is on 53rd street and form the belief that the museum was on 53rd street but that it is no more. But Waldo's beliefs are not updated this way, they are not automatically and unconsciously updated to reflect the new information (2008: 269). Similarly in the pure external case, where both pieces of information are stored externally: Otto has written down on page 10 that the museum is on 53rd street and on page 20 that it has been demolished, and there is no reason to think that a normal consequence of writing down the second sentence is that the first one is erased or that related beliefs are appropriately updated (ibid).

Now, it is of course true that it happens that ordinary beliefs fail to interact the normal way as well. For instance, there are well-known cases of irrationality, where the subject holds a belief which is not sensitive to evidence or where she compartmentalizes a belief from the rest of her beliefs and desires. Indeed, sometimes the compartmentalization can even lead to

contradictory first-order beliefs. However, these cases of irrationality involve familiar psychological failings, ranging from simple distraction to repression and self-deception. As such they can be explained by appealing to ordinary facts about human psychology. The type of irrationality that would result from extended beliefs, however, cannot be explained this way but is wholly *theory-imposed*: We are to ascribe radical irrationalities to subjects who are prima facie rational, not disturbed by any of the ordinary psychological obstacles, simply because it follows from the theory.²¹

Moreover, there are limits to how irrational a state can be and still qualify as a belief. If belief is characterized by its role in folk psychological explanations, then a state that utterly fails to play the role of belief (it is not evidence sensitive and it does not play a role in theoretical and practical inferences) is not a belief. This is something Clark and Chalmers should accept. Delusional beliefs are an interesting case here. Prima facie, these seem to behave in a way that is incompatible with the role of belief in reasons explanations. However, it should be noted that what psychologists characterize as delusional beliefs is not utterly cut off from the ordinary belief role. Thus, delusional beliefs, according to the standard view, are caused by powerful but unusual experiences. For example, cases of Capgras delusion, where people believe that a close family member has been replaced by an impostor, are caused by a very powerful experience that when looking at the family member, for instance one's mother, it does not 'feel' as if it is one's mother. That is, there is a very strong sensory experience providing evidence for the delusional belief (caused by neurological damage relating to the connection between the face-recognition system and the autonomic nervous system).²² Moreover, the belief is not completely inferentially disconnected: It leads the person to form further beliefs and act in certain ways. Thus, the delusional belief shows both some evidence sensitivity and some integration with the subject's further beliefs and desires. This, arguably, suffices for the state to qualify as a belief. What makes it irrational, nevertheless, is that it is disconnected from the subject's *larger* belief system and actions (the person does not take in evidence provided by other people, she does not report her mother missing at the police station, etc.). What seems to happen, therefore, is that the delusional belief becomes part of an alternative, compartmentalized set of beliefs and desires, within which the state functions in a way that is sufficiently similar to regular beliefs for it to qualify as a (delusional) belief.²³ The notebook 'states', by contrast, do not even function like a compartmentalized set of beliefs: It is not as if they belong to a separate part of Otto's belief system within which they exhibit a certain form of evidence sensitivity and a certain (but limited) form of inferential interaction with further beliefs and desires.

If, therefore, we take the folk psychological role of belief to be distinctive of the state, as Clark and Chalmers do, there are very good reasons to question the claim that the notebook information qualifies as belief: If

this information does not interact with further beliefs and desires the way ordinary beliefs do, then the difference between the internal and the external ‘states’ is not, as Clark and Chalmers suggest, shallow—not even by their own lights. Otto may well have the disposition to believe everything in the notebook, but we must not conflate having a disposition *to believe* that *p*, and having the dispositional (standing) *belief* that *p*.²⁴

Given this, let us, finally, return to the problem of cognitive bloat.

4. How Not to Overextend Knowledge

I have argued that the information in Otto’s notebook does not play the role of belief and, more generally, that the subject’s having constant, reliable and easy access to external information does not suffice for this information to play the role of belief. This is not to say that extended beliefs are impossible. As noted above, once we accept a functionalist account of belief there is a principled possibility of external vehicles. However, the functional demands are harder to meet than is usually assumed, and go well beyond the trust and glue conditions. What is required is that the external information plays the role of belief in ordinary folk psychological reason explanations, and this condition is not met simply by the subject’s having easy and reliable access to the information. What is required in addition, is that the external information interacts directly with the subject’s further beliefs and desires. It is possible that in the future a device could be built (perhaps replacing part of a damaged brain) which allows the externally stored information to fully interact with the internally stored information. However, something much more sophisticated would be needed than a simple notebook or an iPhone.²⁵

This, I want to propose, provides us with a solution to the problem of cognitive bloat, since it severely restricts the cases of extended belief: Otto’s use of the notebook does not qualify, nor does Inga’s use of the Encyclopedia or my incessant use of the iPhone. In neither of these cases can the external information be said to be integrated with the subject’s overall system of beliefs and desires in the way required by folk psychology. Moreover, when we consider the role of belief in theoretical reasoning, it becomes clear that the failures of the extended states in these cases are *epistemically relevant* failures. The extended information is not inferentially integrated with the subject’s set of beliefs: It fails to be sensitive to evidence and to cohere with the subject’s existing set of standing beliefs—it will not be supported by the subject’s further beliefs and may even contradict these beliefs. These are precisely the kinds of failures that suffice to undermine the status of a belief as knowledge.

Now, it is of interest to see how this proposal satisfies the central intuitions employed by those who try to solve the problem of cognitive bloat by appealing to ability or virtue theories of knowledge. Thus, despite the talk

of ‘cognitive responsibility’, Roberts’ central claim concerns ‘full cognitive integration’, the idea that external information qualifies as knowledge only if it is properly supported by evidence and the subject is sensitive to inference relations.²⁶ Similarly, Pritchard suggests that what is required for a cognitive ability to be manifested, beyond mere reliability, is that the reliable belief-forming process is “*integrated* within, and therefore part of, the cognitive character of the agent, where an agent’s cognitive character is her integrated web of stable and belief-forming processes” (2010: 136).²⁷ On the view I have suggested the demand on cognitive integration is equally important, only it does not derive from assumptions about knowledge but from assumptions about the nature of belief: To qualify as a belief the external information must be properly integrated, both with the external information as a whole (notebook inconsistencies need be removed) and with the subject’s internally stored beliefs.

Similar remarks apply to another proposal present in the literature, that preventing cognitive bloat requires appealing to the idea that the external information *belongs* to the agent, that the relevant requirement concerns the ‘ownership’ of this information. This idea comes from outside the ability/virtue camp and is defended by those with more externalist leanings, but it serves the same function as the appeal to the subject’s ‘cognitive character’. For instance, Clark (2008b) speaks of the need to find further conditions of proper ownership of a candidate extended process. He continues: “Perhaps we ought rather to speak here of ‘proper inclusion within a distinct cognitive agent’ rather than ‘proper ownership by’ such an agent, so as to avoid giving hostages to internalist prejudice. But either way, the point of all these considerations is really to somehow *tie* the candidate process to a given agent” (2008b: 454).²⁸ The best way to tie the extended information to an agent, I submit, is to require that it is sufficiently inferentially integrated with her internally stored beliefs (and desires, if we also consider the role of belief in practical reasoning).

In the light of this, it is of interest to consider the relation between the view proposed here and epistemological internalism. On the one hand, the possibility of extended knowledge would seem to sit uneasily with certain formulations of internalism. On the other hand, the appeal to inferential integration and evidence seems internalist in spirit, or at least ‘mentalist’ in the sense defended by Earl Conee and Rich Feldman (2001). Is this a real conflict or is it merely apparent? Mentalism, recall, is the claim that epistemic justification depends on the subject’s mental states, such as her beliefs and perceptions.²⁹ Conee and Feldman suggest that mentalism is akin to content internalism, in that it appeals only to what goes on ‘inside cognitive beings’ (2001: 3). It is therefore natural to assume that mentalism is incompatible with extended knowledge, since in the case of extended knowledge not only will the knowledge carrying belief be realized externally, the justification of this belief may also depend on further externally realized beliefs. However,

it all depends on what is meant by ‘internal’. Thus, Conee (2007) argues that according to mentalism, ‘internal’ simply means ‘mental’: “Justification does have to be ‘internal’ to something, namely, the mind” (2007: 59). If this is how ‘internal’ is construed, clearly, then the extended mind thesis is not incompatible with mentalism.

The account of extended belief suggested here, therefore, is perfectly compatible with the claim that all justifiers are mental and in that sense ‘internal’ to the subject. The requirement that the external information is inferentially integrated with the subject’s internally stored beliefs, is fully in line with the idea that the justification of a belief depends on the extent to which it is evidentially supported by the subject’s further beliefs (and experiences). However, Conee and Feldman do not merely claim that all justifiers are mental. They also claim that the *justificatory status* of these mental states (strongly) supervene on the subject’s mental life, such that if two subjects are ‘mentally’ identical, then the justificatory status of their beliefs is the same. I am skeptical of the strong supervenience claim since I believe that evidential relations typically depend on contingent circumstances. If one thinks of the evidence relation as holding between two propositions p and q , such that p provides evidence for q if the truth of p raises the probability of q , then whether p constitutes evidence for q will depend on contingent facts about the world since (in most cases) it is a contingent fact that the truth of p in this sense raises the probability of q . For instance, in some worlds having a perception with the content *X looks like a barn* provides evidence (sufficient for knowledge) for the belief that X is a barn, but in other worlds it does not. If so, then the justificatory status of the subject’s beliefs does not supervene on her mental states: Two subjects may have all the same beliefs and yet one of them will have justified beliefs while the other will not. There is much to be said about this topic and I shall have to leave a further discussion of it to another occasion.³⁰ The important thing to notice here, is just that the failure of supervenience (assuming it does fail) is quite independent of the extended mind thesis: It depends merely on the idea that even if all justifiers are internal states, the justificatory status of these internal states depends on facts beyond the subject’s mental states (whether the latter are internally or externally realized).³¹

Conclusion

I have argued that the worry of ‘cognitive bloat’ can be met by preventing a bloating of belief. Moreover, I have argued that we can prevent this if we take seriously Clark and Chalmers’ suggestion that what is distinctive of belief is the role it plays in ordinary folk psychological reasons explanations. This allows for the principled possibility of extending belief, but only under conditions where the extended state functions as a *candidate* for knowledge

—i.e. when it is sufficiently evidence sensitive and inferentially integrated with the subject's larger set of beliefs. Naturally, not all candidates of knowledge qualify as knowledge, not even if they are true and justified, since there are (plausibly) further conditions that have to be met for a belief to qualify as knowledge. Nevertheless, when the conditions on being a belief is met by an external feature, and the belief *does* qualify as knowledge, there will not be a problem of cognitive bloat: The result will not be overextended knowledge but, simply, extended knowledge.

Notes

1. For a recent discussion see Allen-Hermanson (2012).
2. Clark (2008a) is very explicit on this and argues that we must be able to separate that which belongs to cognition properly and that which merely belongs to the background conditions of cognition.
3. For this line of reasoning see Gertler (2007). Clark and Chalmers agree that consciousness (probably) does not extend, and stress that their concern is with non-occurrent, or standing states (1998: 10). Gertler recognizes that standing beliefs are not conscious and she therefore concludes that they do not belong to the mind. I shall not take a stand on what is distinctive of the mind, but will focus on what is distinctive of *belief*: Since most of our beliefs are standing, clearly, consciousness cannot be construed as a distinctive mark of belief.
4. For this line of response see Adams and Aizawa (2008).
5. See Pritchard (2010) and Roberts (2012). I return to the details of this line of argument in section 2 below.
6. For this reason the term 'vehicle externalism' is very apt for the extended state thesis (see Hurley 1998).
7. In the literature the general assumption is that the extended mind thesis presupposes functionalism. See for instance Weiskopf (2008) and Wheeler (2010). See also Sprevak (2009) who takes the problematic implications of the extended cognition thesis to constitute a counterexample to functionalism.
8. Chalmers (2008) argues that all that is presupposed is the very weak functionalism captured by the Parity Principle, the idea that sameness in role implies a presumption of mentality. But while this may be a weak assumption when it comes to cognitive *processes* it is by no means a weak assumption when it comes to mental states, such as belief. The idea is not only incompatible with identity theories of belief, but with teleological theories as well as normativist theories of belief. For a defense of the teleological theory of belief see for instance Velleman (2000), and for a normativism see Shah (2005).
9. See Weiskopf (2008) who points out that there is a sense in which the extended mind thesis should not be seen as especially radical: "Functionalism has all along been committed to the possibility of extrabodily states playing the role of beliefs and desires" (2008: 267).
10. Chalmers (2008) and Clark (2008a).
11. Gertler (2009) mentions the example of the belief that there are no bicycles on the moon, and she labels these beliefs implicit standing beliefs, to distinguish them

- from those where the content was once explicitly endorsed. In addition, there are all the beliefs that we come to have as part of becoming cognitive agents in the world, beliefs belonging to the background of all action and reasoning (such as the belief that there is a floor under my feet or the belief that there are other people).
12. Notice, also, that even if we hold on to the condition it does not seem to solve the problem of cognitive bloat. We can easily change the example and imagine that Lotte spends a year reading all 37 volumes, consciously endorsing everything in them.
 13. For a classic discussion see Alston (1988a). The problem, note, is not merely that belief is formed *involuntarily*, but that belief formation cannot properly be seen as an intentional *action*. See Glüer and Wikforss (2009) for objections to the idea that belief formation is a form of action.
 14. This is stressed by Adams (2012). He writes: “I can no more voluntarily refuse (nor do I deserve credit for) the beliefs that perception generates in me than the BIV hooked up reliably to the supercomputers” (114). Pritchard suggests that when it comes to belief-forming processes that are present from the start, such as perception, the requirement of cognitive responsibility can be relaxed. However, if cognitive responsibility is a necessary condition on knowledge, then it is difficult to see why the history of the mechanism in question should make a difference.
 15. These problems are discussed in some detail in Glüer and Wikforss (2009). See also Glüer and Wikforss (2014).
 16. Weiskopf (2008) notes this. Beliefs, he writes, are not information itself, “they are bearers of information” (2008: 266, fn 2).
 17. This account of belief will be developed and defended in Glüer and Wikforss (MS). A first stab is made at developing our version of functionalism, what we call ‘reason-providing functionalism’, in a presentation given at the “Aim of Belief” workshop in Oslo, June 2009. For some further preliminary work see Glüer and Wikforss (2013).
 18. This was suggested to me by Nathalie Ashton, in her written comments on an earlier version of this paper.
 19. Chalmers (2008) suggests the proper conclusion is that the classification of states can depend on our explanatory purposes. When we are interested in explaining Otto’s local-scale interactions with the notebook it is natural to deny that he has the extended belief, Chalmers argues, but when we are interested in explaining why Otto walked north it is natural to explain this in terms of his standing, extended belief that MoMA is on 53rd Street (2008: xii). However, this is not compelling. The reason we can give a simpler explanation in the second case, after all, must be that Otto now has an occurrent belief that MoMA is on 53rd Street (or an internally stored standing belief to that effect)—otherwise he would keep looking in the notebook.
 20. Also, contradictory beliefs undermine Otto’s practical reasoning since they block all actions: If Otto desires to go to MoMA he has a reason to go to 53rd Street and (an equally strong) reason not to go to 53rd Street. See also Weiskopf who notes that the attribution of irrational belief systems undermines the explanatory power of the beliefs (2008: 274).

21. In this respect there is a similarity between the extended belief thesis and content externalism, since in both cases it follows that subjects are to be described as holding radically irrational beliefs that cannot be explained by appealing to the ordinary facts about human psychology. See Wikforss (2008).
22. See for instance Davies and Egan (2013).
23. That delusional belief qualifies as belief is not uncontroversial and there is much more to be said about it. For an interesting discussion see Wilkinson (2013).
24. For this distinction, see Gertler (2009).
25. Discussing the case of Otto, Clark (2008b) stresses that the claim was not that an external, passive encoding might somehow behave exactly like the “fluid, automatically responsive resources of internal biological memory”. Rather, Clark says, the claim was that external encodings could become “so deeply integrated into online strategies of reasoning and recall as to be only artificially distinguished from proper parts of the cognitive engine itself” (2008b: 46). This ‘could’-claim, again, is not in dispute. What I dispute is that the information in ordinary notebooks and smart phones are ‘deeply integrated’ into the subject’s reasoning.
26. For instance, Roberts argues that a central failing of reliabilism is that it allows a belief to count as knowledge also when the subject has no evidence for the belief in question (she may even have substantial evidence against it), and stresses that being sensitive to inference relations is central to knowledge (2012: 137–138).
27. See also *ibid.* 147–148 where Pritchard stresses that the external information must be inferentially integrated in the subject’s web of beliefs if it is to qualify as knowledge.
28. See also Mark Rowlands who appeals to the notion of ‘ownership’, suggesting that ownership is a matter of integration into the psychological life of the subject (2009: 17).
29. Mentalism is a species of internalism that is weaker than so called *access* internalism. According to access internalism, justification is determined by the subject’s mental states but it is also required that all the justifiers are cognitively accessible, available to introspection or reflection.
30. For a discussion of epistemic reasons see Glüer and Wikforss, “Epistemic Reasons Without Normativity”, forthcoming in *The Oxford Handbook of Reasons and Normativity*, ed. D. Star.
31. This, in effect, is the position Alston (1988b) labels ‘internalist externalism’, the view that while all justifiers are mental and in this sense internal, their justificatory status depends on external, contingent facts.

References

- Adams, F. 2012. “Extended Cognition Meets Epistemology”, *Philosophical Explorations* 15: 107–119.
- Adams, F. and Aizawa, K. 2008. “Why the Mind is Still in the Head”, in *Cambridge Handbook of Situated Cognition*, eds. P. Robbins and M. Aydede, Cambridge University Press: Cambridge.
- Allen-Hermansson, S. 2012. “Superdupersizing the Mind: Extended Cognition and the Persistence of Cognitive Bloat”, *Philosophical Studies* 158: 1–16.
- Alston, W. 1988a. “The Deontological Conception of Epistemic Justification”, *Philosophical Perspectives* 2: 257–299.

- 1988b. “An Internalist Externalism”, *Synthese* 74: 265–283.
- Chalmers, D. 2008. “Foreword”, in Clark 2008a.
- Clark, A. and Chalmers, D. 1998. “The Extended Mind”, *Analysis* 58: 7–19.
- Clark, A. 2008a. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*, Oxford University Press: Oxford.
- 2008b. “Pressing the Flesh: A Tension in the Study of the Embodied, Embedded Mind?”, *Philosophy and Phenomenological Research* 76: 37–57.
- 2011. “Finding the Mind”, *Philosophical Studies* 152: 447–461.
- Conee, E. and Feldman R. 2001. “Internalism Defended”, *American Philosophical Quarterly* 38: 1–18.
- Conee, E. 2007. “Externally Enhanced Internalism”, in *Internalism and Externalism in Semantics and Epistemology*, 51–67, ed. S. Goldberg, Oxford University Press: Oxford.
- Davies, M. and Egan, A. 2013. “Delusion: Cognitive approaches—Bayesian Inference and Compartmentalization”, in *Oxford Handbook of Philosophy and Psychiatry*, 689–726, Oxford University Press: Oxford.
- Farkas, K. 2012. “Two Versions of the Extended Mind Thesis”, *Philosophia* 40: 435–447.
- Gertler, B. 2007 “Overextending the Mind?”, in *Arguing About the Mind*, 196–205, eds. B. Gertler and L. Shapiro, New York and London: Routledge.
- 2009. “Self-Knowledge and the Transparency of Belief”, in *Self-Knowledge*, ed. A. Hatzimoyisis, Oxford University Press: Oxford.
- Glüer, K. and Wikforss Å. 2014. “Against Belief Normativity”, forthcoming in *The Aim of Belief*, ed. T. H. W. Chan, Oxford University Press: Oxford.
- and Wikforss, Å. 2013. “Aiming at Truth: On The Role of Belief”, *Theorema*, 32: 137–160.
- and Wikforss, Å. 2009. “Against Content Normativity”, *Mind*, 118: 31–70.
- Hurley, S. 1998. “Vehicles, Contents, Conceptual Structure, and Externalism”, *Analysis* 58: 1–6.
- Pritchard, D. 2010. “Cognitive Ability and the Extended Cognition Thesis”, *Synthese* 175: 133–151.
- Roberts, T. 2012. “You Do the Maths: Rules, Extension, and Cognitive Responsibility”, *Philosophical Explorations* 15: 133–145.
- Rowlands, M. 2009. “Extended Cognition and the Mark of the Mental”, *Philosophical Psychology* 101: 389–428.
- Shah, N. 2005. “Doxastic Deliberation”, *Philosophical Review* 112: 447–482.
- Sprevak, M. 2009. “Extended Cognition and Functionalism”, *Journal of Philosophy* 106: 503–527.
- Velleman, D. 2000. *The Possibility of Practical Reason*, Oxford University Press: Oxford.
- Weiskopf, D. 2008. “Patrolling the Mind’s Boundaries”, *Erkenntnis* 68: 265–276.
- Wheeler, M. 2010. “In Defense of Extended Functionalism”, in *The Extended Mind*, ed. R. Menary, MIT Press.
- Wikforss, Å. 2008. “Self-Knowledge and Knowledge of Content”, *Canadian Journal of Philosophy*, 38: 399–424.
- Wilkinson, S. 2013. “Beyond Believing Badly”, *Theorema* 32: 105–120.