

The Origins of Perceptual Knowledge

Susanna Schellenberg (Rutgers)

Every epistemic property has a ground. If a mental state has an epistemic property then the question arises in virtue of what the mental state has that epistemic property. Drawing on work from cognitive psychology, I argue that the ground of the epistemic force of perceptual states lies in properties of the perceptual capacities that constitute the relevant perceptual states. More specifically, I argue that the ground of a perceptual state's epistemic force lies in the function of perceptual capacities to discriminate and single out mind-independent particulars in the environment. Perceptual states have epistemic force due to being systematically linked to mind-independent, environmental particulars via the perceptual capacities that constitute the perceptual states.

I call this view *capacitivism*, since the notion of a capacity is explanatorily basic: It is because a given subject is employing a mental capacity with a certain nature that her mental states have epistemic force. Thus, capacitivism shows how the epistemic force of experience is grounded in metaphysical facts about experience. I show how this view allows us to acknowledge internalist insights by arguing that mental states are constituted by the mental capacities employed which in turn provide the mental state with its epistemic force. I show moreover how this view acknowledges insights of knowledge-first views by arguing that there is a metaphysical primacy of employing mental capacities in the good case over the bad case. In doing so, I develop an epistemic externalist account that entails no commitment to either epistemic disjunctivism or reliabilism.

In Section 1, I develop the notion of capacity and then show how this notion of capacities grounds the epistemic force of perceptual states. In Section 2, I develop a sufficient evidence requirement for perceptual knowledge. In Section 3, I show how the account of perceptual knowledge developed fares in the face of range of cases. Along the way, I show how capacitivism differs from reliabilism, knowledge-first views, and virtue epistemology.

1. Grounding the Epistemic Force of Perceptual States

Perceptual experience is fundamentally a matter of employing perceptual capacities by means of which we discriminate and single out particulars in one's environment. Such discriminatory activity allows for border and edge detection, scene segmentation, and region extraction.¹ A perceptual capacity is a kind of discriminatory, selective capacity that we employ in perception, hallucination, or illusion. Such capacities are understood naturalistically, in terms of their natural function, namely their function to discriminate and single out mind-independent, environmental particulars. These particulars include objects, such as, tigers, trains, or trace elements, but also events, such as, a tiger sprinting towards one, as well as property-instances, such as, a tiger's reddish-orange fur and black stripes. I will argue that perceptual states provide evidence in virtue of the metaphysical nature and function of the capacities that constitute perceptual states.

To set the stage for this view, it will be necessary first to take a closer look at perceptual capacities. An analysis of perceptual capacities will include an analysis of the function of perceptual capacities, the individuation conditions of perceptual capacities, the possession conditions of perceptual capacities, and the relation between employing a perceptual capacity while fulfilling the function of the capacity and employing the same perceptual capacity while failing to fulfill the function of the capacity.

The *function of a perceptual capacity* is to discriminate and single out mind-independent particulars of a specific kind in our environment. For example, I possess the perceptual capacity to discriminate and single out red from other colors. The function of this capacity is to discriminate and single out red surfaces from surround colors. Even though we may employ a capacity while failing to single out any particular that the capacity functions to single out, perceptual capacities have the function to discriminate single out particulars in the environment rather than to fail to do so. 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

¹ For a defense of this idea, see Schellenberg 2016. See also, Julesz 1981 and Krümmenacher and Grubert 2010.

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 having that capacity being reliable.

We can specify the *individuation conditions of a perceptual capacity* as follows: A perceptual capacity is individuated by the mind-independent particulars that the capacity functions to single out. This individuation condition implies that the perceptual capacity that functions to single out instances of red differs from the perceptual capacity that functions to single out instances of vermilion. So there will be a perceptual capacity to discriminate and single out red and a distinct perceptual capacity to discriminate and single out vermilion. Indeed, there will be yet another perceptual capacity to discriminate and single out ruby red and any number of other shades of red. Not only are there perceptual capacities that are more fine-grained than the perceptual capacity to discriminate and single out red, there are perceptual capacities that are more general. For example, most humans possess the perceptual capacity to discriminate colors. Similarly, most humans possess the perceptual capacity to discriminate movement, quantities, magnitudes, as well as the perceptual capacity to discriminate distance. So perceptual capacities can be more or less general and more or less specific. Perceptual capacities differ not only in the determinates they function to single out, they differ moreover in the class of particulars they function to single out. There are perceptual capacities that function to discriminate and single out objects of a specific kind. Others function to discriminate and single out property-instances of a specific kind. Yet others function to discriminate and single out movements and other events.

The *possession conditions of a perceptual capacity* can be given with an asymmetric counterfactual analysis. A subject S possesses a perceptual capacity C if and only if the following counterfactual is true of S: S would be in a position to discriminate and single out a particular of the kind that C functions to single out if S were perceptually related to such a particular, assuming (i) S is

perceptually capable (awake, alert etc.) and (ii) assuming no finking, masking, or other exotic case obtains, and (iii) where S being perceptually related to a particular means that (a) the situational features are such that the particular is perceivable to S (good lighting conditions etc.) and (b) the particular are spatially and temporally related such that S is in a position to gain information about S via her perceptual organs.² The qualification that the subject is perceptually capable rules out cases in which the subject is not at that particular moment capable of employing her perceptual capacity to single out the particular to which she is perceptually related (perhaps because she is intoxicated), even though she is generally capable of employing her perceptual capacity to single out the particular to which she is perceptually related. It is important to note that this asymmetric counterfactual analysis of perceptual capacities allows that a perceiver can possess the capacity to discriminate red from other colors, but never employ the capacity because she is never in an environment that includes red objects.

A perceptual capacity is distinct from the employment or exercise of the capacity. While the capacity is a kind of mental tool, the employment of the capacity is an activity. 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. If a subject possesses a perceptual capacity, she can employ that capacity while either fulfilling the function of the capacity or while failing to fulfill the function of the capacity, such that there is no difference at the level of employing the perceptual capacity, but only a difference at the level of fulfilling the function of the capacity. Not only are perceptual capacities fallible, employing perceptual capacities yields a mental state that is either accurate or defective (and thus guaranteed to be false). In this respect, the view here differs

² The inference from a claim about perceptual capacities to a counterfactual fails in finking, masking and similarly exotic cases. However, all the standard ways of fixing the disposition-to-counterfactual inference can be exploited for the capacity-to-counterfactual inference. See in particular Lewis 1997. Finding a formulation of the capacity-to-counterfactual inference that is indefeasible in light of all possible finking, masking, and similarly exotic cases would be a project of its own. Therefore, I will here work on the assumption that no such exotic cases obtain. This assumption is independently plausible.

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.³

But while employing capacities is fallible, this does not mean that the good case and the bad case are on a par. Any employment of capacities in the bad case is derivative from their employment in the good case. After all, while the employment of a capacity in the bad case fails to fulfill the function of the capacity. By contrast, the employment of the same capacity in the good case fulfills the function of the capacity.

With the distinction between the capacities employed and what, if anything, they single out in place, we can distinguish two ways of individuating perceptual states. On one way of individuation, we focus only on what perceptual capacities are employed. On a second way of individuation, we focus on what perceptual capacities are employed along with the environment in which they are employed. To clarify the distinction, consider two subjects who are each employing the very same perceptual capacity C. Let's call one subject Percy and the other Hallie. Percy is perceiving, while Hallie is hallucinating. So Percy employs the perceptual capacity C and successfully singles out what he purports to single out and what C functions to single out. Hallie employs the perceptual capacity C, but fails to single out what she purports to single out. On the first way of individuation, the relevant features that characterize Percy and Hallie's mental states are the same. The same perceptual capacity is employed. On the other way of individuation, the relevant features that characterize their mental states are different: Percy successfully singles out what he purports to single out, while Hallie fails to single out what she purports to single out.

As I have argued in a series of papers, these two ways of individuating perceptual states allows us to distinguish two ways of individuating perceptual content and yields two kinds of epistemic evidence: phenomenal evidence and factive evidence. *Phenomenal evidence* is determined by how our environment sensorily seems to us when we are experiencing. *Factive evidence* is determined by the perceived particulars such that the evidence is guaranteed to be an accurate guide to the environment.

³ 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.

Employing perceptual capacities yields a mental state that provides us with phenomenal evidence regardless of whether we are in the good or the bad case. Employing such capacities in the good case provides us not only with phenomenal evidence but, moreover, with a second stronger type of evidence, namely factive evidence. So in the good case, perceptual experience provides us with both phenomenal and factive evidence, while in the bad case, perceptual experience provides us only with phenomenal evidence.⁴

Employing perceptual capacities yields a content type. The content type is constituted by the perceptual capacities employed and covaries one-to-one with the phenomenal character of the experiential state. Employing perceptual capacities while either fulfilling the function of the capacities employed or failing to fulfill their function yields a token content. 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 experiencing subject fails to single out what she purports to single out and thereby fails to fulfill the function of the capacity. 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.⁵ This distinction between a content type and a token content allows us to put the distinction between phenomenal evidence and factive evidence on firmer footing. While 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. There is no factive evidence in the bad case because the capacities were not employed successfully and the ensuing token content is defective.

Why does employing perceptual capacities yield mental states with epistemic force? According to capacitivism, phenomenal states provide us with evidence since phenomenal states are

⁴ For a defense of the distinction between phenomenal evidence and factive evidence, see Schellenberg 2013, 2014-a, 2016a, and 2016b. For a critical discussion, see Byrne 2014, McGrath 2016, Pautz 2016, and Neta 2016.

⁵ For a defense of this view of perceptual content, see Schellenberg 2011 and 2014-b.

systematically linked to the particulars that the relevant perceptual capacities function to single out. If a subject's environment sensorily seems to contain F particulars, then she is in a phenomenal state that is constituted by employing perceptual capacities that function to single out F particulars. If a subject is in a phenomenal state that is constituted by employing perceptual capacities that function to single out F particulars, then she is in a phenomenal state that provides phenomenal evidence for the presence of F particulars. This is the case, even if the subject is not in fact perceptually related to an F particular. So employing the capacity that functions to single out F particulars provides the subject with phenomenal evidence as of an F particular even if the subject fails to fulfill the function of the capacity.

Both factive and phenomenal evidence have their rational source in the perceptual capacities employed in experience. So what is new about the account developed here is that it provides a unified account of the internal and external elements of perceptual evidence and their common rational source. More specifically, the source of both kinds of evidence lies in the perceptual capacities employed that in virtue of their function to discriminate and single out particulars are systematically linked to environmental particulars. The notion of systematic linkage in play is understood in terms of a metaphysical and explanatory primacy notion rather than a reliabilist notion. 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 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 the capacity view, 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.

In this way, perceptual capacities are analyzed in terms of their natural function, namely their function to discriminate and single out particulars in the environment. In this way, I am grounding the epistemic force of experience in the metaphysical properties of experience.

Capacitivism entails that having the same sensory experience is compatible with having different evidence. After all, having the same sensory experience is compatible with being in mental states with different content if the relationship between the phenomenal character and the content of a mental state is not identity but rather supervenience. If the phenomenal character of a mental state supervenes on the content of the mental state, there can be differences in content that are not reflected in phenomenal character.⁶ Moreover, if the content of one's mental state determines one's evidence, then a view on which the phenomenal character of a mental state supervenes on the content of the mental state will allow that two beings whose experiences have the same phenomenal character could fail to have the same evidence. Consequently, capacitivism respects the key internalist intuition that mental duplicates will have the same evidence. However, someone in the bad case and someone in the good could never be mental duplicates even though they have the same phenomenal evidence.

One might argue that perceptual states provide us with evidence since the perceptual capacities by which they are constituted are reliable. According to capacitivism, 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. Reliability simply plays no role in my account. This is a good thing, since perception is not a particularly reliable faculty. 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

⁶ For a defense of such a weak representationalist view, see Schellenberg 2010 and 2011.

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 case (Davidson 1987: 443-4).

Swampman. Donald Davidson goes hiking in a swamp and is struck and killed by a lightning bolt. Simultaneously, a second lightning bolt spontaneously rearranges molecules such that they take on exactly the same form that Davidson's body had at the moment of his death. The resulting Swampman behaves exactly like the original author of "Radical Interpretation." His brain is structurally identical to that which Davidson had at the moment of his death. Swampman walks out of the swamp, returns to Davidson's office at Berkeley, and writes the same essays Davidson would have written. Swampman has no causal history.

Does Swampman possess perceptual capacities? Do his perceptual experiences provide him with phenomenal and factive evidence? According to the view developed here, no past experiences are necessary to possess perceptual capacities. Since possessing perceptual capacities does not depend on a history of proper usage—or any usage for that matter—and since perceptual capacities are not understood in an evolutionary way, there is no reason to think that Swampman could not possess the perceptual capacities in play. After all, the condition for their possession is understood counterfactually: if one possesses the capacity to single out red, then one would be able to single out an instance of red, were one related to such an instance. Since Swampman possesses the relevant perceptual capacities, he can through perception gain phenomenal and factive evidence about the swampy world around him. So given the notion of capacity in play, just as Swampman's heart has the function pump blood, his perceptual capacities have the function single out particulars. This is the case even though Swampman himself has no evolutionary history. 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.

By arguing that it is in virtue of the primacy of the good over the bad case that the mental states yielded by employing capacities have epistemic force, capacitivism makes room for an

externalist account of the epistemic role of perceptual experience that does not depend on and does not entail reliabilism (Goldman 1979). Putting this together, capacitivism is an externalist view that does not invoke reliability, remains steadfastly naturalistic, and recognizes a common element in perception and hallucination.

2. A Sufficient Evidence Requirement for Knowledge

We are now in a position to take a closer look at the relation between phenomenal evidence, factive evidence, and perceptual 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 accurate.⁷ 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, capacitivism entails that knowledge is a mental state. In this way, capacitivism 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.⁸

Now let's assume standardly that for subject 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

⁷ For a detailed discussion of the accuracy conditions of perceptual content, see Chapter 5, Section 5.4.

⁸ For the idea that knowledge is a mental state, see Williamson 2000 and Nagel 2013.

position to formulate the following necessary and jointly sufficient conditions for S to have perceptual knowledge that p :

Subject 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 any belief condition on knowledge. Orthodoxy has it that one cannot know that p without believing that p . Capacitivism 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. 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. A second reason for why it is attractive to give up a belief condition on knowledge when the knowledge in question is perceptual knowledge is that an experiential state that falls short of knowledge does not amount to a belief. If we fail to perceive what it seems to us we are perceiving and so fail to acquire knowledge, this does not entail that we believe what it seems to us we are perceiving. We may be hallucinating without forming any beliefs based on our hallucination.⁹

Orthodoxy has it that mental states that fall short of knowledge are justified beliefs. If an experiential state that falls short does not amount to a belief, how then should we think of this mental state? We can think of them as justified mental states. 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

⁹ 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.

But for the reasons given above, it is attractive to give up the belief condition on knowledge.

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. 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.

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 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.

I have argued that we gain knowledge of our surroundings by successfully employing capacities that function to single out particulars in our environment. 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.¹⁰ 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 capacitivism, 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. A perceptual capacity is rather a low-level mental capacity that functions to discriminate, single out, and in some cases classify a particular kind, such as instances of red. Such perceptual capacities are

¹⁰ See Miracchi 2015 for such a view.

determined by general, functional relations between the organism and its environment—for instance, global patterns of the organism’s response to its environment. 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.¹¹ By contrast, the employment of the same capacity in the good case is constitutively a success and provides the subject with factive evidence and thus with knowledge.

According to capacitivism, 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. Insofar as both hallucinations and at least some perceptions provide the experiencing subject with evidence, and so justification for any belief she might form, capacitivism 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, capacitivism 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. After all, both the good and the bad case are brought about by employing perceptual capacities. We get at how the world is via perception in a particular way, namely by employing perceptual capacities. And even when we fail to get at how the world is (and so are in the bad case), we are employing perceptual capacities by means of which we aim to get at how the world is. In this way, capacitivism 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.

It will be helpful to contrast the view developed from Williamson’s knowledge-first view. According to Williamson, evidence is a known proposition and knowledge is a mental state.

¹¹ In this respect among others, the view developed here differs from Sosa 2007, Bergmann 2008, Burge 2010, and Greco 2010.

Evidence is the object of the mental state, namely, a proposition or a set of propositions. Since evidence is a known proposition, there is no room on Williamson's view for evidence provided directly through experience in the bad case. After all, in the bad case there are no true propositions provided directly through experience. On Williamson's knowledge-first view, we have only introspective evidence in the bad case, that is, known propositions about how things seem to us.

Williamson's knowledge-first view is an alternative way of arguing that we have more evidence when we perceive than when we hallucinate than *ceteris paribus* by showing that there are two distinct facts that can figure as the truthmakers of perceptual content: facts about the experience and facts about the environment in which one is experiencing. We have evidence that consists of true propositions when we are hallucinating, namely, introspective evidence about how the environment seems to us. Such an approach restricts the evidence we can gain through perceptual experience to factive evidence; however, the factive evidence includes not just perceptual evidence, but also introspective evidence. So the evidence we gain through perceptual experience is either factive with regard to our environment or with regard to our experience.

On both Williamson's knowledge-first view and the capacity view, we have more evidence when we perceive than when we hallucinate *ceteris paribus*. On Williamson's knowledge-first view, a perceiver has factive perceptual evidence and factive introspective evidence, while a hallucinating subject has only factive introspective evidence. On the capacity view, a perceiver has phenomenal and factive perceptual evidence, while a hallucinating subject has only phenomenal perceptual evidence. Like the knowledge-first view, the capacity view holds that knowledge is a mental state and that we have some evidence in the bad case, but that we have more evidence in the good case. There are five key differences between the two views.

One key difference is that the fundamental explanatory notion of the capacity view 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, the capacity view 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. The capacities employed yield phenomenal evidence that we have in both the good and the bad case.

A third key difference is that on the capacity view, we have at least some evidence in common between the good and the bad case. In both cases, we have phenomenal evidence. So while the capacity view and the knowledge-first view both have it that we have some evidence in the bad case, but that we have more evidence in the good case, on the capacity view the evidence we have in the bad case is evidence that we also have in the good case. Thus the capacity view is non-disjunctivist with regard to the content, the metaphysics, and the epistemology of perception.

A fourth key difference is that on the knowledge-first view, the methods employed in gaining evidence are infallible in that evidence always receives probability 1, so that the mental state yielded is always factive. On the capacity view, the methods employed in gaining evidence are not infallible and while the mental state yielded is factive in the good case, it is not factive in the bad case.

Finally, a fifth 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 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 the capacity view 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, Williamson's knowledge-first view over-intellectualizes the way we get evidence in the bad case.

The second reason why it is more attractive to say that experience provides us with phenomenal evidence directly even in the bad case without retreating to introspective evidence hinges on a second and more pressing over-intellectualization worry with Williamson's knowledge-

first view. According to Williamson, 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. The capacity view does not face these over-intellectualization 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.

The third reason why it is more attractive to say that experience provides us with phenomenal evidence directly even in the bad case without retreating to introspective evidence is that 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. While perceptual evidence stems from perception, introspective evidence stems from introspection. By contrast, the capacity view shows that the source of both factive and phenomenal evidence is our perceptual experience. Indeed, the capacity view 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. In short, with the knowledge-first view, capacitivism explains the bad case in terms of the good case. But against the knowledge-first view (and any other disjunctivist view), capacitivism 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.

3. Gettier Cases: Phenomenal Evidence without Factive Evidence

According to capacitivism, 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_I in our environment and we in fact discriminate and single out that particular at L_I , we gain factive evidence and thus perceptual knowledge of the particular at L_I .

Now, there are, famously, cases in which a subject may form a justified true belief on the basis of her experience without having knowledge. One such case is the following:

Veridical Hallucinations. Vernon suffers a veridical hallucination as of a white cup. He hallucinates a white cup at location L_I and as it so happens, there is a white cup behind a screen just where he hallucinates a white cup to be. Since the actual white cup is behind a screen, he could not be perceptually related to it. So the cup at L_I is not causally relevant in bringing about his hallucination. On the basis of his hallucination, he forms the belief that there is a white cup at L_I .¹²

The standard analysis of such Gettier cases is that the Gettiered subject has a true belief and sufficient evidence for knowledge, but still lacks knowledge. On my view this analysis is too simplistic. I 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.

I argued that if we employ the capacity to discriminate and single out a particular at location L_I and succeed in fulfilling the function of the capacity, then we gain factive evidence and thus perceptual knowledge of the particular at L_I . In a Gettier case, the particular at L_I is not in fact the

¹² This case differs from Bonjour's (1980) Norman case, since there is no assumption that Vernon comes at her hallucination by way of a reliable process. By contrast, clairvoyant Norman comes to his beliefs by way of a reliable process that however is not based on any evidence or reasons.

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 L_1 . That is the case even if it seems to the subject that she is successfully discriminating and singling out that particular at L_1 . Moreover that is the case even if there is a particular at location L_2 that the subject could have successfully singled out but did not in fact successfully single out, as in the Veridical Hallucination case.

Capacitivism provides a way to do justice to internalist and externalist intuitions about the veridical hallucinations case. According to capacitivism, Vernon employs the very same perceptual capacities that he would employ, were he perceiving a white cup. As a consequence, he is in a mental state with content that provides him with phenomenal evidence that there is a white cup at L_1 . So capacitivism can explain why he is not blameworthy for his belief that there is a white cup at L_1 . Not only is Vernon excused in believing that there is a cup at L_1 , since he has phenomenal evidence that there is a cup at L_1 and since phenomenal evidence justifies general propositions, Vernon has some justification for his belief. But Vernon does not have factive evidence, since he is not perceptually related to the white cup at L_1 . Since his evidence is merely phenomenal, he is not justified in believing the singular proposition “that is a white cup.” So 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.

The same analysis can be given of perceptual Gettier cases in which the gettiered subject perceives something that in fact pertains to what she believes. If a subject’s experiential state was not formed on the basis of employing capacities such that those capacities fulfill their function, she will not have factive evidence but will only have phenomenal evidence for what she seems to be perceiving. Since she fails to have factive evidence for what she seems to be perceiving, she will fail to have knowledge. Consider the following case adapted from Chisholm 1966:

Rock Sheep. Malika is in a field looking at something that looks like a sheep. Malika forms the belief “There is a sheep in the field”. What seems to her to be a sheep is in fact a dog disguised as a sheep. If this were a full description of the case, Malika’s belief “There is a sheep in the field” would be false. However, as it happens, there is a sheep

in the field a few feet away hiding behind a rock. So Malika's belief "There is a sheep in the field" happens to be true.

The standard analysis of this case is to say that Malika's belief "There is a sheep 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 sheep in the field and the belief was formed based on ordinary perceptual processes albeit on grounds of perceiving a dog disguised as a sheep rather than the actual sheep on the field. Why think that Malika has sufficient evidence for knowledge? The reasoning seems to stem from the internalist idea that Malika is doing just the same thing from her own perspective as a successful perceiver would be doing in a simple case of seeing a sheep in the field (with no disguised dogs or other complications).

According to capacitivism, the verdict on the case is different. Malika's belief has some justification. After all, Malika employs her capacity to single out a sheep and in virtue of this, she is in a mental state that is intentionally directed at a sheep. Being in this mental state provides her with some evidence for her belief that there is a sheep in the field. It provides her with phenomenal evidence. So Malika has a justified true belief. However, since Malika is not in fact perceptually related to a sheep and so does not single out a sheep (despite it seeming to her that she is doing just that), she does not have *factive* evidence of a sheep. The fact that there happens to be a sheep a few feet away behind a rock does not mean that Malika has *factive* evidence of a sheep on the field, since she did not single out that sheep. She only singled out the dog disguised as a sheep. So Malika has phenomenal evidence but fails to have *factive* evidence that there is a sheep in the field. Since phenomenal evidence is not sufficient evidence for knowledge, Malika does not have sufficient evidence to know that there is a sheep in the field.

What happens to the widespread intuition that Malika has a justified true belief without knowledge? We can acknowledge that Malika's belief "There is a sheep in the field" is a justified true belief. After all, the belief is true since there is a sheep in the field and the belief is justified since it is supported by Malika's phenomenal evidence. More generally, 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, Malika does not believe truly in virtue of successfully employing her capacity to single out a dog.

It will be helpful to compare capacitivism with key commitments of virtue epistemology. Capacitivism 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. 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.¹³ 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.

Like virtue epistemology, capacitivism is a kind of anti-luck theory in that it rejects the JTB+ approach of analyzing knowledge in terms of a combination of independent epistemic properties. In these respects capacitivism and virtue epistemology are structurally similar, but a comparison of key points reveals deep differences. There are four key differences between capacitivism and virtue epistemology.

One key difference is that virtue epistemologists treat knowledge as a particularly successful or valuable case of belief, while according to capacitivism, knowledge is the primary case. In this respect,

¹³ For a reliabilist virtue epistemology, see Sosa 1980, 2007, 2010 and Greco 2009, 2012. For a responsibilist virtue epistemology, see Zagzebski 1996.

capacitivism follows knowledge-first views. According to capacitivism, we know in virtue of successfully employ capacities that function to single out particulars in our environment. By contrast, if we fail to employ those very same capacities successfully, we do not have sufficient evidence for knowledge. When we fail to employ the capacities successfully, we may not be in a belief state. 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 capacitivism and virtue epistemology is over whether the relevant capacities (the successful employment of which generates knowledge) are themselves to be understood in normative or naturalistic terms.¹⁴ For the virtue epistemologist, capacities are understood normatively, as intellectual virtues. While virtue reliabilists can allow that some performances have aims in virtue of having biological functions, the relevant capacities are nonetheless understood normatively. According to capacitivism, 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* believes 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 in responsibilist (Zagzebski 1996) and in reliabilist terms (Sosa 1980, 2007, 2010, Greco 2009, 2012). Among reliabilist interpretations, we can distinguish further between views that understand the “in virtue of” relation causally (Sosa 2007: 95ff) and those that understand it dispositionally (Sosa 2010, 2015). According to capacitivism, 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 the content of *S*'s mental state *p* has the content it has in virtue of *S* having

¹⁴ For an argument that natural functions need not be understood in normative terms, see Broome (2013).

employed her capacity, and in virtue of employing this capacity to successfully single out what she purports to single out.

According to capacitivism, the capacity employed that brings about knowledge that p is not the capacity to believe truly that p or to know that p . It is the capacity to discriminate and single out a particular α in the environment. This capacity will bring about a factive mental state if and only if one is perceptually related to α while employing the capacity to discriminate and single out α , thereby discriminating and singling out α . So by employing this capacity successfully one knows that α is present.

It will be helpful to show how these differences 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 capacitivism, 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 to successfully single out what she purports to single out.

The fourth key difference holds only with regard to reliabilist virtue epistemology. Most versions of virtue epistemology currently on the table develop the notion of competences (or abilities, capacities, virtues) in reliabilist terms (e.g. Sosa 1980, 2007, 2010 and Greco 2009, 2012). Insofar as the reliabilist virtue epistemologist grounds the epistemic force imbued by virtues in the reliability of those virtues, virtue epistemology is subject to all the well-known problems of reliabilism.¹⁵ Since capacitivism is distinctly non-normative and non-reliabilist it avoids these problems.

However, despite these differences, capacitivism is compatible with a broadly virtue-based epistemology. Indeed, by adopting the asymmetric dependence principle of capacitivism while giving

¹⁵ For example, Lehrer's (1990) Truetemp counterexample: Mr. Truetemp has, unbeknownst to him, a temperature-detecting device implanted in his head that regularly produces accurate beliefs about the ambient temperature.

up on any reliabilist commitments, virtue epistemologists could avoid the reliabilist problems of their view. Moreover, while Sosa and other virtue epistemologists do not appeal to the difference between phenomenal and factive evidence, nothing in their view should preclude them from accepting the two levels of evidence that I propose.

Coming back to capacitivism: I argued that 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 her belief but fails to have factive evidence.

There is at least one case that has traditionally been understood to be a Gettier case that would not count as a Gettier case according to capacitivism, namely Ginet's and Goldman's barn façade county case (Goldman 1976: 772-3).¹⁶ But this is a good outcome. To show why let's first consider the case.

Barn Façade County. Henry 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.

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 capacitivism 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 sees a real barn. One might argue that while Henry has perceptual knowledge, he lacks a more sophisticated kind of knowledge. Taking that route

¹⁶ Ginet never put the example in writing, but Goldman credits him with the example.

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 perceptual knowledge, one might argue that the sheer existence of the barn façade entails that it is unreasonable for him to believe that there is a barn, despite the fact that Henry is not aware of the fact that he is in barn façade county. But even if we grant this, it will not affect what factive perceptual evidence he has regarding the barn. It will affect only what he ends up being justified in believing, all things considered.¹⁷ Being in barn façade county might affect the degree to which Henry is *ultima facie* justified in believing that there is a barn in front of him, but given capacitivism, there is no reason to think that it will affect the factive evidence he has of the barn in virtue of seeing the barn. So there is no challenge to Henry's *prima facie* justification for the presence of barn that he gains in virtue of seeing the barn.

Once Henry gains information that he is barn façade county the situation changes. If Henry gains the information that he is in barn façade county, then he gains new evidence. But even though he gains evidence that he is in barn façade county, this does not entail that his perceptual evidence is weakened.¹⁸ After all, even if he has defeaters they 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 in believing, all things considered. So even in this case, while gaining information that he in barn façade county might affect the degree to which Henry is *ultima facie* justified in believing that there is a barn in front of him, there

¹⁷ According to Lackey (2007 and 2009), *S* can know that *p* without deserving credit for truly believing that *p*. If one takes up this distinction between knowing that *p* and deserving credit for truly believing that *p* and if one assumes that successfully employing perceptual capacities is not something one deserves credit for, then one could say that Henry knows that there is a barn in front of him, while denying that he deserves credit for truly believing that there is a barn in front of him. This in turn would allow one to distinguish Henry's epistemic situation from the situation of someone who sees a barn in a county in which there are no fake barns. One could however argue that one does deserve credit for successfully employing perceptual capacities.

¹⁸ See Lasonen-Aarnio (2010) for a similar view. Lasonen-Aarnio argues that one can know even if one's belief is unreasonable.

is no reason to think that gaining this information will affect the factive evidence he has of the barn in virtue of seeing the barn. Once Henry knows that he is in barn façade county, he has good reason to be suspicious of what he learns through perception. Nonetheless, he has factive perceptual evidence of the barn he is seeing in virtue of seeing the barn and in virtue of having factive perceptual evidence, he has perceptual knowledge of the barn. In short, perceptual evidence is so powerful that Henry has perceptual factive evidence regardless of whether he is consciously aware of being in barn façade county.

4. 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*.
- Byrne, A. and Hilbert, D. 2003. "Color Realism and Color Science". *Behavioral and Brain Sciences*, 26: 3-21.
- Chisholm, Roderick 1966. *Theory of Knowledge*. Englewood Cliff, NJ: Prentice-Hall.
- Davidson, D. 1987. "Knowing One's Own Mind". *Proceedings and Addresses of the American Philosophical Association* 60: 441-58.
- Feldman, R. and Conee, E. 1985. "Evidentialism". *Philosophical Studies*, 48: 15–34.
- Goldman, A. 1976. "Discrimination and Perceptual Knowledge". *Journal of Philosophy* 73: 771-91
- 1979. "What is Justified Belief?" In G. Pappas, ed., *Knowledge and Justification*. Dordrecht: Reidel.
- 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Greco, J. 2004. "Knowledge as Credit for True Belief." *Intellectual Virtue: Perspectives from Ethics and Epistemology*, M. DePaul & L. Zagzebski (eds.). Oxford: Oxford University Press.
- 2009. "Knowledge and Success from Ability". *Philosophical Studies* 142: 17–26.
- 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.
- 2012. "A (Different) Virtue Epistemology". *Philosophy and Phenomenological Research* 85: 1–26.
- Hetherington, S. 2001. *Good Knowledge, Bad Knowledge*. Oxford: Oxford University Press.
- Ichikawa, Jonathan and Matthias Steup 2012. "The Analysis of Knowledge". *The Stanford Encyclopedia of Philosophy*: <http://plato.stanford.edu/entries/knowledge-analysis/> <accessed 8/17/15>
- Julesz, B. 1981. "A Theory of Preattentive Texture Discrimination Based on First-Order Statistics of Textons". *Biological Cybernetics*, 41, p. 131–8.
- Klein, P. 1971. "A Proposed Definition of Propositional Knowledge". *Journal of Philosophy* 68: 471–482.
- 1976. "Knowledge, Causality, and Defeasibility". *Journal of Philosophy* 73: 792–812.

- Krummenacher J., Grubert A. et al, 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.
- Lackey, J. 2007. "Why we don't Deserve Credit for Everything we Know." *Synthese* 158: 345-61.
- 2009. "Knowledge and Credit." *Philosophical Studies* 142: 27-42.
- Lehrer, K. 1965. "Knowledge, Truth, and Evidence". *Analysis* 25: 168–75.
- Lewis, D. 1997. "Finkish Dispositions." *Philosophical Quarterly* 47: 143-58.
- Lycan, W. 2006. "On the Gettier Problem Problem". In *Epistemology Futures*, S. Hetherington (ed.). Oxford: Oxford University Press.
- McGrath, M. 2016. "Schellenberg on the Epistemic Force of Experience". *Philosophical Studies*.
- Millar, A. 2008. "Perceptual-Recognitional Abilities and Perceptual Knowledge". *Disjunctivism: Perception, Action, Knowledge*. A. Haddock & F. Macpherson (eds.), Oxford: Oxford University Press.
- Miracchi, L. 2015. "Competence to Know". *Philosophical Studies* 172: 29–56.
- Nagel, J. 2013. "Knowledge as a Mental State." *Oxford Studies in Epistemology* 4: 273-308.
- Neta, R. 2003. "Contextualism and the Problem of the External World". *Philosophy and Phenomenological Research*, 66: 1–31.
- 2008. "What Evidence Do You Have?". *British Journal for the Philosophy of Science*, 59: 89–119.
- 2016. "Perceptual Evidence and the Capacity View". *Philosophical Studies*.
- Plantinga, A. 1993. *Warrant and Proper Function*. Oxford: Oxford University Press.
- Pautz, A. 2016. "What is my Evidence that here is a Cup". *Philosophical Studies*.
- 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, NY: Simon and Schuster.
- Schellenberg, S. 2011. "Perceptual Content Defended". *Noûs* 45: 714–50.
- 2013. "Experience and Evidence". *Mind* 122: 699-747.
- 2014-a. "The Epistemic Force of Perceptual Experience". With a response by Alex Byrne, *Philosophical Studies*.
- 2014-b. "The Relational and Representational Character of Perceptual Experience". *Does Perception have Content?* ed. B. Brogaard, New York: Oxford University Press.
- 2016-a. "Phenomenal Evidence and Factive Evidence." Symposium with comments by Matt McGrath, Ram Neta, and Adam Pautz, *Philosophical Studies* 173: 875-96.
- 2016-b. "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.
- 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. New York: Oxford University Press.
- 2010. "How Competence Matters in Epistemology". *Philosophical Perspectives* 24: 465–475.

- 2015. “The Unity of Action, Perception, and Knowledge”. *Judgment and Agency*, E. Sosa. 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.