

Belief may not be a necessary condition for knowledge¹

Katalin Farkas, Central European University

Forthcoming in *Erkenntnis*. The final publication is available at Springer via <http://dx.doi.org/10.1007/s10670-014-9620-2>

Abstract: Most discussions in epistemology assume that believing that p is a necessary condition for knowing that p . In this paper, I will present some considerations that put this view into doubt. The candidate cases for knowledge without belief are the kind of cases that are usually used to argue for the so-called 'extended mind' thesis.

Keywords: analysis of knowledge, dispositions and beliefs, extended mind

I. Internal and external components in knowledge

We acquire much of our knowledge of the external world by exercising our cognitive abilities: we open our eyes and ears, and without seemingly trying to do anything, we arrive at some knowledge. In some cases, more effort is required: we seek evidence, we carefully adjust our beliefs. We can thus contribute a lot to acquiring knowledge, both by exercising our abilities and by making visible efforts. But knowledge of the external world is never entirely under our control in the following sense: no matter how much we contribute in the ways just described, we always have to rely on the cooperation of the environment. As Wittgenstein puts it: “It is always by favor of Nature that one knows something (OC, §505)”. For example, nature does us a favor – or rather refrains from doing a disfavor – in sparing us from the deceptions of an evil demon. Further, objects stay stable enough to be explored, and they exert all sorts of causal impact upon us. Their surfaces reflect light, they release compounds that stimulate our olfactory and gustatory senses. There is air around us to carry sound waves. Our fellow human beings do not constantly try to mislead us. And so on.

“Internal” and “external” factors in knowledge are balanced differently on different theories of knowledge, and there are differing views on how much exercise of a cognitive ability or wielding of cognitive effort is required for knowledge. But the majority of even the most committed internalists delegate *some* work to a cooperating external world. Almost everyone will agree that there could be subjects who are in some sense in the same position as we are “internally”, and yet lack knowledge: for example, our brain-in-a-vat counterparts. In what sense these subjects are “internally” the same is a matter of controversy. Some will say that the ignorant subjects have the same mental states or the same justification or evidence, some say that they have the same physical states, some others will say that their position is simply subjectively indiscriminable from ours. Be

¹ Talks bearing more or less close relation to this paper were presented at the 2013 Pacific APA meeting, at the 2013 Swedish Congress of Philosophy and at a workshop on the Varieties of Externalism in Edinburgh. Thanks to Brie Gertler, Sandy Goldberg and Tim Crane for discussion, and to Timothy E. Kunke for commenting on the talk in Edinburgh. Two anonymous referees for *Erkenntnis* provided very constructive suggestions for revision.

that as it may, most will agree that one crucial difference between us and our ignorant counterparts is supplied not by our best efforts, or by our abilities narrowly construed, but rather by a relation to a cooperating – or at least non-hostile – environment.

The proposal I am trying to explore in this paper is that with the advance of technology and the increase in division of labour, we could delegate a somewhat *increased* share of the effort of obtaining and sustaining knowledge to a cooperating environment. There are parts of knowledge that are too tedious to acquire and retain in our head: remembering phone numbers or birthdays, for example. A student studying for a case law exam may understand all the intricate details of relevant precedents, but has a hard time in remembering whether the case was, say, Smith vs. Jones 1998, or Jones vs. Brown 1989. I propose that in such and similar cases, we can summon a cooperating environment to help us to store information, essentially in the way depicted in the so-called “extended mind” scenarios. I shall argue that we can naturally *extend* the application of the concept of knowledge to such cases. Andy Clark and David Chalmers, who introduced extended mind scenarios (1998), put forward a similar argument for beliefs: they argued that given our notion of belief, we can naturally extend the application of the notion to these scenarios. My proposal is that this works better for knowledge than for belief: extending the concept of knowledge is more straightforward and faces fewer objections. Hence in the envisaged scenarios, we have *prima facie more* reason to attribute knowledge than to attribute belief.

The paper is structured as follows. Section II gives the briefest sketch of some conceptions of knowledge that could underlie the proposal that knowledge without belief is possible. These conceptions rely on at least two ideas proposed in contemporary epistemology: that knowledge comes first, and that the function of the concept of knowledge is to flag reliable informants. Section III presents some cases which have recently been claimed to involve knowledge without belief which are nonetheless quite different from the current proposal. In section IV, I shall describe a putative extended mind scenario. In section V, I note that in ordinary situations, it is quite natural to attribute knowledge in some of the extended mind cases, and that certain conceptions of knowledge fit these cases. In section VI I recall some considerations against regarding extended mind scenarios as cases of belief, and in sections VII and VIII I argue that the reasons against attributing beliefs are not, at the same time, equally good reasons against attributing knowledge in these scenarios. Section IX considers an objection that applies equally against extended beliefs and extended knowledge, and touches upon the epistemic significance of consciousness. Section X considers another objection, based on the facticity of knowledge, that can be brought against various proposals that countenance knowledge without belief, and attempts to answer the objection.

Section XI presents the conclusion that in “extended mind” scenarios, we have *prima facie*

more reason to attribute knowledge than to attribute belief. Section XII points out that the gap between attributing knowledge and belief leaves open at least the following alternatives. Some philosophers will want to preserve a strong connection between knowledge and belief, and will argue that if we have knowledge, then we must have the corresponding beliefs. These philosophers may have an additional argument for attributing beliefs in extended mind scenarios. Other philosophers will have strong independent reasons to resist the suggestion that subjects in “extended mind” scenarios actually have beliefs, but are convinced by the arguments in favour of attributing knowledge. On this view, may have cases where knowledge is possible without belief

II. Knowledge first, and belief may not even be second

In this paper, I will try to explore the possibility of knowledge without belief. If all defensible accounts of knowledge obviously and directly rely on belief being a necessary condition for knowledge, then this project is doomed: it would be perhaps like searching for a red object without colour. So I will need at least a sketch of a plausible account of knowledge that does not immediately rely on belief as a necessary condition for knowledge.

The standard analysis of knowledge begins with the observation that knowledge entails truth: if S knows that p , then p (or ' p ' is true). This is accepted: when we are interested in knowledge, we are interested in the search for truth. Then clearly, a knowing subject must have some sort of relation towards a truth. Many analyses of knowledge then quickly proceed to state that this relation must be an attitude, the object of attitude a true proposition, and the attitude in question must be belief. This gives rise to the idea that knowledge must be true belief plus something. In fact, these moves are too quick.

To resist the idea that knowledge must be true belief plus something, we can recruit some ideas from Timothy Williamson’s “knowledge-first” epistemology (Williamson 2000). Williamson asserts the primacy of the concept of knowledge, and denies that it can be given a reductive analysis in terms of notions like belief, truth and justification. Williamson does not deny that an interesting and illuminating explanation can be given of the concept of knowledge; indeed he himself offers such an explanation. However, he argues that this explanation will not proceed through identifying components of knowledge like belief. Williamson himself holds that knowing itself is a mental state, but a knowledge-first epistemology need not agree with this. For example, one of Williamson's central theses can be reformulated as the claim that knowledge is the most general factive state, without necessarily committing us to the claim that it is a *mental* state. A number of arguments given by Williamson can be reworked to support the view that knowledge is a state (again, not necessarily a mental one) which cannot be analysed as true belief plus something.

Another possibly promising line of inquiry would focus on the *function* of attributing knowledge. Edward Craig also advises against developing a theory of knowledge through providing necessary and sufficient conditions. Instead

We take some prima facie plausible hypothesis about what the concept of knowledge does for us, what its role in our life might be, and then ask what a concept having that role would be like, what conditions would govern its application. (Craig 1990 p.2)

In pursuing this project, Craig further assumes that “[k]nowledge is not a given phenomenon, but something that we delineate by operating with a concept which we create in answer to certain needs, or in pursuit of certain ideals.” (ibid, p.2) The function of the concept of knowledge, according to Craig, is to flag reliable informants. He sets up the idea in the following way: human beings need to learn about the environment to guide their actions. They have “on-board” sources of information, including their senses, and their powers of reasoning. But the information they can access this way is limited, so they will do well to rely on others in their community. Hence they will have an interest in evaluating others as sources of information, and this is when the concept of knowledge comes handy: we attribute knowledge to approved sources of information. We can gain further insight into the nature of knowledge by considering which features would be most useful to identify approved sources of information. For example, according to Craig, reliable informants will typically believe the information they are imparting, but the account allows for exceptions, and indeed this is considered as an advantage of the theory (p.14).

These considerations are admittedly very sketchy, and on a final and detailed presentation of a theory of knowledge, I would depart from Williamson or Craig on many points. However, perhaps these brief remarks are sufficient to show, for the purposes of this paper, that there are various lines of inquiry into the nature of knowledge which do not necessarily invoke the concept of belief right away, or indeed allow for the possibility of knowledge without belief. We can form at least a preliminary understanding of what knowledge is, and what it is for, without assuming that the relation to truth involved in knowledge is necessarily the same as the attitude of belief. In section V and VIII, I will further elaborate how the extended mind scenarios may fit various conceptions of knowledge.

III. Inconsistent dispositional profiles

Eric Schwitzgebel and Blake Myers-Schulz (2013) recently argued that there may be cases of knowledge without belief. They present five scenarios where it is arguable that people are more

willing to attribute knowledge that p , than to attribute belief that p . The first scenario derives from a suggestion by Colin Radford (1966): it involves a person who feels insecure about her knowledge of British history, and feels that she is guessing when answering questions, yet she gives the correct answers most of the time.

The other four scenarios are familiar from recent discussions on the nature of belief (see for example Gendler 2008, Schwitzgebel 2010, Gertler 2011). Many philosophers agree that beliefs (conceived as standing states rather than episodes in the stream of consciousness) entail certain dispositions: dispositions to act, and dispositions to undergo certain kinds of conscious episodes. The problematic cases are those where someone manifests dispositions which partly match the dispositional profile associated with believing that p , but partly match the profile associated with believing not- p . The four scenarios presented by Myers-Schulz and Schwitzgebel all fall in this category, and they are instances of the following types:

1. Instinctive emotional reactions which are not endorsed. Someone may have a fearful reaction to something that, on reflection, they declare to be harmless.
2. Routine behaviour in the face of contrary information. Someone may learn that the usual route home is closed, and on reflection, they assent to the updated information, but when their mind is on something else, they act out the old routine.
3. Unconscious bias. Someone openly asserts the equality of genders, races, etc, and can even support the claim with arguments and data; yet many of their immediate reactions betray a prejudice.
4. Some cases of self-deception. The stock case is the husband who denies that his wife is unfaithful, yet performs a pattern of behavior that suggests that he suspects that she really is.

In the first three cases, the candidate for knowledge without belief is the proposition endorsed on reflection: that the thing causing fright is actually harmless; that the usual road is closed; that genders, races, etc. are equal. What may cast doubt on the claim that the subject *believes* these propositions is the fact that she manifests behavior that is consistent with believing the opposite. In the fourth case, the candidate proposition is the one suppressed and not consciously endorsed, namely, that the husband's wife is unfaithful. Here the lack of open endorsement may cast doubt on the subject's believing the proposition.

In the debate about the nature of beliefs, a number of different accounts of these cases have been put forward. Schwitzgebel (2010) argued that these are cases of 'in-between belief': the subject does not have either of the two beliefs that compete for explaining the behavioral dispositions.

Different accounts were given for example by Tamar Szabo Gendler (2008) and Brie Gertler (2010) at least for the first three types of cases: on these accounts, the subject does believe the proposition consciously endorsed. Therefore it is worth noting that the final verdict (the verdict that goes beyond recording intuitions, and is backed up by a theory) on whether we can have cases of knowledge without belief depends on our philosophical conception of belief.

The aim of this paper is also to suggest that belief may not be a necessary condition for knowledge, but my scenarios derive from the debate on the so-called 'extended mind hypothesis', and are quite different from those put forward by Myers-Schulz and Schwitzgebel.

IV. Extended mind scenarios

Andy Clark and David Chalmers introduced the now well-known case of Inga and Otto in their 1998 paper "The Extended Mind". Inga would like to go to the Museum of Modern Art (MoMA); she recalls that the MoMA is on 53rd street, and she sets off accordingly. Otto suffers from severe memory loss and therefore he keeps all important information recorded in a notebook which he carries with him all the time. When he decides to go to MoMA, he looks up the whereabouts of the museum, finds it's on 53rd street, and then he sets off. Many people agree that Inga had had the belief that the Museum of Modern Art was on 53rd street even before the issue came up in connection to her current visit. But Clark and Chalmers claim that if Inga has the belief, so does Otto, even before he looked up the information in his notebook. Otto has reliable, constant and easy access to the contents of his notebook, and he endorses the contents of his notebook automatically. This, according to Clark and Chalmers, is enough to qualify him as having the belief. Metaphorically, we may say that Otto's mind "extends" to his notebook.

Clark and Chalmers's proposal generated a lot of discussion, but my aim here is not to review this. Instead, I will put forward an alternative analysis of the case: that Otto *knows* that MoMA is on 53rd street, but he *may not believe* it. The idea is that certain kind of access to information gained and stored in an appropriate way may be sufficient for knowledge. Belief is one way to gain, store and have access to information in the appropriate way, but not the only one, and extended mind scenarios, though fall short of the requirements for belief, can still provide the kind of access to information that will suffice for knowledge. I will not claim that it's *clear* that Otto does not believe the proposition; I will only argue that there may be *more* reason to attribute knowledge to Otto than to attribute him belief. Incidentally, this conclusion is similar to the one offered by Myers-Schulz and Schwitzgebel: their claim is that the scenarios they present are more clearly cases of knowledge than cases of belief. Otto's case is obviously different from the above four scenarios discussed by Myers-Schulz and Schwitzgebel: Otto manifests a lot of dispositions that would be

characteristic of someone who believes that the MoMA is on 53rd street, but he does not manifest any dispositions that would indicate his believing the contrary. So if what I say in this paper is right, then we have some cases, quite different from those described by Myers-Schulz and Schwitzgebel, where we have more reason to attribute knowledge than to attribute belief.

V. EM scenarios as knowledge

Notice that in some everyday contexts, it is very natural to attribute knowledge to subjects who are in Otto-type situations. You ask me if I know NN's phone number, and I say “sure”, reaching for my smartphone. Although I don't know it *by heart* (which sounds like a particular, but not exclusive way of knowing something), I wouldn't hesitate to answer yes when I'm asked if I know it. Or I could be wondering about when the train is leaving tomorrow given the new summer schedule, and you propose asking the train inquiry service, since they will certainly know². It seems that it is more natural to ask questions and give answers in these situations in terms of what we know, rather than in terms of what we believe. For example, if someone asked whether I have beliefs concerning NN's phone number, I wouldn't even be sure what to answer. The question sounds more like a philosophical or theoretical enquiry, but then it would presuppose a certain theoretical or philosophical view of belief.

The observation about everyday usage is not decisive: maybe we speak loosely in these situations.³ Still, the examples are worth pursuing. Note that for example the conception of knowledge sketched in Section II fits the extended mind scenarios. Otto is a reliable informant on many issues, and so am I about phone-numbers. If the function of attributing knowledge is to flag reliable informants, then, as long as the usual extended mind conditions apply (eg. the information is easily accessible, reliably available and is automatically endorsed), both Otto and me will be approved sources of information, and hence on this conception, it will be natural to attribute knowledge both of us.

Note also that Craig's arguments for this theory are entirely independent from considerations

² Note that the attribution in these cases typically takes the form of a know-wh construction. Jonathan Schaffer (2007) notes that ‘know-wh’ attributions are ubiquitous (seem to be more frequent than ‘know-that’ attributions). Different analyses of know-wh ascriptions exist: some reduce it to know-that ascriptions, some – like Schaffer – offer a more complex analysis. Still, most proposals agree that know-wh ascriptions, just like know-that ascriptions, identify informational knowledge rather than skill-related knowledge (if that distinction makes sense at all).

³ One may make similar observations about everyday usage in the cases presented by Myers-Schulz and Schwitzgebel. Myers-Schulz and Schwitzgebel then conducted a survey among undergraduates of their university on whether protagonists in their stories know or believe certain things, and they used the results of the survey to support their conclusion. So one way to develop the thesis of the current paper would be to survey people for their responses to ordinary extended mind scenarios. However, I have some reservations about this argumentative strategy, so I will not pursue this direction in this paper.

about extended mind scenarios⁴. As I mentioned above, Craig does allow for the possibility of knowledge without belief. As examples, he mentions the Radford cases involving the person who gives reliably correct answers at an exam, but feels that she is guessing (these cases were already described above in section III, in the discussion of the paper by Myers-Schulz and Schwitzgebel). These cases are obviously different from extended mind scenarios. Craig actually does not rely on the specificities of the cases, but simply argues that though a reliable informant will typically have a belief concerning the subject matter she is informing us about, in certain circumstances, the belief condition may be missing. This conclusion is established without an explicit or implicit reliance on anything having to do with extended mind scenarios. However, when considering Otto's case and my case with phone numbers, we find that they fit this independently established conception.

Craig's theory is a natural ally in exploring the possibility of knowledge without belief, since his is one of the few theories that does not start an account of knowledge from assuming that knowledge is true belief plus something. But once we start to think of Otto's case as a case of knowledge, we can discover that it shows features that are analogous with a number of other theories of knowledge. Duncan Pritchard (2010) investigates the question of whether extended mind scenarios can be regarded as cases of knowledge under the proposal that knowledge involves a true belief that is due to cognitive ability. The production of a belief is due to a cognitive ability if a “reliable belief-forming process is *integrated* within, and therefore a part of, the cognitive character of the agent” (Pritchard 2010, p. 136). In this form, the proposal does not suit us, because it assumes that belief is a necessary condition for knowledge. However, we can easily modify the condition according to the idea broached in the previous section: we could require that instead of having a belief, the agent has access to information whose production and storage “is integrated within, and therefore part of, the cognitive character of the agent.”

Pritchard notes that Otto himself decided to extend his cognitive processes, thereby showing considerable epistemic virtue: he realised his own capacities were failing, and decided to support them appropriately. There are ways for Otto to enter information into the notebook, to maintain his records, to consult the notebook which show further epistemic virtues. This actually requires filling in some further details of the story that were not present in Clark and Chalmers's original account. For example, if Otto is to count as a knower, it seems to matter a great deal whether Otto keeps his entries in the notebook up-to-date, whether he bases the entries on reliable sources, and so on. But if we grant all this, we can conclude with Pritchard that on a plausible construction of cognitive ability, Otto's cognitive success *is* due to his cognitive ability, since he successfully integrated this information resource to his cognitive character. Pritchard's arguments apply to Otto's case even

⁴ Incidentally, Craig's book was published years before the extended mind hypothesis became a subject of discussion.

without the assumption that he possesses the appropriate beliefs. Otto's acquiring, storing and accessing information from his notebook can be still epistemically virtuous and form an integrated part of his cognitive character. On the basis of this, it is *prima facie* plausible to attribute him knowledge. As we shall see for example in section VIII, there are other epistemic conditions, like sensitivity and safety, which are applicable to Otto's case – and to my case with my phone – even without the assumption that we have the appropriate beliefs.

So far we have the following reasons to support that we have knowledge in “extended mind” scenarios: we naturally attribute knowledge in many extended mind scenarios, and there are several directions for developing the concept of knowledge that fits with these practices. A full defense of the position will require more work. However, the aim of this paper is simply to suggest this alternative analysis of the extended mind scenarios, and thereby present cases, hitherto not considered in this context, where we may have more reasons for attributing knowledge than for attributing beliefs.

VI. Belief or no belief?

It is sometimes claimed that what is most important in the debate about extended minds is that extended minds are actual, rather than merely possible. Certainly, *if* someone agrees with Clark and Chalmers that the fictional Otto has the beliefs in question, they could further argue that *our* own predicament is in some respects similar to that of Otto. For example, people who carry their smartphones with them all the time and rely on them heavily may be regarded as having their mind extending to their smartphone. However, many of those who oppose the extended mind thesis object to the very *possibility* of someone like Otto's having beliefs merely in virtue of his relationship to his notebook. They argue that Otto's predicament vis-a-vis the proposition that MoMA is on 53d street (call this proposition 'M') is sufficiently different from Inga's, so that we cannot regard him as believing that M.

Let me give a brief summary of two of the most prominent arguments for this claim. Adams and Aizawa (2010) argued that mental states need to have non-derived intentionality, and the record of M in the notebook lacks this. They further argue that defenders of the extended mind thesis make the mistake of confusing something that causally supports a mental state with a constituent of that state: entries in the notebook cause, and not constitute Otto's beliefs. Robert D. Rupert (2010) claimed that extended systems like Otto-plus-the-notebook are too unstable to form a persisting, unified system that deserves to be called 'cognitive'. Each of these arguments has been met by counterarguments. I cannot settle these debates, but I note that either or both of these considerations may give someone a reason to deny that Otto believes that M. In the next two sections I will argue

that notwithstanding this denial, one could still accept that Otto *knows that M*.

VII. Differences between belief and knowledge

In section V, we saw that it is *prima facie* plausible to attribute knowledge to Otto. In section VI, we presented some reasons to deny that Otto *believes* that M. Are these also reasons to deny that Otto *knows* that M? If someone assumes that belief is a necessary condition for knowledge, then the answer is yes. But we should not do this: the purpose of this paper is precisely to test this assumption. We should rather see whether the particular arguments involving belief also apply to knowledge. We will see that they do not, at least not straightforwardly. And the differences between knowledge and belief will further support the claim that extended mind cases could be usefully understood as cases of knowledge.

The debate about extended minds is often phrased as a debate about the *vehicles* of mental states. Clark and Chalmers argued that if the information-storing equipment of the notebook was placed inside Otto's head, we would not hesitate to regard it as part of a cognitive process, and hence, by parity, we should regard it as part of the cognitive process even if it is outside the head. The processes inside the head are naturally thought of as vehicles or constituents of the mental state, so the issue seems whether anything outside the body is a constituent of the mental state. This way of formulating the issue has puzzled some commentators. If Otto's mind is located partly in his head and partly in the notebook, does that mean that Otto's mind has big holes? That it is a scattered object? (Fodor 2009). One may feel that this kind of objections miss the point; but then it would be better to *reformulate* the question, for example in the following way: what conditions are necessary in order to attribute Otto a certain belief? This talk would switch from talking about constituents to talk about obtaining of certain necessary conditions. Once we start viewing the cases as cases of knowledge rather than belief, this switch makes even more sense.

It is generally accepted that states of knowledge involve relations to the external environment. Even those who resist various forms of externalism about belief would have little problem with conceding that knowing is an external state. Knowledge often involves appropriate relations to the environment, not only to the object of knowledge, but also to the environmental conditions that enable a subject to acquire knowledge. These relations are not expected to be entirely stable: conditions for knowledge may obtain for a while and then cease to obtain.

Take the example of seeing a picture. One's visual experience causally depends on the intensity and nature of light, but many would be reluctant to regard the source of light as part of an extended cognitive system (for example Rupert 2010, p. 329). The presence of light causally

contributes to certain states the perceiver has – a visual experience, and possibly a subsequent perceptual judgement – but on many theories, it is not constitutive of these states. If someone wanted to say that the presence of light is constitutive of a state like a visual experience or a judgement, they might be committing the fallacy of mistaking a cause for a constituent, which was Adams and Aizawa's charge against defenders of the extended mind thesis about beliefs.

In contrast, it seems to be *less* problematic to regard the enabling environmental conditions as necessary conditions for seeing, or perceiving, and thereby *knowing* that the picture is there, since perceptual knowledge obtains when certain conditions are in place. Mistaking a cause for a constituent does not apply to the state of knowing: many factors that merely causally contribute to mental states like perceptual experiences and judgements, are necessary conditions for having knowledge. We do not need to talk about constituents, but rather of obtaining of some necessary conditions, and in the case of knowing, these will involve conditions outside our organic bodies. Since these conditions are not anymore deemed to be constituents of mental states, we do not need to require that they exhibit original intentionality.

Even if external and causally efficacious conditions are often necessary for state of knowing, this may not be enough to dismiss the worries about attributing knowledge to Otto. Perhaps knowing is external in the sense that it involves relations to one's environment, but, someone may suggest, it isn't completely external: it has an internal component (or at least it entails having an internal component, even if the state is not the sum of an internal and external component). And in Otto's case, the suggestion continues, the notebook plays the role that an internal component would normally play. But the notebook isn't in fact internal, though it should be, if it is to make up or accompany a proper state of knowing. Hence Otto does not have knowledge after all.

I have two points in reply to this objection. First, I am not suggesting that non-conscious information stores like phones or databases *in themselves* can count as knowing agents. The only cases I consider as candidates for extended knowledge are human beings who are clearly capable of acquiring knowledge. They probably need to have cognitive abilities integrated into a cognitive character in the way it is described in section V. They will need to relate entries in databases to conscious events in the way it is explained in section IX below. So only agents with an appropriate general "internal" setup will be candidates for knowing in the first place. The proposal does not extend the realm of intentional agents beyond the agents currently acknowledged. It merely adds to the external conditions that enable these agents to have knowledge, and thereby extends the range of cases where the same agents possess knowledge.

Second, the idea that knowing requires having an internal state that has certain features not satisfied by the Otto-plus-notebook scenario is just the idea that belief (understood in a certain way)

is necessary for knowledge, expressed in a more general terminology. When I claim that belief may not be necessary for knowledge, I rely on a particular conception of belief: belief as a certain kind of internal state. So in fact, my point could be also expressed by saying that “a certain kind of internal state may not be necessary for knowledge”, though some *other* kind of internal state may be necessary, if we are moved by the considerations concerning cognitive ability (note that we saw that cognitive abilities can be awarded even in the absence of belief). Obviously, on other conceptions of belief, my point would not be valid: anyone who is convinced by Clark and Chalmers's original argument and is ready to attribute beliefs to Otto will not recognize his case as knowledge without belief. But this is well understood, and will apply to other putative arguments for separating knowledge and belief. For example, we saw also in the argument offered by Myers-Schulz and Schwitzgebel that you need certain theories of belief to demonstrate that a case involves knowledge without belief.

I started in section 1 with the observation that it is always by favor of Nature that we know something. This idea seems central to a widely accepted conception of knowledge. With the emergence of various forms of externalist theories about the mind, philosophers have become increasingly friendly towards the idea that at least in some cases, it is by favor of nature that we believe certain things. For example, it may be proposed that it is by favor of Nature that we have beliefs like water quenches thirst; our Twin-Earth counterparts could not have this belief, because Nature did not supply them with certain circumstances. However, these externalist suggestions are still encountering resistance. Moreover, even if some philosophers accept a certain version of externalism (for example about mental content), they may be reluctant to accept other forms, for example the type involved in extended mind scenarios. My point is that as long as we have agents with an appropriate internal constitution, we should expect *less* of this kind of resistance to the proposal that *knowing* something involves facts outside the subject: and this opens the way to interpreting extended mind scenarios as cases of knowledge without belief.⁵

VIII. Reliance on the world

Another worry, raised by Rupert and mentioned in section VI, was the lack of stability about the system of Otto-plus-the-notebook. Rupert thinks that in order to form a cognitive system, the components of the system have to have a stable connection to each other, and this is not satisfied by Otto and the notebook: the notebook is too easily detachable from Otto. The worry does not obviously arise for knowledge: a state of knowing has “components” that do not have a stable

⁵ I cannot see a similar move in the case of externalism about mental content, so my remarks here are confined to the kind of externalism involved in Otto-type cases.

connection to each other. My current state of knowing that a vase is on the table plausibly depends, among other things, on the vase and the presence of light. The vase could be removed, the light in the room could be switched off, and the state of knowledge would cease to obtain. But this does not cast doubt on the claim that when all components *were* in place, the state of knowledge did obtain.

It is true that knowing may require other sorts of stability. This is expressed, for example, in various conditions like sensitivity or safety: beliefs should track the truth, and they should not easily go wrong. If the world was in so much flux that we needed a large amount of luck to have true beliefs, we would not have knowledge.

This observation raises an interesting point. Something like sensitivity and safety seem to be applicable to Otto's case as well, independently of whether we attribute beliefs to Otto or not: the notes in the notebook have to be kept in such a way that they track truth and do not easily go wrong. For example, Otto has to make sure that he keeps his notebook up-to-date; consulting a 40 year-old entry for someone's current address may not be a good idea. The idea that records track the truth or would not easily go wrong is perfectly comprehensible, even though we made no mention of beliefs. Now, given that we are questioning that Otto has the relevant beliefs, the conditions cannot be put in the usual terms, with reference to features of belief. However, it is quite plausible that the very same ideas of sensitivity and safety could be translated to terms adequate for Otto's situation; so some notebook or smartphone owners will lack knowledge, while others will have it, depending on whether they satisfy these conditions. This seems to show that the state of a subject in an extended mind scenario, even without belief, could satisfy the conditions of stability that we require from knowledge.

Hence, *prima facie*, at least some of the reasons to claim that Otto lacks the belief that M are not, at the same time, reasons to say that Otto lacks the knowledge that M (unless we already assume that belief is a necessary condition for knowledge.)

IX. Extending too far

There is at least one natural objection that people raised against the original extended belief thesis which applies equally to the proposal to treat extended mind scenarios as cases of knowledge. The objection is that once we allow to extend knowledge or belief to the notebook, we will be unable to stop it from extending too far. Suppose that we accept that Otto has knowledge in virtue of having the information stored and available in his notebook. Doesn't it follow that when Otto visits a library, or browses the internet, his knowledge suddenly extends to all the information contained in the books and on billions of webpages? And isn't that a *reductio* against all theses about extension?

Any account of extended mind cases must be able to answer this challenge: it must provide

some conditions that stop extending knowledge or belief too far. Clark and Chalmers propose that while there may be no categorical answer to the challenge, focusing on the central features of Otto's situation may offer a guidance. They list four such central features: one, that Otto constantly relies on the notebook: if there is a situation where information in the notebook is relevant, he rarely takes action without consulting the notebook. Second, he has easy access to the contents of the notebook. Third, content retrieved from the notebook is automatically endorsed by him. Fourth, everything in the notebook was consciously endorsed by him at some point in the past, and is in fact in the notebook because of that endorsement. The fourth condition, as Clark and Chalmers note, may in fact be too strong, and may even disqualify ordinary cases of belief. It will disqualify also one of my own examples given above: the person in the directory enquiries knowing when the trains leave. So the fourth condition may have to be abandoned or refined in some way.

The interesting thing about the four features is that they all have to do with how entries in the notebook (or information stored in the brain) relate to episodes in the stream of consciousness. Decisions are preceded by retrieval of information, accessing information implies that a content is entertained in the stream of consciousness, endorsement upon retrieval or in the past are conscious acts. The whole problem of extension arises only with respect to standing states of subjects. However, the need to stop extension going too far shows that standing states must retain an intimate connections to episodes in the stream of consciousness.

One clarification may be in order. As I already mentioned, the whole proposal applies only to beings who are sufficiently similar to us, in that they have a stream of consciousness and standing states, and the point is that given that we do have consciousness, it has the significance for knowledge indicated above. One could imagine an information processing creature which displays analogous phenomena without consciousness. The creature stores information in various forms, and the information can be mobilised for active computations, which would then result in some form of "acceptance" – for example, displaying a result on a screen. So the claim is not that such a structure would be impossible without consciousness. The claim is rather that given that we are in fact endowed with consciousness, the episodes in the stream of consciousness are recruited to play a certain role.

X. Objection: knowledge is factive

The following objection may be brought against the suggestion that subjects in extended mind scenarios have knowledge: as the example is constructed, there is nothing to guarantee that Otto gets the *correct* information from his notebook: he may have entered a mistaken address, say that MoMA is on 54th street. The same goes, of course, for actual extended mind cases. However – this

objection continues – even if Otto had an incorrect entry, his relationship to the proposition that MoMA was on 54th street would be importantly similar to his relationship to the correct proposition in the original example: he relies on this proposition in his actions, he gives it as answer to relevant questions, etc. The causal role of the entry is hence the same whether it is correct or not, and this is precisely the reason why the attitude should be called belief, rather than knowledge, since belief does not entail correctness – so goes the objection.

Notice that exactly the same objection can be brought against Myers-Schulz and Schwitzgebel's proposal or Radford's example. As the examples are constructed, the propositions that are candidate objects for knowledge without belief are, of course, all true. But relevantly similar cases could be easily designed with false propositions. The cases would be relevantly similar in that the protagonists would manifest the same dispositions in relation to the false propositions, so the false propositions would plausibly play the same role in explaining their actions. And this, one may suggest, supports the claim that these are cases of belief, and not knowledge.

The answer to the objection is as follows (and the same answer is available also for the scenarios presented by Myers-Schulz and Schwitzgebel). We do indeed need a name for the state of someone who is in an extended mind scenario, but the information available to her is mistaken. In some other languages, there is an expression that is suitable for covering this case: for example in Hungarian, one would say that “Otto rosszul *tudja*”, which literally means that Otto *knows* something “badly” or perhaps “mistakenly”. Occasionally, we can find the same use in English. Witness the following description from Ellis Peters' *The Holy Thief*: “He had always (...) this conviction that he owns the girl, *because he believes, no, God be good to him, poor soul, he mistakenly but surely knows!* that he brought her here. God forbid that he should ever find out the truth, that she is far away in her own chosen place ...” (emphasis added). “Mistakenly but surely knowing” is state that the protagonist, and at least some of those in her environment would be likely to take as knowledge, if asked; a state that from the internal point of view, plays the same role in the subject's life as knowledge does, with the difference that the proposition in question is false.

It would require further argument to show that this state is a state of believing, which is a common factor to the veridical and the falsidical case. There is a similar issue discussed in relation to disjunctivist theories of perception. Take a veridical perception of an object, and the corresponding hallucination which is indiscriminable from the perceptual state but lacks veridicality. It requires a further argument to show that there is a kind of mental state, a visual experience that is a common factor to these two cases._

XI. Conclusion

It is natural to say that in some Otto-type situations, we have knowledge. There is a conception of knowledge that supports this practice: knowers are sources of reliable information, who can integrate their storage of information within their cognitive character, where within contextually defined boundaries, it does not matter how the information is retrieved, and what external aid is used in the process. Many agree that knowing depends on external relations and the obtaining of appropriate conditions, so in Otto type cases, there is no prima facie reason to deny knowledge from Otto just because this depends on an appropriate causal connection to an external enabling condition, that is, the presence of his notebook. The entries in the notebook can exhibit a number of features that are widely thought to be important for knowledge, such as sensitivity and safety. At the same time, there are various reasons to hold that in these cases, subject lack the relevant beliefs. Hence we may have cases where knowledge is possible without belief.

XII. Coda: possibly an additional argument for the extended mind thesis

Throughout the paper, I concentrated on arguments that would motivate attributing knowledge rather than belief to the protagonists of “extended mind” scenarios. I started with conceptions of knowledge that do not rely on belief as a necessary condition for knowledge, and emphasised how Otto-type cases fit the widely accepted notion that knowledge depends on external enabling conditions. The idea was to establish that extended scenarios involve knowledge rather than belief, by relying on features of knowledge that are commonly thought to distinguish it from belief. However, once we get to this point, the argument could also proceed further in the opposite direction (calling in mind the observation that one philosopher's modus ponens is another philosopher's modus tollens). Suppose someone accepts that when I answer your question of whether I know NN's phone-number, I speak truthfully, because I am a reliable source of information on this matter. But this person also feels strongly about the connection between belief and knowledge: some sort of acceptance attitude, she may hold, is necessary for having knowledge and in fact this is what belief is; or she may think that “belief” is the best name for the attitude that is like knowledge except it's not factive. Consequently, this person has now an extra reason to accept the extended mind thesis: since Otto has knowledge, and belief is necessary for knowledge, she will conclude that Otto has the relevant beliefs. Hence we have a further argument for the extended mind thesis.

References:

- Adams, Fred and Ken Aizawa (2010) "Defending the Bounds of Cognition" in Menary (2010): 67-80
- Clark, Andy and David Chalmers 1998 "The Extended Mind" *Analysis*, 58, pp. 7-19
- Craig, E. (1990), *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, Oxford: Clarendon Press.
- Fodor, J. (2009) "Where is my mind?" *London Review of Books*, 31(3), 13-15.
- Gendler, Tamar Szabó (2008). 'Alief and Belief', *Journal of Philosophy* 105, pp. 634-663.
- Gertler, Brie (2011). 'Self-Knowledge and the Transparency of Belief', in Anthony Hatzimoysis (ed.), *Self-Knowledge*. Oxford: Oxford.
- Menary, R. (Ed.) (2010). *The extended mind*. Cambridge, Massachusetts: The MIT Press
- Myers-Schulz, Blake and Eric Schwitzgebel, 2013 "Knowing That P Without Believing That P" *Nous* 47, pp. 371–384.
- Pritchard, Duncan (2010) "Cognitive Ability and the Extended Cognition Thesis", *Synthese* 175 (2010), pp. 133-51
- Radford, C. (1966). "Knowledge – By Examples," *Analysis*, 27, 1-11.
- Rupert, Robert D. (2010) "Representation in Extended Cognitive Systems: Does the Scaffolding of Language Extend the Mind?" in Menary (2010) pp. 325-54
- Schaffer, J. (2007), "Knowing the Answer", *Philosophy and Phenomenological Research*, Vol. LXXV No. 2, 1-21.
- Schwitzgebel, Eric (2010) "Acting Contrary to Our Professed Beliefs" *Pacific Philosophical Quarterly*. 91 531–55
- Williamson, Timothy (2000) *Knowledge and Its Limits*, Oxford: Oxford University Press.
- Wittgenstein, Ludwig (1969-1975) *On Certainty* ed. G.E.M. Anscombe and G.H. von Wright
Translated by Denis Paul and G.E.M. Anscombe. Basil Blackwell, Oxford