

4

Perceptual Capacities, Knowledge, and Gettier Cases

Susanna Schellenberg

Abstract

I will exploit the basic commitments of capacitism to develop a distinctive externalist view of perceptual knowledge. The basic idea of capacitism is that perception is constitutively a matter of employing perceptual capacities that function to discriminate and single out particulars in our environment. It is because a given subject is employing a perceptual capacity with a certain function that her mental states have epistemic force. Employing such perceptual capacities constitutes a mental state that provides us with phenomenal evidence, and employing such capacities in the good case also provides us with knowledge-worthy factive evidence. In perceptual Gettier cases the subject only has phenomenal evidence and so lacks sufficient evidence for knowledge. This approach is distinctive in three respects: it groups perceptual Gettier cases in with hallucinations, it treats both via a lack of sufficient evidence (rather than through invoking some sort of fourth condition), and it divides up perceptual Gettier cases from fake barn cases and broken clock cases. The underlying picture of perceptual knowledge avoids the pitfalls of both externalist disjunctivist views and internalist views, while revealing what is right in both externalist and internalist approaches. Capacitism is an externalist view that does not invoke reliability, remains steadfastly naturalistic, and by recognizing a metaphysically substantive common element between perception and hallucination avoids any commitment to disjunctivism.

I offer a view of perceptual knowledge that is externalist but not disjunctivist, and show how it helps handle perceptual Gettier cases. The basic idea is that in experience we employ perceptual capacities that function to discriminate and single out particulars in our environment. It is because a given subject is employing a perceptual capacity with a certain nature that her mental states have epistemic force. Employing such perceptual capacities yields a mental state that provides us with a weak type of evidence that I call *phenomenal evidence*, and employing such capacities in the good case

provides us, moreover, with a second strong type of evidence that I call *factive evidence*. As I will argue, while factive evidence is sufficient evidence for knowledge, phenomenal evidence is not.

In perceptual Gettier cases, it is standardly thought that the subject has sufficient evidence for knowledge, but fails to know for some other reason. Once we recognize the distinction between phenomenal evidence and factive evidence, we can say that in perceptual Gettier cases, the subject has mere phenomenal evidence; but since she does not have factive evidence, she fails to have sufficient evidence for knowledge. Thus, I offer a way of handling perceptual Gettier cases without appeal to any factor beyond that of evidence. On my view, perceptual Gettier cases receive the same underlying treatment as hallucinations.

To keep the discussion tractable, I will focus on perceptual knowledge and perceptual Gettier cases. The lessons I wish to draw, however, reach beyond perception. Towards the end of the chapter, I will consider the extent to which my account generalizes to non-perceptual cases. To keep the discussion tractable, I will moreover focus on developing my own positive view, rather than criticizing others.

1 Perceptual Experience and Perceptual Evidence

In order to give an analysis of perceptual Gettier cases, it is useful to first clarify some basic features of perceptual experience.

1.1 *The nature and epistemic force of perceptual experience*

Let's first look at what happens when we are in the good case and so are perceptually related to what we purport to single out in our environment. When we perceive we employ perceptual capacities by means of which we discriminate the objects, events, and property-instances in our environment. Say we see that the book on the red desk is green. We employ our capacity to discriminate the green color of the book from the surrounding red color of the desk and, in doing so, we are able to single out the green book from its surrounding.

More generally the idea is that perception is fundamentally a matter of employing perceptual capacities whose natural function is to discriminate and single out particulars in our environment. The relevant particulars are mind-independent objects, events, and property-instances, such as cups, cups falling off tables, and the color or shape of a cup. When we perceive what we purport to single out, we in fact single out the particular in our environment.

Sensory states are understood as ensuing from employing such capacities in a sensory mode, that is, modes such as seeing, hearing, touching, smelling, or tasting. So if the same perceptual capacities are employed in the same sensory mode, the same sensory state is constituted. In this way, the employment of perceptual capacities in a sensory mode grounds the sensory state.

The perceptual capacities of concern here are individuated by the types of particulars they function to single out. For present purposes, we can understand types of particulars as natural kinds in the environment. Such natural kinds include shapes and more controversially colors, as well as kinds of objects and events.¹ So perceptual capacities are individuated by types of mind-independent particulars, while sensory states are in turn individuated by the perceptual capacities employed. In this sense, the account provided is an externalist account of sensory states. As I will argue shortly, it is modestly externalist insofar as the perceptual capacities can be employed while failing to single out any particular of the type they function to single out.

This account of the fundamental structure of perceptual experience yields a natural explanation for why perceptual experience has epistemic force. The key idea is that perceptual experiences have epistemic force in virtue of the nature of the perceptual capacities employed in experience. Perceptual capacities function to single out particulars in the environment. In virtue of functioning to single out particulars in our environment, perceptual states are systematically linked to what they are of in the good case, that is, the case of a successful perception. And in virtue of being systematically linked to what they are of in the good case, such perceptual states provide evidence for what they are of in the good case.

Consider a perceiver who sees a white cup. Let's call him Percy. Percy is in the good case. His perceptual state is systematically linked to the white cup he is seeing in that he is perceptually related to that very cup. His perceptual state is of the white cup and it provides him evidence that the particular cup is present, because of the systematic linkage between his perceptual state and the state of his environment. I refer to this view as *capacitism*, since the fundamental structure of perception is analyzed in terms of employing perceptual capacities.

When we suffer a hallucination, we employ the very same perceptual capacities that we would be employing were we enjoying a subjectively indistinguishable perception—albeit failing to single out the relevant particulars. The perceptual capacities are, even when employed in the bad case, systematically linked to what they are of in the good case. After all, they still function to do what they do in the good case, namely discriminate and single out particulars in the environment. The problem is simply that the environment is not playing along. As I will argue for in more detail shortly, when we hallucinate, we have some evidence for believing that things are how they seem to us in virtue of employing capacities that are systematically linked to what they are of in the good case. But this is a merely weak sort of evidence that is not sufficient evidence for knowledge.

Now one could argue that hopes, fears, and imaginations are, like perceptions, a matter of employing mental capacities. If that is right, why does this account of perceptual experience not overgeneralize to these other mental states? The key idea is that perceptual experiences have epistemic force in virtue of the fact that the perceptual capacities

¹ For a detailed development of such a color realist view, see Byrne & Hilbert (2003).

employed have the function to single out particulars in the environment. It is because the capacities constituting perceptual experience function differently than the capacities in play in hoping, fearing, and imagining that perceptual experiences have epistemic force and hopes, fears, and imaginations do not. Perception is our primordial connection to particulars in our environment. The fundamental difference between perception, on the one hand, and hopes, fears, imaginations, and also beliefs, on the other, is that perceptual capacities function to single out particulars to which we are perceptually related, while the capacities employed in those other mental states do not have this function.

1.2 Phenomenal and factive evidence

This account of the fundamental structure of perceptual experience and its epistemic force suggests a distinction between two ways of individuating perceptual states, two ways of individuating perceptual content, and two forms of perceptual evidence.

To see how this account of the fundamental structure of perceptual experience provides two ways to individuate perceptual states, compare Percy, who perceives a white cup, with Hallie, who hallucinates a white cup. While perceiving Percy is perceptually related to a white cup, hallucinating Hallie is not. It seems to Hallie that she is perceiving a white cup, but she is not. Let's assume Percy and Hallie both form a belief that they express with 'that cup is white.' We can leave open for now whether Percy's belief has the same content as Hallie's belief. The important point for present purposes is that Percy and Hallie would express their beliefs in the same way. After all, their environment seems the same way to them.

We can analyze the perceptual capacities Hallie employs with respect to the conditions for which they function (good cases), despite the fact that Hallie is not currently in such a condition. There is nothing wrong with Hallie's perceptual system. The problem is that her environment is not playing along. Since her perceptual capacities are working well with respect to the conditions for which they function, Hallie's perceptual state has at least some epistemic merit: it is a product of employing her perceptual capacities. The way in which Hallie is failing is simply that she is not singling out relevant particulars. So both Percy and Hallie employ perceptual capacities, and both of their perceptual capacities are systematically linked to what they are of in the good case in that they function to single out particulars. But only Percy is actually singling out relevant particulars.

The distinction between the capacities employed and what, if anything, they single out brings into focus two ways to individuate perceptual states. On one way of individuation, the relevant features that characterize Percy's and Hallie's mental state are the same: the same perceptual capacities are employed. On the other way of individuation, the relevant features that characterize their mental states are different: Percy successfully singles out environmental particulars; Hallie fails to single out what she purports to single out. Capacitism unifies these two ways of individuating

the relevant mental states in that both are analyzed in terms of the perceptual capacities employed.

Associated with these two ways of individuating perceptual states, there are two ways of individuating perceptual content. Assuming standardly, though not uncontroversially, that perceptual states have representational contents, we can say that employing such perceptual capacities yields a content type.² The content type is constituted by the perceptual capacities employed and covaries one to one with the sensory state. Percy's and Hallie's mental states are the same regarding their content type. The token contents of their mental states, however, differ: the token content of perception is a singular content that is constituted by successfully employing these capacities in an environment, thereby singling out particulars in said environment. The token content of hallucination is defective insofar as the capacities are employed baselessly. Hallie fails to single out what she purports to single out. One way to analyze what it means for the content to be defective is to say that it is gappy. While the token singular content covaries with the environment in which the relevant capacities are employed, the content type does not covary with the environment in which they are employed.

Finally, and most crucially for what will come, these two ways of individuating perceptual states and perceptual content yield two kinds of epistemic evidence: phenomenal evidence and factive evidence. Phenomenal evidence is individuated by the content type that is in turn individuated by the perceptual capacities employed. Factive evidence is individuated by the token content that ensues from employing these capacities successfully in a particular environment. In the good case, perceptual experience provides us with both phenomenal and factive evidence. In the bad case, perceptual experience provides us only with phenomenal evidence. There is no factive evidence in the bad case because the capacities were not employed successfully and the ensuing token content is defective.

By introducing the notion of phenomenal evidence, we can explain what evidence Percy and Hallie have in common. Moreover, we can explain in virtue of what hallucinating Hallie is not simply blameworthy for her belief that there is a white cup on her desk. She has reason for believing that there is a white cup on her desk. After all, for all she can tell, there is a white cup on her desk. She has phenomenal evidence that supports her belief.

By introducing the notion of factive evidence, we can explain in virtue of what Percy is in a better evidential position than Hallie. While his evidential position may seem to him to be indistinguishable from that of Hallie's, Percy has, unbeknownst to him, additional factive evidence and thus evidence that the particular white cup to which he is perceptually related is in fact before him. Thus, Percy has evidence that supports a singular thought about his environment. As I will argue below, only Percy has sufficient evidence for knowledge.

² For a defense of this idea, see Schellenberg (2011a) and (2014a).

By unifying the two ways of individuating perceptual states via the perceptual capacities employed, capacitism can account for externalist and internalist intuitions about perceptual experience: there is a metaphysically substantial common element between the good case and the bad case, thus avoiding any disjunctivist implications, while acknowledging that the content of perception is singular.

In what follows I will assume this view of perceptual experience, and the associated two-level view of perceptual evidence.³ In the rest of this chapter, I will discuss the implications of this view for perceptual Gettier cases.

2 Perceptual Knowledge and Perceptual Gettier Cases

In Gettier cases it is standardly thought that the subject has a true belief and sufficient evidence for knowledge, but still lacks knowledge. On my view this analysis is too simplistic. I will argue that the subject in a perceptual Gettier case lacks sufficient evidence for knowledge: she has phenomenal evidence, but fails to have factive evidence. This allows me to explain why the subject in a perceptual Gettier case lacks knowledge, without requiring anything beyond a sufficient evidence requirement. But first let's take a closer look at the relation between phenomenal evidence, factive evidence, and perceptual knowledge.

2.1 *Perceptual evidence, justification, and perceptual knowledge*

Given the distinction between phenomenal and factive evidence, and given the orthodox assumption that sufficient evidence is required for knowledge, I am now in a position to ask: which of either form of evidence is sufficient evidence for knowledge? As I will argue, factive evidence is sufficient evidence for knowledge, but phenomenal evidence is not.

The capacity view explains in virtue of what one is in a better epistemic position when one perceives than when one hallucinates. Consider again Percy and Hallie, who both form the belief 'that cup is white.' Insofar as Percy's belief is supported by more and better evidence than Hallie's, Percy's belief is better justified than Hallie's. More generally, when one forms a belief on the basis of perceiving the particular α , one's belief is better justified than when one forms a belief on the basis of suffering a hallucination that is subjectively indistinguishable from perceiving α . One's evidence in the good case justifies any supported belief to a higher degree than does one's evidence in the bad case.⁴

Now if one cannot tell that there is a difference in evidence, let alone what the difference is, what effect will the difference in evidence have for one's cognitive life? In response, we can say that the difference in evidence will have repercussions for what

³ For a detailed defense of this view, see Schellenberg (2013, 2014b, 2016a, 2016b). For a critical discussion, see Byrne (2014), McGrath (2016), Pautz (2016), and Neta (2016).

⁴ An interesting question is what the connection is between the strength of the evidence we have for a proposition and our confidence in that proposition. For a discussion of the relationship between having evidence for p and having confidence in p , see Neta (2003, 2008) and Silins (2005).

one is justified in believing. Factive perceptual evidence is evidence of particulars in a perceiver's environment and so justifies singular thoughts about her environment. Phenomenal evidence is not evidence of particulars in a perceiver's environment. More generally, we can say that any perception can give rise to a number of different beliefs, including singular beliefs and general beliefs. The factive evidence provided by my perception will give support to both singular beliefs and general beliefs. Phenomenal evidence, by contrast, supports only general beliefs. The point generalizes beyond experience: if you have propositional justification for the singular belief, 'that cup is white,' you will also have propositional justification for the general belief it entails, for example, the belief, 'There is a white cup.' That follows from a relatively simple schema for deductive closure.

In short, while factive evidence supports singular beliefs as well as general beliefs, phenomenal evidence supports only general beliefs. So while one cannot tell whether one is in the good case or the bad case and so cannot tell whether one has factive evidence in addition to phenomenal evidence, it makes an epistemic difference to have factive evidence in addition to phenomenal evidence, in that more of one's beliefs will be justified.

Moving from evidence to knowledge: phenomenal evidence is not sufficient evidence for knowledge since having mere phenomenal evidence is compatible with suffering a hallucination. In hallucination, the subject has a justified mental state (such as a belief) that falls short of knowledge. Factive evidence, on the other hand, is sufficient for knowledge. After all, the token content of perception that determines factive evidence is determined by and covariant with the perceiver's environment and so guaranteed to be true. Hence factive evidence is infallible. So if one thinks of sufficient evidence for knowledge in terms of safety guarantees or in terms of low fallibility, then one should agree that factive evidence is sufficient evidence for knowledge.

If this is right, then we gain perceptual knowledge by successfully employing perceptual capacities, that is, capacities that function to discriminate and single out particulars in our environment. Insofar as capacities are mental tools by means of which we relate to our environment, the states yielded by employing such capacities are mental states. Moreover, insofar as successfully employing capacities yields perceptual knowledge, capacitism entails that knowledge is a mental state. In this way, capacitism gives an explanation of what it means for knowledge to be a mental state. Knowledge is a mental state in virtue of being constituted by employing mental capacities.

Now let's assume standardly that for S to have perceptual knowledge that p, S must have sufficient evidence that p. Given this assumption and given the analysis so far, we are now in a position to formulate the following necessary and jointly sufficient conditions for S to have perceptual knowledge that p:

S has perceptual knowledge that p if and only if p is true, S employed a capacity to single out what she purports to single out, and S's mental state has the content it has in virtue of S having successfully employed her capacity to single out what she purports to single out.

This analysis of perceptual knowledge is neutral on the belief condition on knowledge. Orthodoxy has it that one cannot know that *p* without believing that *p*. capacitism is neutral on whether there is any such belief condition on knowledge. This is attractive, since arguably, we know that *p* simply in virtue of seeing that *p*. By contrast, we do not believe that *p* simply in virtue of seeing that *p*. After all, I can see that *p* without forming any beliefs. Moreover, a perceptual state that falls short of knowledge does not amount to a belief. Any perceptual belief will of course be grounded in a perceptual experience. But being grounded in a perceptual experience is not the same as the perceptual state constituting a perceptual belief.⁵ For these reasons, I will talk about justified mental states rather than justified beliefs. These justified mental states may be beliefs, but they need not be.

Justified mental states that fall short of knowledge can be true or false. Such mental states share certain features with knowledge states: they are states in which perceptual capacities are employed and states that have at least some justification. Perception yields knowledge states. Hallucination and perceptual Gettier cases yield justified mental states that fall short of knowledge. The relevant difference between justified mental states that fall short of knowledge and knowledge states is that the latter but not the former are ones in which capacities are employed successfully. But beyond the capacities being employed such that they do what they function to do, there is nothing additional in play in the case of knowledge that is not in play in the case of a mere justified mental state.

2.2 *Perceptual Gettier cases, first pass: phenomenal evidence without factive evidence*

So far I have introduced a distinction between factive evidence present only in the good case, and merely phenomenal evidence present as a common element in both the good and the bad cases, and I have sketched a conception of knowledge on which factive evidence constitutes sufficient evidence for knowledge while phenomenal evidence does not. I am now in a position to discuss the consequences of capacitism for perceptual Gettier cases. I will argue that the subject in a perceptual Gettier case lacks sufficient evidence for knowledge. This allows me to explain why the subject in a perceptual Gettier case lacks knowledge, without appealing to anything beyond a sufficient evidence requirement. Consider the following case:

Robot Dog. James is relaxing on a bench in a park and sees what he takes to be a dog in a nearby field. On the basis of his perception, he believes, ‘There is a dog in the field.’ As it happens, the apparent dog is actually a robot dog. The robot is so well

⁵ The arguments provided here need to be modified only slightly such that a belief condition on knowledge is met. If we add a belief condition on knowledge, the necessary and jointly sufficient conditions for *S* to have perceptual knowledge that *p* would be:

S has perceptual knowledge that *p* if and only if *p* is true, *S* employed a capacity to single out what she purports to single out, *S* believes that *p*, and her belief has the content it has in virtue of *S* having successfully employed her capacity to single out what she purports to single out.

constructed that it could not be distinguished from an actual dog— at least not from where James sits. James does not know that such robot dogs exist. So if that was a full description of the case, James's belief 'There is a dog in the field' would be false. However, a few feet away from the robot dog, there happens to be a real dog in the same field, concealed from James's view. So James's belief 'There is a dog in the field' happens to be true.⁶

What should we make of this case? The standard approach is to say that James's belief 'There is a dog in the field' is both true and supported by sufficient evidence for knowledge. The motivation for this approach is that there is in fact a dog in the field and the belief was formed based on ordinary perceptual processes albeit on grounds of perceiving a robot dog rather than the actual dog on the field. Why think that James has sufficient evidence for knowledge? The reasoning seems to stem from the internalist idea that James is doing just the same thing from his own perspective as a successful perceiver would be doing in a simple case of seeing a dog in the field (with no robots or other complications).

The capacivist's verdict on the case is different. James's belief has some justification. After all, James employs his capacity to single out a dog and in virtue of this, he is in a mental state that is intentionally directed at a dog. Being in this mental state provides him with some evidence for his belief that there is a dog in the field. It provides him with phenomenal evidence. So James has a justified true belief. However, since James is not in fact perceptually related to a dog and so does not single out a dog (despite it seeming to him that he is doing just that), he does not have factive evidence of a dog. The fact that there happens to be a dog a few feet away out of James's view does not mean that James has factive evidence of a dog in the field, since he did not single out that dog. He only singled out the robot. So James has phenomenal evidence but fails to have factive evidence that there is a dog in the field. Since phenomenal evidence is not sufficient evidence for knowledge, James does not have sufficient evidence to know that there is a dog in the field.

What happens to the widespread intuition that James has a justified true belief without knowledge? We can say that in perceptual Gettier cases, the belief has some justification (due to the gettiered subject having phenomenal evidence), but that it does not have knowledge-level justification. After all, the subject does not believe truly in virtue of her mental state having been arrived at by employing capacities that function to do what they are supposed to do. For example, James does not believe truly in virtue of *successfully* employing his capacity to single out a dog.

So I am arguing that perceptual knowledge is best understood in terms of employing perceptual capacities that function to discriminate and single out particulars in the environment, such that when these capacities are employed successfully the subject enjoys a mental state that is both true and justified. The key is that knowledge differs

⁶ See Ichikawa & Steup (2012: sec. 4). This case is structurally the same as Chisholm's (1966: 23, fn. 22) sheep-shaped rock case.

from mere justified true mental states in that the capacities employed in knowledge in fact succeed in serving their natural function, whereas in mere justified true mental states, the capacities are employed without singling out what the subject purports to single out.

2.3 *The metaphysical and explanatory primacy of the good case*

I have argued that we gain knowledge of our surroundings when we successfully employ capacities that function to single out particulars in our environment. These capacities are understood naturalistically, in terms of their natural function. So I am grounding the epistemic force of experience in the metaphysical properties of experience. In order to give support to this line of reasoning, it will be necessary to take a closer look at the notion of perceptual capacities in play. I will then be in a position to take a closer look at Gettier cases.

Insofar as capacities to discriminate and single out particulars in our environment yield knowledge of those particulars, one might argue that these capacities should simply be analyzed as capacities to gain knowledge of those particulars. But that would be to put the cart before the horse. It is unclear what the explanatory gain would be of analyzing knowledge in terms of capacities to know. Indeed, an account that would analyze perceptual knowledge in terms of capacities to know would be circular. According to the view developed here, the perceptual capacities in play are not analyzed as capacities to know: one neither employs the capacity to know when one is in the bad case, nor when one is in the good case.

How then should we understand perceptual capacities? A perceptual capacity is a mental capacity that functions to discriminate, single out, and in some cases classify a particular kind, such as instances of red (Julesz 1981 and Krummenacher, Grubert, & Müller 2010). Such perceptual capacities are determined by general, functional relations between the organism and its environment—for instance, global patterns of the organism's response to its environment. If we possess the perceptual capacity that functions to discriminate and single out instances of red, we are in a position to discriminate instances of red from other colors in our environment and to single out instances of red. More generally, to possess a perceptual capacity is to be in a position to discriminate and single out the type of particulars that the capacity is directed towards, were one related to such a particular.

For present purposes we can remain neutral on whether perceptual capacities are low-level (yet personal level) discriminatory capacities or more high-level concepts. That said, insofar as perception is a low-level ability that beings have which are significantly less cognitively sophisticated than we are, there are reasons to think that the perceptual capacities that constitute perceptual states are low-level discriminatory capacities rather than more high-level concepts. However, nothing here hinges on the issue.

Now insofar as one can employ capacities to single out particulars or employ them while failing to single out the particular one purports to single out, they are fallible.

Indeed, they yield states that are either (guaranteed to be true) or defective (guaranteed to be false). In this respect, the view here differs fundamentally from Williamson's (2000) view. According to Williamson, the methods employed in gaining evidence are infallible and the mental state yielded is always factive.⁷

But while employing capacities is fallible, this does not mean that the good case and the bad case are on a par. Capacitism is not neutral between the employment of capacities in the good case and the employment of capacities in the bad case. As I will argue, any employment of capacities in the bad case is derivative from their employment in the good case. The employment of a capacity in the bad case fails to fulfill the function of the capacity and therefore fails to yield factive evidence and knowledge.⁸ Thus we have mere phenomenal evidence. By contrast, the employment of the same capacity in the good case is constitutively a success in that it yields factive evidence and so allows the subject to gain knowledge.

Given that perceptual capacities can be employed while failing to single out the particulars they function to single out, one might ask why their function is to successfully single out particulars, as opposed to having the function of failing to single out particulars. What explains the asymmetry? In response, any plausible account of natural function will support the idea that the heart has the function to pump blood rather than the function to fail to pump blood. This is so even though hearts may fail to pump blood. Likewise, perceptual capacities have the function to single out particulars in the environment rather than to fail to do so. This is even though they may fail to single out a particular. An evolutionary account of function would posit that perceptual capacities evolved for the purpose of singling out particulars rather than for the purpose of failing to single out particulars: they were selected to single out particulars. However, there is no need to explain the asymmetry in evolutionary terms or in any other reliabilist terms. On any plausible account of natural function, we can say that perceptual capacities function to single out particulars rather than fail to do so. In this sense, the idea of a natural function is not tied into the idea of these functions being reliable or the conditions explanatory of a system constituted by these functions being reliable.

More specifically, we can say that the perceptual capacities employed in the bad case are explanatorily and metaphysically parasitic on their employment in the good case. There is an explanatory primacy of the good over the bad case since one can give an analysis of the perceptual capacities employed in the bad case only by appealing to their role in the good case. This explanatory primacy is licensed by a metaphysical primacy of the good over the bad case: there is such a metaphysical primacy insofar as perceptual capacities are determined by relations between perceivers and their environment. On one way of understanding metaphysical primacy, we can associate things

⁷ In a similar vein as Williamson, Millar (2008) argues that abilities are infallible and yield mental states that are factive. I am here assuming that Williamson's methods and Millar's abilities can be treated as analogous to my notion of capacity.

⁸ In this respect among others, the view developed here differs from Sosa (2007), Bergmann (2006), Burge (2010), and Greco (2010).

with natures and see if the nature of one thing makes reference to another. If so, the latter will be said to be relatively primary and the former secondary. We can then construct chains so that if the nature of *A* makes reference to *B*, and the nature of *B* makes reference to *C*, then *C* will be primary, *B* secondary, and *A* tertiary. According to capacitism, the bad case is by nature a case brought about by the subject employing her perceptual capacities, and these capacities are by nature defined in terms of success in the good case. So on this pattern, we get the good case coming out relatively primary and the bad case coming out secondary. This is not the only way to use talk of metaphysical primacy but it is one plausible way of using the term.

Phenomenal evidence and factive evidence are epistemically united insofar as both are provided by mental states that are constituted by employing perceptual capacities. In showing that both kinds of evidence have their rational source in employing perceptual capacities, the suggested view provides a unified account of perceptual evidence—one which is distinctively capacity-based, non-normative, and non-reliabilist.⁹

Factive evidence provides additional evidence that is different from phenomenal evidence. It is evidence of a different kind insofar as the systematic linkage to the environment is ideal, such that the evidence provided is factive, infallible, and perfectly safe, and hence sufficient evidence for knowledge. So factive evidence provides a rationality boost beyond the one that a perceiver already has from phenomenal evidence. This explains why Percy is in a better evidential position than Hallie. Now from the first-person perspective, one may not be able to tell the difference between a hallucination in which one has only phenomenal evidence, and a perception in which one has both phenomenal and factive evidence. But we need not think that what is accessible from the first-person perspective dictates what is rational to heed.

While this view of perceptual evidence is externalist, it makes room for a phenomenal conception of evidence. So, in contrast to externalist views such as Williamson's, capacitism shows that we have at least some evidence provided directly through experience in the bad case: we have phenomenal evidence.¹⁰ In contrast to evidential internalist views (cf. Pollock 1974, Feldman & Conee 1985, and Pryor 2000), capacitism shows that we have more evidence in the good than the bad case: we have additional factive evidence. So the defended view provides us with something that neither factive evidentialists nor evidential internalists can supply.

The fact that there is such an explanatory and metaphysical priority of the good case does not imply that there is no common element between the good and the bad case. Both the good and the bad case are brought about by employing perceptual capacities. We get at the truth in a particular way, by employing perceptual capacities. And even when we fail to get at the truth (and so are in the bad case), we are employing perceptual capacities by means of which we aim to get at the truth.

⁹ I will offer more detailed comparisons with disjunctivist, normative, and reliabilist views below.

¹⁰ According to Williamson, we have only evidence provided by an appearance proposition in the bad case. For discussion of the problems with this, see below.

So there is an epistemically relevant and metaphysically substantial common factor between the good and the bad case, namely the capacities employed. As a consequence, capacitism rejects the basic commitments of epistemic disjunctivism. Indeed, in virtue of the capacities employed, the subject has phenomenal evidence—regardless of whether she is in the good or the bad case. However, as argued above, while there is such a metaphysically substantial common factor between the good and the bad case, there is an explanatory and metaphysical priority of the good over the bad case. In virtue of this explanatory and metaphysical priority of the good over the bad case, there is an asymmetry between the good and the bad case.

So capacitism holds that a mental state can be justified by a perception or a perceptual experience that falls short of perception, as is the case in a hallucination, an illusion, or a misperception. In any case in which perceptual capacities are employed, a sensory state is yielded that at the very least provides phenomenal evidence.¹¹ Insofar as both hallucinations and at least some perceptions provide the experiencing subject with evidence, and so justification for any belief she might form, capacitism entails that justification is common to both cases of knowledge and mere beliefs. More generally we can say that insofar as both hallucinations and at least some perceptions provide the experiencing subject with evidence for her mental state, capacitism entails that justification is common to both cases of knowledge and mental states that fall short of knowledge. Since perceptual capacities function to single out particulars, their employment yields states that are prone to yield factive evidence and knowledge, even though the environment does not always play along. In this way, capacitism provides an explanation of perceptual justification and the way justification is on the one hand necessary for knowledge, but why mere justified mental states are nevertheless metaphysically and epistemically dependent on mental states that amount to knowledge.

So with the knowledge-first view, capacitism explains the bad case in terms of the good case. But against the knowledge-first view (and any other disjunctivist view), capacitism has it that there is a metaphysically substantial common element between the good and the bad case, namely the capacities employed. This common element explains how it is that we have at least some justification in the bad case. In virtue of this, the view provides a unified account of the internalist and externalist elements of perceptual knowledge and evidence.

¹¹ It is worth highlighting that I am not arguing that all capacities employed in perception have repercussions for our sensory states and the epistemic force of experience. The idea is not that whenever we use, say, a discriminatory, selective capacity, we are in a sensory state. The visual system makes use of many capacities on a subpersonal level. The thesis in play is rather that the nature of sensory states is best understood in terms of employing perceptual capacities, rather than in terms of awareness relations to strange particulars, such as sense data or qualia, or abstract entities, such as properties or propositions. For the contrast between analyzing sensory character in terms of awareness relations to peculiar entities, such as strange particulars or abstract entities, on the one hand, and understanding sensory character in terms of a mental activity, such as employing perceptual capacities, on the other, see my (2011b). We can accept this thesis while acknowledging that there are many capacities the employment of which have no repercussions for our phenomenal lives.

2.4 *Gettier cases again: separating the barns from the sheep (and the clocks)*

Now let's get back to perceptual Gettier cases. As I have argued, according to capacitism, we gain knowledge of our environment if we single out particulars in our environment by employing perceptual capacities that function to single out those very particulars. So when we employ the capacity to discriminate and single out a particular at location L in our environment and we in fact discriminate and single out that particular at location L, we gain factive evidence and thus perceptual knowledge of the particular at location L.

In a Gettier case, the particular at location L is not in fact the particular that the subject purports to single out. So in this case, the subject employs capacities that fail to discriminate and single out the particular at location L. That is the case even if it seems to the subject that she is successfully discriminating and singling out that particular at location L. Moreover that is the case even if there is a particular at location L* that the subject could have successfully singled out but did not in fact successfully single out, as in the robot dog case where there is a real dog a few feet away from the robot, and in Chisholm's original case of the sheep-shaped rock with a real sheep hidden behind it.

When a true mental state is justified, this is due to employing capacities that function to single out what seems to be present. In perceptual Gettier cases, the subject fails to have knowledge despite having a justified true mental state since she is not appropriately related to what it seems to her she is related to. More specifically, she fails to have knowledge since the capacities she employs do not single out what they function to single out and it seems to her she is singling out. So she has phenomenal evidence for belief but fails to have factive evidence.

In this sense, perceptual Gettier cases are similar to hallucinations in that capacities are employed without fulfilling their function. In contrast to hallucinations, however, there is an external, mind-independent ground for employing the capacity elsewhere in the environment. But the relevant perceptual capacity is not directed towards the right particular.

There is at least one case that has traditionally been understood to be a Gettier case that capacitism would not count as a Gettier case, namely, Ginet's and Goldman's barn façade county case (Goldman 1976: 772–3).¹² But this is a good outcome. Here is why.

Barn Façade County. Consider Henry, who is driving down the road in barn façade county. The county is peppered with barn façades: from the road they look just like barns, but they are in fact structures that only look like barns seen from the road. Viewed from any other angle, one would immediately be able to tell that they are mere façades. Henry is looking at the one and only barn in barn façade county and forms the belief 'that is a barn.' His belief is justified and true.

¹² Ginet never put the example in writing, but Goldman credits him with the example.

The standard verdict of this case is to say that since the truth of the belief is a result of luck, Henry does not know that it is a barn. Henry's belief is a result of luck since it is false in most of the closest non-actual cases. The verdict of capacitism is different and arguably more plausible: Henry employs his capacity to single out the barn he sees, and since he is perceptually related to that very barn, he has both phenomenal and factive evidence of the barn. In virtue of having factive perceptual evidence, he has perceptual knowledge. After all, he is directly seeing a real barn. One might argue that while Henry has perceptual knowledge, he lacks a more sophisticated kind of knowledge. Taking that route would follow Sosa's analysis of the case (Sosa 2007: 96, fn. 1). Sosa argues that Henry has animal knowledge but lacks reflective knowledge. We can remain neutral here on whether Henry lacks such reflective knowledge and whether we should distinguish more primitive from a more sophisticated kind of knowledge.

Either way, even though Henry has factive evidence and so sufficient evidence for knowledge, his belief may still suffer defeaters. After all, he is in barn façade county. One might argue that the sheer existence of the barn façade constitutes a defeater, regardless of whether Henry knows about them. But even if we grant this, the defeaters are not undercutting defeaters, but rather mere rebutting defeaters. So although Henry may have such defeaters, it will not affect what evidence he has regarding the barn. It will affect only what he ends up being justified to believe, all things considered.

The fake barn case should be sharply distinguished from Russell's broken clock case (Russell 1948: 154).

Broken Clock. Sophie glances at a clock showing 3 p.m. and forms the justified true belief that it is 3 p.m. But unbeknownst to Sophie the clock is stuck on 3 p.m., and she just glanced at it in one of the two moments in the day when it coincidentally showed the right time.

In this case, Sophie has factive evidence that the clock shows 3 p.m., and (unproblematically) perceptually knows that the clock shows 3 p.m. On that basis she forms the inferred belief that it is 3 p.m., but her inference is subject to defeat (after all, the clock is broken), and is in any case no longer a perceptual matter. So here I am suggesting that we move beyond the taxonomy of 'Gettier cases,' and separate out such cases as the robot dog case (perceptual Gettier case), the fake barn case (successful perception), and the broken clock case (inferential Gettier case).

Of course once the topic of inferential Gettier cases comes into view, the worry arises that my view is specially designed for perceptual Gettier cases, insofar as it is based on an understanding of the nature of perceptual experience in terms of singling out particulars, which is not apt for inference, memory, testimony, or other bases for knowledge. So one might object that there is unified phenomenon of 'Gettier cases' across perception and these other bases for knowledge which calls for a unified solution, when my account is only apt for perceptual cases.

I have two very different responses to this worry. First, there are reasons to doubt that there is a unified phenomenon of Gettier cases. From an epistemic perspective,

perception may have special features (as may inference, and memory, and testimony). The robot dog and broken clock cases, for example, have in common that there is justified true belief without knowledge, but they differ in many other details, such as whether one even lays one's eyes on the relevant particular. Both cases need to be resolved, but it is not obvious that they need to be—or ought to be—resolved in the same way.

The second response to the worry is more ambitious. There are reasons to think that the underlying account of epistemic force I offer generalizes in unifying ways. On my account, perceptual experience has evidential force because of the systematic linkage between perceptual capacities and the good case of successful perception. The natural generalization of this view of evidential force to, say, inference would be that inferred beliefs preserve evidential force because of the systematic linkage between inferential capacities and the good case of successful inference. On this view, the broken clock case might then be analyzed as an inferential bad case in which the environment was not playing along, with the result being that Sophie only winds up with phenomenal inferential evidence that it is 3 p.m. and not with factive inferential evidence that it is 3 p.m. It would lead too far astray to work through such a generalized account of knowledge here. I will reserve that for another occasion.

3 Contrast with Related Alternative Strategies

In order to explain more specifically the strategy of capacitism, it will help to contrast it with related alternative strategies: reliabilism, virtue epistemology, and the knowledge-first view.

Let's start with reliabilism. Reliability simply plays no role in my account. The epistemic force of perceptual experience neither relies on perceptual capacities being reliable nor on the reliability of conditions explanatory of one's having the capacity. On the suggested capacity view, if a subject is in a sensory state that is determined by employing perceptual capacities that function to single out *F* particulars, then she is in a sensory state that provides evidence for the presence of *F* particulars in virtue of the metaphysical nature and function of these capacities. It is rational to heed the testimony of our senses since sensory states are systematically linked to the particulars that they are of in the good case. The notion of systematic linkage in play is understood in terms of metaphysical and explanatory primacy rather than reliability. Now, the perceptual capacities employed in perception may happen to be reliable. However, even in this case it is the metaphysical and explanatory primacy of the good over the bad case that gives experience its epistemic force.

So in speaking of it being the function of perceptual capacities to single out the relevant particulars, I do not mean to speak of their reliability but rather of how they are to be understood metaphysically. In other words, I am not speaking of their actual track record whatever that might be, but their metaphysical nature (what they are).

Part of what is at issue in whether or not one invokes reliability is what one can say about Davidson's Swampman cases (1987), in which an atom-by-atom duplicate of a human being forms spontaneously when lightning strikes a swamp. Intuitively one wants to say that Swampman's experience can provide him with evidence. I can say this, provided it is agreed that Swampman's perceptual capacities function to single out particulars just as ours do. So, I can give the intuitive response that Swampman has evidence even though he has no past interactions with anything and lacks ancestors. The reliabilist, however, is forced to deny that Swampman has evidence.¹³

Turning to virtue epistemology: virtue epistemologists hold that to know is to believe truly because you believe virtuously. More specifically, the idea is that knowledge is a true belief that is not accidentally true, but rather true due to the subject's dispositions, competence, abilities, or virtues.¹⁴ As Greco (2004: 111) puts it: "To say that someone knows is to say that his believing the truth can be credited to him. It is to say that the person got things right due to his own abilities, efforts and actions, rather than due to dumb luck, or blind chance, or something else." So on such a view, you deserve credit for what you know, since what you know is due to your intellectual virtues. Sosa (2007) develops his version of virtue epistemology within the framework of his AAA model of assessment. Performances can be assessed for accuracy (truth), adroitness (manifesting intellectual competence), and aptness (being true in virtue of being competent). Knowledge is identified with apt belief. This provides a way to say that knowledge is non-accidentally true belief, while allowing that one might know, even if one might easily have been wrong.

The capacity view shares with virtue epistemology the guiding idea of explaining knowledge in terms of a mental activity that has certain distinctive properties in the good case. Moreover, like virtue epistemology, capacitism is a kind of anti-luck theory in that it rejects the JTB+ approach of analyzing knowledge in terms of a truth-functional combination of independent epistemic properties.

There are three key differences between capacitism and virtue epistemology. One key difference is that virtue epistemologists treat knowledge as a particularly successful or valuable case of belief, while according to capacitism, we know when we successfully employ capacities that function to single out particulars in our environment. Failing to employ those very same capacities successfully does not yield a belief with sufficient evidence for knowledge. So contrary to the virtue epistemological approach, knowledge is not analyzed as a particularly successful or valuable case of belief.

A second key difference between the capacities view and virtue epistemology is over whether the relevant capacities (the successful employment of which generates

¹³ This requires me to commit to an account of function that allows Swampman's heart to have the function of pumping blood, and his perceptual capacities to have the function of singling out particulars, even though Swampman himself has no evolutionary history. I take it to be a constraint on an adequate account of natural function that it gets this case right.

¹⁴ For a reliabilist virtue epistemology, see Sosa (1980, 2007, 2010) and Greco (2009, 2012). For a responsibilist virtue epistemology, see Zagzebski (1996).

knowledge) are themselves to be understood in normative or naturalistic terms. For the virtue epistemologist, these capacities are understood normatively, as intellectual virtues. For the capacity theorist, these capacities are understood naturalistically, in terms of their natural function. So I am ultimately grounding the epistemic force of experience in the metaphysical properties of experience.

A third key difference is in the analysis of knowledge. According to virtue epistemology, a subject *S* knows that *p* if and only if *p* is true, *S* knows that *p*, *S* exercised a competence to believe truly in believing that *p*, and *S* believes truly due to *S* having formed her belief in virtue of exercising the competence to believe truly. Versions of virtue epistemology differ in how they understand the ‘in virtue of’ relation. It has been understood causally (Sosa 2007: 95ff), dispositionally (Sosa 2010, 2015), and in reliabilist terms (Greco 2012). According to capacitism, by contrast, *S* knows that *p* if and only if *p* is true, *S* employed a capacity to single out what she purports to single out, and *S*’s mental state has the content it has in virtue of *S* having employed her capacity and in virtue of employing this capacity successfully singled out what she purports to single out.

How does this play out for Gettier cases? In Gettier cases, the virtue epistemologist argues that while *p* is true, *S* believes that *p*, and *S* exercised a competence to believe truly in believing that *p*, *S* does not believe truly due to *S* having formed her belief by exercising the competence to believe truly. So *S* fails to know, since her belief is not in the right way due to her competence. In short, *S* does not believe aptly. According to capacitism, the Gettier case is characterized in the following way: *p* is true, *S* employed a capacity to single out what she purports to single out, but the content of *S*’s mental state does not have the content it has in virtue of *S* having employed her capacity and in virtue of employing this capacity successfully singled out what she purports to single out.

With respect to the knowledge-first view, as developed prominently by Williamson (2000), capacitism follows the knowledge-first view in holding that knowledge is a mental state and that the methods employed in gaining evidence are infallible and the mental state yielded is factive.

One key difference is that the fundamental level of analysis of capacitism is capacities rather than knowledge. A second and related difference is that the knowledge-first view is disjunctivist, and recognizes no epistemically relevant common element between the good case and the bad case. As I argued earlier, capacitism is not disjunctivist, and recognizes an epistemically relevant and metaphysically substantial common element between the good case and the bad case, namely the capacities employed.¹⁵

A third key difference is that the knowledge-first view requires positing that we do not get evidence directly through our experience when we hallucinate, but only

¹⁵ For a view that synthesizes the advantages of virtue epistemology with the advantages of a knowledge-first view, see Miracchi (2015).

through introspection. Arguably, however, experience provides us with evidence directly—even when we hallucinate. The notion of phenomenal evidence that I have developed makes room for experience providing us with phenomenal evidence directly even in the bad case without retreating to introspective evidence. This is an important advantage of capacitism over Williamson's version of the knowledge-first view for three reasons. First, introspection is a sophisticated intellectual activity, yet even subjects who do not have sophisticated intellectual abilities can get evidence through hallucination. By relying on subjects attending to how things seem to them, the Williamsonian view overintellectualizes the way we get evidence in the bad case.

A second and more pressing overintellectualization worry is that on the Williamsonian view, the evidence we have in the bad case is an appearance proposition. Appearance propositions involve appearance concepts and some sort of self-reference. However, non-rational animals hallucinate and, presumably, they gain evidence in virtue of hallucinating even though they are not capable of being in mental states that are constituted by appearance propositions. Capacitism does not face these overintellectualization problems, since we have phenomenal evidence in the bad case in virtue of being in a sensory state: there is no need to introspect or attend to our experience to have phenomenal evidence. On the view developed, we can have phenomenal evidence even if we have no ability to refer to ourselves and do not possess appearance concepts.

Third, a view on which we get evidence only through introspection in the bad case, but directly through perceptual experience in the good case, requires positing that the source of our evidence differs at least in part in the good and the bad case. By contrast, capacitism shows that the source of both factive and phenomenal evidence is our perceptual experience. Indeed, capacitism provides for a unified account of perceptual evidence by revealing the common rational source of the evidence one has in perception and the evidence one has in a subjectively indistinguishable hallucination.

So while I am following the knowledge-first approach in arguing that we have a kind of evidence in the good case that we do not have in the bad case, contra Williamson I am not rejecting the phenomenal conception of evidence. Moreover, the notion of evidence in play is not understood as identified with knowledge. We should not and need not retreat to the idea that experience provides us only with introspective evidence in the bad case. Doing so would undermine the epistemic force of experience.

Putting this together, capacitism is an externalist view that does not invoke reliability, remains steadfastly naturalistic, and recognizes a common element in perception and hallucination. I cannot here argue in more detail that these are the right choices (though I think they are), but will rest with just identifying respects in which capacitism is distinctive.¹⁶

¹⁶ Additional alternative strategies include proper functionalism (Plantinga 1993, Bergmann 2006), arguing that a 'no defeaters' approach (Lehrer 1965, Klein 1971, 1976) solves the Gettier Problem (Lycan 2006), or identifying justified true belief with knowledge (Hetherington 2001, Weatherson 2003).

4 Conclusion

I have offered a distinctive externalist view of perceptual knowledge and shown how it helps handle perceptual Gettier cases. The basic idea of capacitism is that in experience we employ perceptual capacities that function to discriminate and single out particulars in our environment. It is because a given subject is employing a perceptual capacity with a certain nature that her mental states have epistemic force. Employing such perceptual capacities yields a mental state that provides us with phenomenal evidence, and employing such capacities in the good case also provides us with knowledge-worthy factive evidence.

In perceptual Gettier cases the subject only has phenomenal evidence and so lacks sufficient evidence for knowledge. This approach is distinctive in three respects: it groups perceptual Gettier cases in with hallucinations, it treats both via a lack of sufficient evidence (rather than through invoking some sort of fourth condition), and it divides up perceptual Gettier cases from fake barn cases and broken clock cases. The underlying picture of perceptual knowledge avoids the pitfalls of both externalist disjunctivist views and internalist views, while revealing what is right in both externalist and internalist approaches.

References

- Bergmann, M. (2006). *Justification without Awareness*. Oxford: Oxford University Press.
- Burge, T. (2010). *Origins of Objectivity*. Oxford: Oxford University Press.
- Byrne, A. (2014). Perception and Evidence. *Philosophical Studies* 170(1): 101–13.
- Byrne, A. & Hilbert, D. (2003). Color Realism and Color Science. *Behavioral and Brain Sciences* 26: 3–21.
- Chisholm, R. (1966). *Theory of Knowledge*. Englewood Cliffs, NJ: Prentice-Hall.
- Davidson, D. (1987). Knowing One's Own Mind. *Proceedings and Addresses of the American Philosophical Association* 60: 441–58.
- Feldman, R. & Conee, E. (1985). Evidentialism. *Philosophical Studies* 48: 15–34.
- Goldman, A. (1976). Discrimination and Perceptual Knowledge. *Journal of Philosophy* 73: 771–91.
- Greco, J. (2004). Knowledge as Credit for True Belief. In DePaul, M. & Zagzebski, L. (Eds), *Intellectual Virtue: Perspectives from Ethics and Epistemology* (pp. 111–34). Oxford: Oxford University Press.
- Greco, J. (2009). Knowledge and Success from Ability. *Philosophical Studies* 142: 17–26.
- Greco, J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.
- Greco, J. (2012). A (Different) Virtue Epistemology. *Philosophy and Phenomenological Research* 85: 1–26.
- Hetherington, S. (2001). *Good Knowledge, Bad Knowledge*. Oxford: Oxford University Press.
- Ichikawa, J. & Steup, M. (2012). The Analysis of Knowledge. In Zalta, E. N. (Ed.), *Stanford Encyclopedia of Philosophy*, URL: <<http://plato.stanford.edu/entries/knowledge-analysis/>>.

- Julesz, B. (1981). A Theory of Preattentive Texture Discrimination Based on First-Order Statistics of Textons. *Biological Cybernetics* 41: 131–8.
- Klein, P. (1971). A Proposed Definition of Propositional Knowledge. *Journal of Philosophy* 68: 471–82.
- Klein, P. (1976). Knowledge, Causality, and Defeasibility. *Journal of Philosophy* 73: 792–812.
- Krummenacher, J., Grubert, A., & Müller, H. J. (2010). Inter-Trial and Redundant-Signals Effects in Visual Search and Discrimination Tasks: Separable Pre-Attentive and Post-Selective Effects. *Vision Research* 50: 1382–95.
- Lehrer, K. (1965). Knowledge, Truth, and Evidence. *Analysis* 25: 168–75.
- Lycan, W. (2006). On the Gettier Problem Problem. In Hetherington, S. (Ed.), *Epistemology Futures* (pp. 148–68). Oxford: Oxford University Press.
- McGrath, M. (2016). Schellenberg on the Epistemic Force of Experience. *Philosophical Studies* 173(4): 897–905.
- Millar, A. (2008). Perceptual-Recognitional Abilities and Perceptual Knowledge. In Haddock, A. & Macpherson, F. (Eds), *Disjunctivism: Perception, Action, Knowledge* (pp. 330–47). Oxford: Oxford University Press.
- Miracchi, L. (2015). Competence to Know. *Philosophical Studies* 172: 29–56.
- Neta, R. (2003). Contextualism and the Problem of the External World. *Philosophy and Phenomenological Research* 66: 1–31.
- Neta, R. (2008). What Evidence Do You Have? *British Journal for the Philosophy of Science* 59: 89–119.
- Neta, R. (2016). Perceptual Evidence and the Capacity View. *Philosophical Studies* 173(4): 907–14.
- Pautz, A. (2016). What Is My Evidence that Here Is a Cup. *Philosophical Studies* 173(4): 915–27.
- Plantinga, A. (1993). *Warrant and Proper Function*. Oxford: Oxford University Press.
- Pollock, J. (1974). *Experience and Justification*. Princeton, NJ: Princeton University Press.
- Pryor, J. (2000). The Skeptic and the Dogmatist. *Noûs* 34: 517–49.
- Russell, B. (1948). *Human Knowledge: Its Scope and Limits*. New York: Simon and Schuster.
- Schellenberg, S. (2011a). Perceptual Content Defended. *Noûs* 45: 714–50.
- Schellenberg, S. (2011b). Ontological Minimalism about Phenomenology. *Philosophy and Phenomenological Research* 83: 1–40.
- Schellenberg, S. (2013). Experience and Evidence. *Mind* 122: 699–747.
- Schellenberg, S. (2014a). The Relational and Representational Character of Perceptual Experience. In Brogaard, B. (Ed.), *Does Perception Have Content?* (pp. 199–219). New York: Oxford University Press.
- Schellenberg, S. (2014b). The Epistemic Force of Perceptual Experience (with a response by Alex Byrne). *Philosophical Studies* 170(1): 87–100.
- Schellenberg, S. (2016a). Phenomenal Evidence and Factive Evidence. Symposium with comments by Matt McGrath, Ram Neta, and Adam Pautz. *Philosophical Studies* 173: 875–96.
- Schellenberg, S. (2016b). Phenomenal Evidence and Factive Evidence Defended: Replies to McGrath, Neta, and Pautz. *Philosophical Studies* 173: 929–46.
- Silins, N. (2005). Deception and Evidence. *Philosophical Perspective* 19: 375–404.
- Sosa, E. (1980). The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge. *Midwest Studies in Philosophy* 5: 3–25.

- Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. New York: Oxford University Press.
- Sosa, E. (2010.) How Competence Matters in Epistemology. *Philosophical Perspectives* 24: 465–75.
- Sosa, E. (2015). The Unity of Action, Perception, and Knowledge. In Sosa, E. (Ed.), *Judgment and Agency* (pp. 7–33). Oxford: Oxford University Press.
- Weatherson, B. (2003). What Good Are Counterexamples? *Philosophical Studies* 115: 1–31.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Zagzebski, L. (1996). *Virtues of the Mind*. Cambridge: Cambridge University Press.