

THE RELIABILITY OF MEMORY: AN ARGUMENT FROM THE ARMCHAIR

Ali Hasan
ali-hasan@uiowa.edu

ABSTRACT

The “problem of memory” in epistemology is concerned with whether and how we could have knowledge, or at least justification, for trusting our apparent memories. I defend an inductive solution—more precisely, an abductive solution—to the problem. A natural worry is that any such solution would be circular, for it would have to depend on memory. I argue that belief in the reliability of memory can be justified from the armchair, without relying on memory. The justification is, roughly, that my having the sort of experience that my apparent memory should lead me to expect is best explained by the hypothesis that my memories are reliable.

My solution is inspired by Harrod’s (1942) inductive solution. Coburn (1960) argued that Harrod’s solution contains a fatal flaw. I show that my solution is not vulnerable to Coburn’s objection, and respond to a number of other, recent and likely objections.

*It seems to me I've heard that song before
It's an old familiar score
I know it well, that melody*
—Frank Sinatra

1. The Problem of Memory

Consider the following skeptical argument. Many, arguably most, of our beliefs depend on memory for their justification, either directly or indirectly. In order to be justified in believing something on the basis of memory, I must have independent justification for thinking that memory is reliable or trustworthy—i.e., some reason to think that my memories are correct or likely to be correct, or that they are good guides to the truth regarding the past. Perhaps I can construct an argument to the effect that my memories and apparent perceptions cohere very well with each other, and that the best explanation for this coherence is that these memories and apparent perceptions are more or less correct.¹ But even if there is a good argument or inference of this sort, it is bound to be rather complex, complex enough that I would not be able to use it without relying on

¹ Brandt (1955) offers an argument of this sort.

some of my memories.² And so, that we could have some inferential justification for belief in the reliability of memory without relying on memory seems implausible; some have even complained that they “have no idea how such an inference would go” (Huemer 1999: 347). But any attempt to justify one’s reliance on memory by relying on memory seems viciously circular. It is tempting to conclude that there is no justification for relying on memory, and given the extent of our reliance on memory, this threatens a radical, near-global, skepticism.

There are different anti-skeptical responses to the argument. Perhaps the most popular is that memory can provide, or preserve, justification for a belief, without one’s having any reason to think that one’s memories are trustworthy.³ A second is that the sort of circularity involved is benign, not vicious. And a third is that beliefs regarding the reliability or correctness of memory are somehow non-inferentially justified, or are “entitlements” of some sort that can make a difference to justification.

I provide a different response: there is an abductive justification for belief in the reliability of memory that does not depend on memory. This should be of interest to those traditional foundationalists who accept that some sources—introspection and reason—can be sources of foundational justification, but deny that other sources like perception, memory, and testimony have a similar status. Such views are often criticized on the grounds that they lead to skepticism regarding perceptual and external-world beliefs, but the problem of memory is an even more fundamental threat. While some have developed responses to external-world skepticism that appeal to inference to the best explanation (IBE) or abduction, these responses seem to all depend on the subject’s having justification for trusting her memories regarding her own experiences.⁴ The solution should also be of interest to those (whether they identify as traditional foundationalists or not) who hold that justification ultimately depends on some kind of explanatory fit or explanatory coherence between one’s beliefs and one’s experiential or conscious mental states. But it should be of interest to others as well, for it is surely an interesting question whether we do or even could have non-circular reasons to think that memory is reliable, even if we doubt that such reasons are required for justification.

To my knowledge, R. F. Harrod’s “Memory” (1942) was the first serious attempt to provide a non-circular, inductive solution to our problem. Robert Coburn (1960) argued that Harrod’s solution contains a fatal flaw. I present a similar solution and explain why it is not vulnerable to Coburn’s objection. I also respond to more recent objections from BonJour (2010), Huemer (1999), and Moon (2017), and other likely objections. I conclude by considering the worry that our beliefs are rarely based on the sort of argument given here. My main goal is to show that an account of memorial justification along these lines is worthy of serious consideration.

² See BonJour (2010). As he puts it, we depend on memory for “assembling and keeping track of the resources for a justification of any but the simplest, most immediate beliefs” (169-70).

³ On this view, the *absence* of a reason to think that one’s memories are *unreliable* might be required, but not possession of or access to a reason to think that memories are reliable.

⁴ See for example, McGrew (1995), BonJour (1999), Vogel (2008), McCain (2014), and Hasan (2017). My task here is not to defend IBE or abduction as a form of inference, but to show that if it is an *a priori* justified form of inference (as proponents of the IBE response to skepticism tend to hold), then it can be used to defeat skepticism about memorial justification as well.

2. BonJour on the Problem of Memory

Before considering Harrod's solution and defending my own, I want to discuss BonJour's response to the problem. This will help make the problem more vivid, especially as it arises for the traditional foundationalist, and show that certain tempting ways to respond to the skeptical problem, or limit its significance, are unsuccessful.

Like other traditional foundationalists, BonJour denies that memory can be a source of basic or foundationally justified belief. This is because memory is not "justificationally self-sufficient" in the way that introspection and reason arguably can be, the basic idea being that the latter sources themselves involve or bring with them a reason to trust their deliverances, whereas perception and memory do not (BonJour 2010: 170). Traditional foundationalists might elaborate on and defend this in different ways. Some, for example, would argue that only in the case of introspection and reason do we ever have direct awareness of or acquaintance with what our beliefs purport to be about.⁵

BonJour does not provide what we might call a direct solution to our problem, i.e., one that explains how we could have or acquire some reason to think that memory is reliable. He does not deny that some particular beliefs about the past could be justified by inference from, or because they cohere with, other beliefs. But he claims that any such justification would still depend on memory, at least of a fundamental sort—the sort needed for "assembling and keeping track of the resources for a justification of any but the simplest, most immediate beliefs" (BonJour 2010: 160-70). In response to the question of whether we are left with skepticism, he concedes: "In a way, I think that we are" (BonJour 2010: 170). But he seems to want to limit or contain the scope and significance of this concession. He says that this kind of skepticism is "best viewed not as challenging the justification of particular beliefs," but as challenging "our very existence as temporally integrated cognitive agents"; it is the former "narrower" sorts of issues with which "epistemologists have primarily been concerned" (BonJour 2010: 171).

Why should we think that this kind of skepticism is limited or can be contained in this way? Indeed, if we do not have justification for thinking that we are "temporally integrated cognitive agents" with memories that are to some extent reliable, why shouldn't this threaten our justification for particular beliefs that depend on memory?

A skeptic about induction or about the external world or about other minds questions whether the specific sorts of evidence that we have for those kinds of claims is really enough to make it likely that they are true—with the implication being that we should perhaps stop accepting those claims on that sort of evidence. But skepticism about the most fundamental sort of memory does not really so much imply that there are certain sorts of claims that we should cease to accept, since to follow and know that we are following such a policy would itself depend on memory. ... Instead, it is better viewed as challenging whether we are indeed the sorts of ongoing cognitive agents, integrated over time, that could deliberately follow any such policy or have good or even bad reasons for most of their beliefs. [BonJour 2010: 170-1]

Philosophers sometimes claim that skepticism about memory—the view, roughly, that we are not justified in trusting or relying on our memories—is epistemically self-defeating: if

⁵ For discussion of such views, see Hasan and Fumerton (2017).

it is true then one cannot be justified in accepting it, for one cannot grasp the entire skeptical argument (such as the argument presented in the first section above) “at once” and so must rely on memory. Bonjour seems to be saying something very similar: if skepticism about memory is true then it does not follow that one should—presumably, in an epistemic sense of “should”—stop holding particular beliefs that are based on memory. Why? Because this “should” applies only if one knows, or at least is justified in believing, that memories are not to be trusted and that these particular beliefs are based on memory. But these beliefs about memory would themselves arguably depend on memory. So, even if skepticism about memory is true, it does not follow that we would be justified in heeding its normative implications by suspending beliefs that are based on memory, or that we should (in an epistemic sense) give up these beliefs.⁶ In contrast, skepticism about the external world perhaps does imply that we should (in an epistemic sense) cease believing certain things on the basis of apparent perception. Managing and revising our beliefs in accordance with external world skepticism might require the use of introspection, reason, and memory, but if perception is not required then this handling of our doxastic states might be justified.

Let us grant for the sake of argument that Bonjour is right that, in the sense just discussed, skepticism about memory is self-defeating, and that he is right that the other forms of skepticism—about induction, about the external world, and about other minds—are not.⁷ Even so, skepticism about memory does challenge the justification of beliefs that are based on memory. First, the fact that some form of skepticism is self-defeating does not show or give us a reason to think that it is false, or that any argument for it must be unsound. That skepticism about memory is self-defeating does not provide us with a reason for thinking that we are, as Bonjour puts it, “ongoing cognitive agents,” or that our memories are trustworthy. Second and more importantly, even if the skeptical argument is epistemically self-defeating, if we find ourselves accepting the premises of the skeptic’s argument and find no reason to reject the skeptic’s reasoning, it is hardly unproblematic to go on accepting beliefs on the basis of memory. Fumerton (1995: 50-1) gives a helpful example to illustrate the problem. Suppose there is a culture that takes the eight ball to be a reliable guide to the truth. Suppose that a skeptic about this method of belief—call it “eight-ball reasoning”—gets the idea to ask the eight ball whether eight-ball reasoning is reliable, and the eight-ball answers “No.” Let’s add that the skeptic has no other, independent reasons to reject the reliability of eight-ball reasoning. The eight-ball reasoners might point out that the skeptic’s position here is epistemically self-defeating: if the conclusion is true then the skeptic is not justified in accepting that conclusion. But it is obvious that there is a problem with the eight-ball

⁶ Why does Bonjour say that memory skepticism *does* challenge the view that we are “ongoing cognitive agents”? Wouldn’t an attempt to argue that we are *not* such ongoing cognitive agents also depend on memory, with the result that we would not be justified in suspending belief that we are such agents? Bonjour says that memory skepticism in some sense questions or challenges “our very existence” as ongoing cognitive agents (2010: 171); he does not explicitly say that it challenges our justification for thinking that we exist as cognitive agents. But it’s not clear to me what other sense there could be to the idea that the skeptic challenges our existence in this context.

⁷ One might plausibly argue in a parallel fashion that skepticism about induction does not so much imply that there are particular claims we should not accept, since to follow and know one is following such a policy would itself depend on induction. But let’s set this worry aside.

reasoners' position, and that they can't simply ignore the fact that the eight-ball delivered a negative verdict on its own reliability.

As we have seen, BonJour denies that memory is an adequate, "self-sufficient" source of foundational belief in the way that introspection or a priori reason might be. But perhaps BonJour could change his mind about that. Perhaps the fundamental sort of memory that we rely on to keep track of our own experiences and reasons functions, like consciousness or direct awareness, as a fundamental form of access. For the access internalist, one must have access to the relevant sorts of reasons; one need not also have access to the fact that one has access to these reasons.⁸ If this is right, then the fact that one has access, via memory, to one's reasons for belief need not itself be something to which one has access. BonJour does elsewhere characterize memory as playing an essential role in providing or preserving access to the premises of one's arguments or reasoning (1998: 126). Perhaps he can offer something like this as an account of the epistemic role of memory, and deny that one must have a reason to think one's memory is reliable.

There remains a significant tension, however, between this view of the epistemic role of memory and the traditional view (which BonJour accepts) of the epistemic role of perception. If we allow our (at least veridical) memories to provide a kind of access to past experiences or thoughts and allow them to play this role independently of any justification for trusting memory, why not allow (at least veridical, non-hallucinatory) perceptual experiences to provide a kind of access to the world, independently of any justification for trusting perceptual experience? There is no clear rationale for including memory as an epistemically fundamental form of access on a par with acquaintance or direct awareness, while denying that perception is also a fundamental form of access.

In light of these problems, it is not clear what account of our justification for relying on memory is being offered by BonJour, and if the purpose is not to offer such an account, why we should think the skeptical problem does not have the force and scope it appears to have. In any case, BonJour clearly does not offer what I have called a direct solution to the problem, i.e., a (non-circular) reason to think that memory is reliable or trustworthy—indeed, he denies that any such solution can be given. Let us now turn to Harrod's attempt to provide such a solution.

3. Harrod's Argument and Coburn's Objection

My own solution to the problem is inspired by Harrod's (1942), so it makes sense to begin with a brief presentation of his argument and Coburn's (1960) objection to it. Harrod's attempt to justify the reliability or "informativeness" of memory depends on the following claim: "If we start from the position that any one of an infinite number of things may happen at any moment, the fulfillment of a particular thing verifies in high degree the grounds for the prediction" (1942: 64). So, if the hypothesis that my memories are reliable supports a relatively specific prediction about my experiences, one that conflicts with infinitely many other things I could have experienced, and the prediction is fulfilled, then that is very good evidence that memory is reliable. Let's consider, then, whether there are any such cases of fulfilling a prediction that would provide evidence for the reliability or trustworthiness of memory.

⁸ Some internalists might hold the stronger requirement that one have access to the fact that one has access to good reasons. But this is arguably too demanding, and many contemporary internalists would deny it.

Let's start with an example that is problematic but instructive. I have a vivid memory of sitting in a living room for some time now, looking at a particular arrangement of furniture and house plants, and listening to the sound of crickets. I make a prediction, based on the hypothesis that my memories are reliable or generally correct, that I will continue to see and hear these things. Turning my attention back to my surroundings, I find that I do indeed see a similar arrangement of furniture, and hear crickets. The prediction is fulfilled, and so, the hypothesis that my memory is reliable is confirmed.

There are at least two potential problems with such examples, problems discussed by Harrod and Coburn themselves. First, it seems that I am relying on beliefs that depend on memory for their justification, which renders this attempt to justify the reliability of memory circular. My belief that I am perceiving these items of furniture, house plants, etc., and am not merely dreaming, arguably depends in part on memory for its justification. Moreover, the claim that the correctness of my memory of being in the living room makes probable that I find myself in the living room seems to depend on background assumptions—e.g., that such things as walls and furniture tend to persist (unlike, e.g., sneezes and snowflakes), that no one stealthily moved me or these objects, etc. These background assumptions would seem to depend on memory for their justification.

We might try to avoid this problem by focusing on memories *of or about one's experience*, and bracket the question of whether these or any other experiences ever correctly represent the external world.⁹ Let us assume that I have no knowledge or even justified belief about the existence of an external world; the relevant hypothesis is that my memories are informative or reliable *with respect to my past experiences*. However, without the above background assumptions, why should we think that the truth of the hypothesis that my memory is reliable makes certain predictions probable? Harrod's answer is, in short, *by induction*.¹⁰ The central idea is summarized nicely by Coburn:

...the probability of any belief concerning the continuation into the future of any characteristic type or pattern of experience will be directly proportional to the duration of the type or pattern of experience and inversely proportional to the length of the predicted continuation—and this is irrespective of the credibility of any other beliefs concerning matters of fact. [Coburn 1960: 83]

If my memory is correct, and so, I have indeed been undergoing living-room-type experiences for an extended period of time, then, in line with induction, it is highly likely, *a priori*, that I continue to have a similar experience, at least for a short while.

A second problem is that the argument might still depend on memory, and so be circular, for another reason: In order to be justified in trusting my memory by now observing the fulfillment of a prediction I had made, wouldn't I have to rely on my

⁹ Many of our memories—at least many of our “episodic memories”—could be understood as memories *of experiences* of one kind or other, and we could try to justify our reliance on episodic memories as a guide to past experiences, bracketing for now the issue of whether these experiences ever correctly represent the external world, as well as the status of other sorts of non-experiential or non-episodic, “semantic memories” (e.g., I once learned and now “remember” that the American Civil War lasted from 1861 to 1865). I discuss these distinctions briefly below, in the next section.

¹⁰ Harrod (1942) takes induction to be justifiable *a priori*.

present memory of having made that prediction? If the goal is to avoid relying on memory in justifying my reliance on memory, this doesn't help. According to Harrod, this difficulty can be avoided because it is possible to entertain the *prediction* and its *fulfillment* in the "specious present":

It appears to me that predications are often made and fulfilled within a specious present, so that there is no need to resort to memory to find instances of predictions fulfilled. For instance, I may predict continuity for the ordinary objects around me and the prediction may be fulfilled within one present totality. And if a flash of lightning occurs, I may predict and experience its end within the same totality. Only if memory is informative should I have grounds for predicting continuity for the chairs and tables and death for the lightning flash; the success of the prediction within the specious present is therefore strong evidence of the informativeness of memory. [Harrod 1942: 65]

Coburn admits that "it is by no means obvious...that the specious present is not really long enough to accommodate both a predication concerning the continuation of a sound, say, and the fulfillment of such a predication" (1960: 83).

Coburn objects, however, that some skeptical or delusive hypotheses—i.e., hypotheses that entail that one's memory is mistaken—would also be confirmed, even more strongly, by the same experience.¹¹ Suppose that I have been experiencing a specific sort of sound for regular ten-minute intervals, separated by ten-minute-long intervals without this sound, and that this pattern started exactly ten hours and thirty seconds ago. And, in line with this pattern, I started experiencing the sound again thirty seconds ago. Let's call this hypothesis SK. Suppose, however, that I have a mistaken memory that I've been experiencing this arrangement for several minutes. SK, if true, would make it highly probable that I experience the sound now—indeed, if we follow Harrod's reasoning, it would make this even more probable than would the hypothesis that my memory is correct and that I have been experiencing the sound for several minutes. For recall that, according to Harrod, the probability of any belief regarding the continuation of any pattern of experience will be directly proportional to the duration of the past pattern of experience and inversely proportional to the length of the predicted continuation. So, experiencing the sound in the present confirms SK more strongly than it confirms the hypothesis that my memory is informative. Indeed, given any experience that would seem to confirm the hypothesis that memory is informative, we can come up with another hypothesis that entails that memory is delusive and that the experience confirms to a higher degree, if we follow Harrod's reasoning. Harrod's solution thus contains an "obvious and quite serious flaw" (Coburn 1960: 82).

Coburn briefly discusses the attempt to rule out hypotheses like SK by appeal to simplicity: the hypothesis according to which the sound stretches for a short interval of time prior to the present moment is simpler than the one according to which the sound is experienced for ten-minute intervals over several hours. But, he complains, "it is unclear how this approach could be worked out in such a way as not to eliminate hypotheses concerning the reliability of ostensible memories..." (1960: 84). Many of our memories are, after all, memories of complex experiences.

¹¹ For a more recent critique of Harrod that raises the very same objection, among others, see Bernecker (2008: 99-100).

4. The Abductive Solution

There are at least two respects in which my argument, while inspired by Harrod, is different: First, my argument takes the form of an *inference to the best explanation*, and while Harrod's might suggest such an inference, he does not explicitly present it as such an inference. Second, my argument depends crucially on taking one's apparent memory to be part of the data, while Harrod's does not. The basic argument form is this: I have apparent memory M, and of all the infinitely many experiences I could have, I am having just the sort of experience E that I would have if my apparent memories were reliable. The best explanation of my having M and E is that my apparent memories are reliable. So, my apparent memories are reliable. When I say that apparent memories are "reliable" or "trustworthy," I just mean that they are probably correct or likely to be correct,¹² or, if you prefer, that they are good guides to the truth regarding the past.

Before defending this solution, I should clarify that I do not intend the argument to apply directly and equally to "memory" in every sense of the term. As others have pointed out, talk of 'memory' simpliciter can obscure the importantly different kinds of memory, or distinct senses of the term (Frise 2018). Philosophers and psychologists commonly distinguish between "declarative memory" and "procedural memory" (roughly, memory of information and events vs. skill-related memory or know-how); there's a difference, for example, between remembering or seeming to remember that there is a guitar in the next room, and remembering how to play it. We are concerned here with declarative memory. Under declarative memory, we can distinguish between merely semantic and episodic memory. Making the distinction precise can be tricky, but an intuitive grasp of it is good enough for our purposes. My memory that the Ancient Greek word for guitar is 'kithara' and that that Abraham Lincoln was assassinated are examples of merely semantic memory, not episodic memory. Even though the latter example is in a sense memory of an episode, it doesn't count as episodic memory: I do not remember Lincoln's being assassinated, only that he was assassinated. In contrast, my memory of having eggs for breakfast this morning and of seeing the guitar in the next room are cases of episodic memory.

Harrod is focusing on episodic memories, and so am I. More specifically, I am interested in justifying reliance on a kind of mental state with two features: some, possibly faint, experienced quality, structure, or content, and something like a conscious propensity or inclination to take this quality or structure to have already been experienced, or to take the content to be a correct representation of some past experience or conscious state. To borrow a phrase from Harrod, memory, in the sense in which I am interested in here, is something like *a strong propensity to predicate liveliness to faint elements in experience* (1942: 57). It is natural to call these *apparent* memories, to emphasize both that it is a conscious state and that it could be mistaken, but for simplicity I sometimes drop the qualifier. In what follows, when I talk of *memory* or *memories*, unless I state otherwise, this should be taken to stand for this familiar, conscious propensity or inclination.

I want to leave open whether there is something underlying this conscious inclination, or entirely distinct from it, that we might also want to call a memory, and what might justify our reliance on it. I want to leave open whether any of my non-

¹² On some interpretations of probability, this will be understood in terms, roughly, of the ratio or frequency of correct apparent memories in the class of all apparent memories. But there are other interpretations of probability, and discussing which is correct, or epistemically relevant, would take us too far afield.

episodic, semantic memories are reliable.¹³ And I want to also leave open whether there is a specific cognitive system, or set of them, that might be called “memory,” and what might justify our reliance on such a system. The core argument is thus limited in scope. The task is to explain how we could have justification for believing that apparent memories, the sort of states characterized in the previous paragraph, are likely to be correct. We can bracket questions regarding the veridicality of experiences regarding the external world, as well as questions about the existence or reliability of cognitive systems. If the argument provides some degree of justification for relying on apparent memories, for thinking that these states are likely to be correct or to be a good guide to the truth, we will have some reason to trust our ability to keep track of our own experiences and thoughts over time. We might then attempt to build on this and acquire justification to believe the conclusions of more complex arguments that depend on apparent memories, arguments that might further raise the degree of justification we have for trusting these memories and, ultimately, provide justification to trust semantic or other kinds of memory.

Let us consider some examples. I’ll start by explaining how I could acquire justification for trusting a particular memory, and then explain why the justification generalizes—i.e., why a single instance can provide justification for thinking other memories are likely to be correct. This will allow us to articulate and directly address likely concerns with the argument.

I seem to recall experiencing what sounds like the chirping of crickets for some time now, up to the present. Take this memory to concern the phenomenological or internal features of the experience only; we can suppose that I know nothing about the actual external cause of such experiences.¹⁴ If my memory of this experience is correct, then (given induction) it is probable that I experience the chirping sound in the present, and highly improbable otherwise. At the same time, I find that I do indeed experience a chirping sound. So, the fact that, of all possible experiences, I am having just the sort of experience that my memory should lead me to expect gives me a good reason to trust my memory on this occasion.

Here is a second, slightly more complicated example. I seem to recall hearing a verse of an apparently familiar song, and I expect it to continue in a specific way. Once again, we can assume that I know nothing about the external world, and take the memory to concern my experiences only, and not the existence of any performer, performance, musical instrument, radio, etc., outside my mind. If this song is indeed familiar and my memory of hearing a verse of it is correct then it is probable that the song continue in just the way I expect, and improbable otherwise. At the same time, I find that it does indeed continue in this way, and this gives me a good reason to trust my memory. In this case, it gives me some reason to trust both that I just experienced a verse and that I’ve experienced this song before. In some cases, the beginning part and the continuing part of the familiar pattern could both be experienced in the specious present. For example, upon hearing the very beginning of Beethoven’s *Fifth Symphony*, I might

¹³ Aren’t some semantic memories needed to understand the words and concepts involved in the argument I am presenting? And if so, isn’t reliance on semantic memories inevitable? I address this objection, among others, in section 5.

¹⁴ Indeed, as it turns out, I thought I was hearing crickets when in fact they were little frogs that sound a lot like crickets. But none of this matters for the argument.

seem to recall that it continues in *this* specific way, and I also experience it continue in just that way in the specious present. This gives me a reason to trust that I've experienced this musical pattern before.¹⁵ These musical examples (and Harrod's example of the lightning flash in the passage quoted in the previous section) illustrate that support for the correctness of a memory by appeal to a present experience is not limited to experiences of continuity or sameness, and not limited to memories of experiences had just prior to the present. But, for simplicity, I'll often focus on cases in which the ostensible memory is of a recent experience that is similar to the present one.

It is possible, but implausible, that this sort of match between the memory and present experience in such examples is a mere coincidence. The best explanation seems to be that the memory is correct—and not merely *accidentally* correct of course, but correct due to some real connection between the past experience and the present memory, or else we would merely be replacing one coincidence with another.

This is a good point to consider Coburn's objection and see how the present solution fares against it. Recall the skeptical hypothesis SK, according to which I have been experiencing a specific sort of sound for regular ten-minute intervals, separated by ten-minute-long intervals without this sound, for ten hours and thirty seconds. According to Harrod, the probability of any belief regarding the continuation of any pattern of experience will be directly proportional to the duration of the past pattern of experience and inversely proportional to the length of the predicted continuation. So, experiencing the sound in the present confirms SK more strongly than it does the hypothesis, which is in line with what I seem to remember, that I have been experiencing the sound for a few minutes. We can reply, however, by pointing out that sensory experience is not the only relevant data. I have the *apparent memory* of having a similar experience for some time now, and can access both the sensory experience and this memory in the specious present. SK explains *why I have a sensory experience E*, but it does not explain *why I experience what my apparent memory (if correct) should lead me to expect*; it makes probable that I have *these sensory experiences*, but not that I have *the sorts of experiences I seem to recall*. In raising his objection against Harrod, Coburn does not consider taking the memory to be part of the data to be explained. To be fair, this is probably because, as we have seen, Harrod's argument explicitly involves comparing "a prediction" with its "fulfillment" and does not, at least not explicitly, take one's memory to be part of the data. Coburn's skeptical hypothesis thus fails to account for the data. The objection might work against Harrod's argument, but not against a version of the argument that takes the memory to be part of the data.

A worry might remain that is related to Coburn's objection and his consideration of the appeal to simplicity. We can put it in the form of a dilemma: either the relative simplicity of hypotheses is relevant to deciding between them, or it isn't. If it isn't, then we can still come up with delusive hypotheses that are more strongly confirmed by the experience, even while taking memory to be part of the data. All we have to do is complicate our hypothesis so that, given the hypothesis, it is very highly likely that I have this experience-memory pair. If on the other hand simplicity is relevant, then we won't be able to justify relying on complex memories, memories of complex patterns of experience. My response is to grasp the second horn. Perhaps, other things being equal, memories are less likely to be correct the more complex the experience apparently recalled. But this is compatible with the argument's applying to some complex memories. What is given in

¹⁵ See Harrod (1942: 65) for a similar example.

one's experience (and this includes one's apparent memories) can be complex, and complex data might require a more complex hypothesis to account for it. Examples like those involving memories of songs or pieces of music discussed above are cases in point. Moreover, even if the argument of this section (including the attempt to generalize to some other memories coming up below) gives us little or no reason to trust very complex memories, we might be able to acquire more justification by a more complicated inference, after we have some minimal justification for relying on some of our memories.

Perhaps other skeptical hypotheses can do better than SK in accounting for the match between memory and experience in the specious present. But there is a general reason to think they are inferior. The non-skeptical explanation invokes a connection between my current memory and past experience, and induction does the rest, raising the probability of my having a certain experience now, given the nature of that past experience; there is no need to invoke an *independent* connection between past experiences and current experiences to account for the match. Contrast this with the skeptical alternative: what sort of skeptical hypothesis would make it non-accidental that the experience and memory match? Any such hypothesis would have to posit some thing (or things) in the past—some state, agent, mechanism, or what have you—and a connection between it and the present experience on the one hand, and between it and the memory on the other. The memory's being caused by the past experience is thus a simpler explanation of the data than the presence of some common cause—another state, mechanism, deceiving agent, or what have you.¹⁶

But how can I generalize from such a case so as to conclude not only that this memory is correct, but that my memories more generally are likely to be correct? Well, if this sort of match is not a fluke or coincidence, if there is some causal explanation for it, some real process or connection between the content of my memory and the content or character of my present experience on a particular occasion, then memories with alternative contents are likely to match as well. This is a crucial point, and so worth motivating further. Consider the analogy: I think of a number between zero and a billion, and Raymond tells me, *correctly*, that I am thinking of 801,044,193! That would give me some justification for thinking that he has some way to read my mind or, perhaps, influence my mind, such that the match would have existed even if I had thought of a different number. Similarly, the fact that, of all the endless kinds of experiences I could have, I am having the sort of experience I should have if my memory were correct gives me a reason to think that it would have been correct even if the memory had a different content.

If what I have said in this section is right, then I can have at least some justification for taking my memories to be reliable or likely to be true, where this justification does not itself depend on memory. It must be admitted that, thus far, the argument might provide only a modest degree of justification—and, as already mentioned, only memory understood (roughly) as an inclination to regard some quality,

¹⁶ What about the possibility that my present experience is actually what gives rise to a corresponding delusive memory? This skeptical hypothesis would explain why I have an apparent memory of having just had a similar experience—e.g., the apparent memory of hearing the same chirping sound that I hear in the present. But it would not explain my apparent memory of having other experiences—e.g., it doesn't explain why I have the apparent memory of hearing the first part of a song that is playing now, and the apparent memory of pulling the chair to sit here, and placing this cup of coffee on that coaster. This response might raise difficult questions about the limits of what we can attend to in the specious present. I discuss some of these questions in section 5.

structure, or content as having been experienced. I will attempt to address some related worries in the objections that follow. But the point I want to stress now is that the argument provides some justification for relying on the fundamental sort of memory needed for keeping track of my experiences and thoughts over time. With this justification in hand, I might be able to acquire some justification for trusting the more substantive conclusions of more complex arguments. These might include arguments to the effect that various apparent perceptions and memories cohere very well with each other, and that the best explanation of this impressive coherence is that many of these apparent perceptions and memories are likely to be correct. My justification for believing that my memories are trustworthy might thus be increased after I have some initial justification to trust them.

5. Replies to Some Objections

First Objection: *Even if particular instances don't rely on memory, generalizing from particular instances requires that they be remembered.*

BonJour very briefly considers a solution at least roughly like the one I have offered: A particular instance of justifying a belief about a past experience need not depend on memory, “since the content of the memory to be justified can be compared directly with a present perception” (2010: 169). He objects, however, that “any attempt to generalize from such particular justifications will require that they be remembered” (169). I don't think so, for reasons already given above. From the fact that, of all the experiences possible a priori, the actual experience is just the sort I seem to recall, or just the sort I should expect given what I seem to recall—a fact I can recognize within the specious present—I can infer not only that I happen to be recalling correctly, but that my recalling this correctly *is not a mere accident or coincidence*. But then, some degree of trustworthiness or reliability can be ascribed to my memories, and not just to this particular one or memories with the very same content. I can thus rely on a single instance as support for the general claim, and do so within the specious present, and so not have to recall it.

Second Objection: *You can't get much justification from a single case.*

Perhaps BonJour's underlying worry, raised explicitly by Moon (2017), is that one cannot get *much* justification from a *single* case. Here's a way to motivate this objection.¹⁷ Suppose that it suddenly seems to Norman that the President is in New York. He turns on the TV and finds, on the basis of live news coverage, that the President is indeed in New York. If the argument for the reliability of memory works, then a similar argument with the conclusion that Norman has a reliable clairvoyant ability works as well. But it is surely implausible that this single case could provide Norman with much, if any, justification for believing that he has a reliable clairvoyant ability.

This may be a legitimate concern; perhaps the degree of justification is not likely to be very high. But I want to offer four reasons to think the argument given here is still

¹⁷ Thanks to Nevin Climenhaga for raising this example, which is a variation on BonJour's (1980) famous case of Norman the Clairvoyant.

significant.

First, as we've seen, it's not obvious that the justification is bound to be pretty small, since this depends on how low the prior probability of a match is. If this prior probability is extremely low, as it is when the possible alternatives are extremely large or endless, a single case can yield very strong confirmation. The example already given above involving Raymond's report on what number I am thinking of is just one simple illustration of this. It's true that we normally don't take a single prediction to provide a high degree of support for a hypothesis, but this is arguably because we normally make predications in a context in which we assume some background knowledge that narrows the possibilities down considerably, so that the predicted observation is one of a small number of possibilities.¹⁸ For this reason, normally, that a single prediction turns out to be correct is not very surprising. This explains why we are inclined not to take Norman's experience as providing much support for clairvoyance: while there are many places that the President could be, given Norman's background knowledge, it's not that surprising that the President should turn out to be in New York.

Second, if I can be aware of more than one sort of match between memory and experience within the specious present, or a match of a more complex sort, as I am inclined to think, that could raise the degree of justification significantly. This could benefit from more extended discussion, but it is plausible that what we are aware of in the present moment is richer than the simple examples I have used for illustration, even if there is a firm limit to the number or complexity of things of which we can be aware in the specious present. For example, I presently have the apparent memory of having sat down to type on a laptop computer, thinking and writing about the problem of memory, and hearing the chirping of crickets and the hooting of an owl. I find, at the same time, that I am having the experience as of typing on my laptop computer, and thinking and writing about the problem of memory. I still seem to hear the chirping of crickets, though not the hooting of an owl—though that too is not surprising, for I seem to recall that owls don't hoot continuously. While it takes some time to describe all this, or express it in language, it is plausible that much or all of this content is available in the specious present.

Third, as already discussed, if the argument provides some degree of justification for relying on the fundamental sort of memory needed for the construction of more complex arguments, it is possible that these arguments further raise the degree of justification we have for trusting memories.

Fourth, even if some of these points are questionable, it is still a significant and interesting result that we have an argument that yields even some degree of justification, where most seem to have thought there could be none. To use an obvious analogy: if one is trying to lift oneself up by one's own bootstraps, the really major accomplishment is to do it at all, even by a millimeter, with the actual height achieved being very secondary.¹⁹

Third Objection: *You cannot arrive at reasons to think that the hypothesis is the best explanation without relying on memory.*

BonJour also objects that "it is doubtful that the reasons, whatever exactly they

¹⁸ In a footnote, Harrod makes a similar point (1942: 64, n. 1)

¹⁹ Thanks to Larry BonJour for suggesting this analogy.

may be, for regarding the explanatory [hypothesis] ... as providing the *best* explanation can be arrived at without reliance on memory” (2010:169), and others have since raised similar objections (e.g., Moon 2017). The need for memory seems inevitable if the reasons are many, or are of a very complex sort, as is plausible when it comes to abductive arguments for the reliability of perceptual experiences. But it’s less clear that this is the case for the abductive argument given here. First, the hypothesis of the reliability of memory is quite minimal. It does not involve specifying a mechanism or process that accounts for this reliability. It invokes a connection between something in the past—my past experiences—and my current memories. Second, as already discussed, we have a straightforward reason to think skeptical explanations are inferior: we get an explanation of the match between present experience and memory “for free,” so to speak, thanks to induction. The skeptical hypotheses don’t get this correlation or match “for free.” We need not go through a complicated set of reasons, or eliminate numerous alternatives piecemeal, and this at least makes it more plausible that we can have justification for taking the hypothesis to be the best explanation without relying on memory.²⁰

Fourth Objection: *Isn’t memory needed to (a) understand the words or at least (b) employ the relevant concepts involved in your abductive argument?*²¹

My response to the previous objection might suggest an even more fundamental objection to the effect that the use of memory is inevitable: we depend on semantic memories of some sort in order to understand what words mean, and what others are saying to us. You have to remember what words mean in English in order to understand the abductive argument presented in this paper. Doesn’t this show that the proposed justification depends on memory after all? Even if the dependence on language can be avoided, there is a similar objection that we arguably rely on memory even in the mere employment of concepts. These concepts are stored, preserved, or maintained in some

²⁰ One might object that the subject still needs to attend to or have in mind a number of different items: the explanans, the explanandum, a reason to think the one explains the other, and a reason to think it is better than skeptical explanations. That is already four items, and arguably each of these items contain some further complexity (e.g., the explanandum involves both the apparent memory and the corresponding experience). So, the subject must have at least four, arguably more, items in mind at once. Isn’t this already too much? Indeed, don’t scientific studies show that the number of things we can attend to at once is limited to about four? (E.g., Cowan 2001) Thanks to Jack Lyons for raising a similar worry.

I cannot discuss the relevant psychological literature here, and will make only the following points in response. While I want to allow for the possibility that my argument is defeated by empirical data about the limits of attention, it is difficult to find studies that clearly fit the bill in the scientific literature. First, the scientific studies concern limits to the number of “chunks” (which might be simple or complex) that we can attend to at once. It might be that we can attend to complexes, items that are grouped or chunked together, with the effect that the limit is not as restrictive as it may initially seem. If I am able to attend to an explanatory or probabilistic relation between A and B, or infer A from B, am I attending to one item, or two, or three? If this counts as one item, we might be able to access the relevant reasons while remaining within what cognitive scientist’s hypothesized capacity limits. Second, the studies tend to focus on limits to “working memory,” usually understood as a kind of short-term memory in the service of ongoing cognitive tasks, with only tentative implications for limits on attention or conscious awareness itself, let alone for how complex an argument must be, and by what measure, in order for it to “fit” in the specious present.

²¹ I’m grateful to an anonymous referee for raising this objection.

way by my cognitive system, and my ability to continue to employ these concepts depends on the existence of this cognitive system and the availability to me of its concept-retrieving and concept-deploying functions.

In response, I should begin by emphasizing two points: my task in this paper is to explain how one might be justified in believing that one's apparent memories (in the sense discussed) are reliable or likely to be correct, *not* to explain how one might be justified in believing that one's semantic memories are correct, that one is using language correctly, or that one has a particular sort of cognitive system. Second, while a belief or inference might depend causally on the fact that one is using language correctly or that one has a particular sort of cognitive system, it does not follow that the belief or inference must depend essentially, for its justification, on the fact that that one is using language correctly or that one has such a cognitive system. Indeed, the abductive argument I offered here requires no assumption, on the subject's part, regarding the reliability of any form of memory, whether linguistic or semantic, or episodic—or, for that matter, procedural.

It may help to elaborate on the response a bit, for versions (a) and (b) of the objection respectively. While you must understand the English language in order to read my explanation of the central argument, and while the use of language is typically indispensable in thinking various problems through and increasing our understanding, it doesn't follow that one could not grasp the same sort of argument without relying on a more or less correct memory of what words conventionally mean, or without assuming that one's memory of how words are used is correct. Indeed, one can coherently assume that one does not know whether one is using words correctly—i.e., using them in ways normally used, whether by others or even by one's previous self—but still understand the argument and appreciate its force. A simple example may help illustrate the point: I can be justified in believing that circles have no corners, even if I am inclined to express this by saying “snurkles hafenot snorners,” and even if I am inclined to think, incorrectly, that this is in line with the way that the thought is normally expressed, by others or my previous self.

We can give a similar response to the second version of the objection. I agree that, as a matter of fact, the mere employment of concepts (at least typically) uses memory, in the sense that it (at least typically) relies on the existence of a cognitive system that stores, preserves, or maintains these concepts or conceptual abilities. But this is not essential for justification. What matters is that I can employ the relevant concepts, and we can leave open where I got these concepts from and what maintains them. A simple example may help illustrate the point: my occurrent belief that circles have no corners does not depend for its justification on how I acquired the concepts that I am employing, or whether and how these concepts are stored or maintained over time in my mind. I can be justified in believing that circles have no corners even if God had just now put these concepts in my mind, or that in some other way I am grasping and employing them for the very first time. Similarly, even if God had just put the concepts of memory, experience, etc., in my mind, this would not defeat my justification for thinking that memory is reliable.

Fifth Objection: *Sometimes our present experiences don't match what our memories lead us to expect. Wouldn't this be a reason to distrust our memories, with strongly skeptical consequences?*²²

²² Thanks to an anonymous referee for raising this worry.

If certain experiences can, by my argument, provide a reason to trust our memories, wouldn't certain other experiences give us a reason to distrust our memories? Suppose I have a memory of experiencing a chirping sound for some time now, but there is in fact no chirping in the present. Doesn't that give me a reason to distrust my memories? Such experiences could easily undermine our reliance on memory, with radically skeptical consequences.

In response, I agree that if I am in a situation in which this sort of mismatch occurs, that could give me some reason for distrusting my memories. But two points limit the significance of this sort of case. First, it is plausible that what we are typically aware of in the present is richer than the simple examples I have used for illustration, and that some core features of our experiences that match our apparent memories are pervasive. I typically have a rich experience as of having a more or less specific sort of body, and as of being in a more or less specific sort of environment, and many of these features experienced are as expected. Even if the chirping sound ceases, my experience may yet be largely in conformity with what I would expect given my apparent memory. Second, even if there are cases where I have a disconfirming experience and no confirming experiences, or cases in which the disconfirming experiences (or experiential features) eclipse or outnumber the confirming ones, with the result that I have no good reason to trust any of my memories at that moment, it may be that such cases are rare and short-lived. I might still *typically* have access to good reasons to trust my memories.

Sixth Objection: *Wouldn't you have to be using the argument every time you formed a belief on the basis of memory?*

Michael Huemer raises a different objection. If my justification for belief in the reliability of memory were inferential, "I would have to be in some sense using the argument *every time* I had a justified memory belief. It would not be enough for me to go through the argument once, and thenceforth merely remember that I had demonstrated the reliability of memory" (1999: 347). After all, if this were enough, then why, asks Huemer, couldn't I justify my belief that the sun is 93 million miles away by merely remembering that it is 93 million miles away? If I can rely on memory in this way, I wouldn't need to already have justification for thinking that memory is reliable.

I agree that it is not enough that I merely remember that I had demonstrated the reliability of memory. But suppose that, unlike the example involving the sun (where I no longer have access to the initial evidence or source for my belief), I am able to remember the argument itself—or that I have reflective access to such an argument, whether due to memory or not. This doesn't seem problematic in the way that remembering merely that I had demonstrated the reliability of memory would be. I suggest that this is because, in the former case but not the latter, I have access to a good reason to believe that memory is reliable.

There is, however, a related objection that has to do with the distinction between propositional and doxastic justification. I discuss this objection in the final section of the paper.

6. Having a Cognitive Home

Some may object that we rarely, if ever, believe that memory is reliable on the basis of an argument of the sort given here. Being *justified in believing* that *p* ("doxastic justification")

arguably requires not only *having good reasons or justification to believe that p* (“propositional justification”); it also requires *believing that p*, and doing so *on the basis of good reasons*. Even if we have good reasons, of the sort offered here, to believe that memory is reliable, it is implausible that we ever believe this on the basis of those reasons.

First, I am not convinced that doxastic justification is ruled out. Why should we think it is? The thought might be that being justified in believing something at some time requires using an argument or inference in support of one’s belief at that time, or at least having used one in the past. But I doubt that having doxastic (non-foundational) justification requires this. I believe that it is now approximately 7 p.m. on the basis of looking at my watch, and it does not seem that I rely on an actual inference in arriving at this belief. And yet it makes sense to say that I am justified in believing that it is 7 p.m., and that part of my justification for believing this is my belief that this watch is reliable or trustworthy. Thus, at least if we want to allow that such beliefs be doxastically justified, we should not hold that this requires the *actual use*, now or in the past, of a relevant argument or inference.²³

Second, even if the reliability of memory is rarely, if ever, doxastically justified, propositional justification is still significant. I am primarily interested here in defending the view that we have good, non-circular reasons to take our memories to be trustworthy. For this purpose, it is enough that some such reason or argument is *available* to us. Whether we actually use the argument is less interesting to me than whether we could have used it.²⁴

I am sometimes asked why we should care about whether we can justify our reliance on memory from within the specious present, and similar epistemological inquiries of the traditional Cartesian sort, if I admit that Hume may very well have been right that our ordinary beliefs are typically not the product of rational argument. At least part of the answer is that we care about rationality, and rationality remains essentially tied to reasons for belief that are available from the first-person perspective, even if that perspective is often focused on other things. Rationality might have to do with what we are able to do, or where we can go, by relying on the first-person perspective, and not necessarily with where we already are; various memorial, perceptual, and other beliefs might be rational in the sense that I am able to justify them, to access and appreciate reasons for their truth. My being rational consists, so to speak, in having a cognitive home, and like any home, what makes it my home is that I can return to it, not that I am already there or that I rarely leave.²⁵

²³ In any case, there are other accounts of doxastic justification that reject the requirement that one use such an inference. Perhaps the use of such an inference should not be required as part of the basing relation. See Korcz (2015) for competing accounts of the basing relation. Or if it is required as part of the basing relation, perhaps doxastic justification, one’s being justified in believing something, doesn’t always require satisfying the basing relation.

²⁴ But do ordinary subjects have access to the sort of argument I have presented here? This is a difficult question. Non-philosophers don’t discuss their justification for relying on memory; they are not in the business of reconstructing or explaining our justification. But that does not show that they aren’t able to grasp it, in at least a rough, implicit way: ordinary subjects can recognize that they are having experiences with features that are familiar and expected, and seem quite capable of understanding that this gives them some reason to trust memory. I do not take myself to have established that ordinary subjects do indeed have justification for trusting memory, but it is far from obvious that they don’t have it.

²⁵ I borrow this useful metaphor of a “cognitive home” from Williamson (1996) and some of his critics. Williamson argues that we are “cognitively homeless” in the sense that we cannot satisfy the standard of

Acknowledgments

Early versions of this paper were presented at the Central APA (Chicago 2018), the CSPA (St. Louis 2017), Northwestern University (Evanston 2017), and the Russell XIV Conference (Healdsburg 2017). Many thanks to Larry Bonjour, Richard Fumerton, Andrew Moon, Geoff Pynn, Kevin McCain, James Beebe, Jonathan Weisberg, Lars Enden, Jim Stone, and an anonymous referee for very helpful comments on early versions of the paper. I also benefited from discussions with Baron Reed, Sandy Goldberg, Nik Maggos, Bruce Russell, Matthias Steup, Alastair Norcross, Mylan Engel, Hans Hansen, Nevin Climenhaga, Al Casullo, Jack Lyons, Nathan Weston, Trevor Nyman, Finnur Dellsén, Bryan Appley, Landon Elkind, David Stern, Casey Swank, Gregory Landini, and Jennifer Kayle. Finally, this paper is dedicated in memory of Robert C. Coburn (1930-2018). I am grateful to Robert for introducing me to the epistemology of memory in graduate school, and for many delightful and engaging discussions of this problem and much else.

REFERENCES

- Bernecker, S. 2008. *The Metaphysics of Memory*. Dordrecht: Springer
- Bonjour, L. 1980. 'Externalist Theories of Empirical Knowledge.' *Midwest Studies in Philosophy*, 5 (1):53-73.
- 1998. *In Defense of Pure Reason*. Cambridge University Press.
- 1999. 'Foundationalism and the External World.' *Philosophical Perspectives*, 13: 229-249.
- 2010. *Epistemology: Classic Problems and Contemporary Responses*, 2nd Ed. Rowman & Littlefield Publishers.
- Brandt, R. 1955. 'The Epistemological Status of Memory Beliefs.' *Philosophical Review*, 64 (1): 78-95.
- Coburn, R. 1960. 'A Defect in Harrod's Inductive Justification of Memory.' *Philosophical Studies*, 11 (6): 81-85.
- Cowan, N. 2001. 'The Magical Number 4 in Short-Term Memory: A Reconsideration of Mental Storage Capacity.' *Behavioral and Brain Sciences*, 24: 87–185.
- Frise, M. 'The Epistemology of Memory.' *The Internet Encyclopedia of Philosophy*. ISSN 2161-0002, <https://www.iep.utm.edu/home/about/>, 10/10/2018.

"luminosity," which he associates with traditional Cartesian internalism. Baron Reed (2006), among others, has argued that Williamson's standard is excessively strong, and that a standard that is weaker than this is unaffected by Williamson's argument and yet substantial enough to satisfy the Cartesian. While this paper does not concern Williamson's luminosity argument, the point I am making here is similar to Reed's claim that there is "shelter for the cognitively homeless."

- Fumerton, R. (1995). *Metaepistemology and Skepticism*. Rowman & Littlefield.
- Harrod, R. 1942. 'Memory.' *Mind*, 51 (January): 47-68.
- Hasan, A. 2017. 'Skepticism and Spatial Objects.' *International Journal for the Study of Skepticism*. DOI: 10.1163/22105700-20171199
- Hasan, A. and Fumerton, R. 2017. 'Knowledge by Acquaintance vs. Description.' *The Stanford Encyclopedia of Philosophy* (Fall 2017 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2017/entries/knowledge-acquaindescrip/>>.
- Huemer, M. 1999. 'The Problem of Memory Knowledge.' *Pacific Philosophical Quarterly*, 80 (4):346–357.
- Korcz, K. 2015. 'The Epistemic Basing Relation.' *The Stanford Encyclopedia of Philosophy* (Fall 2015 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2015/entries/basing-epistemic/>>.
- McCain, K. 2014. *Evidentialism and Epistemic Justification*. Routledge.
- McGrew, T. 1995. *The Foundations of Knowledge*. Rowman and Littlefield Publishers.
- Moon, A. 2017. 'Skepticism and Memory.' In S. Bernecker & K. Michaelian (eds.), *The Routledge Handbook of Philosophy of Memory*. Routledge.
- Reed, B. 2006. 'Shelter for the Cognitively Homeless.' *Synthese*, 148 (2): 303-8.
- Vogel, J. 2008. 'Internalist Responses to Skepticism.' In John Greco (ed.), *The Oxford Handbook of Skepticism*. Oxford University Press.
- Williamson, T. 1996. 'Cognitive Homelessness.' *Journal of Philosophy*, 93 (11): 554-73.