

Schellenberg – Experience and Evidence

Brian T. Miller

Nov 14, 2016

Introduction

Percy the perceiver has a veridical experience as of a cup on the table.

Hallie the hallucinator has a non-veridical hallucination as of a cup on the table.

Question: what evidence do Percy and Hallie have for believing that there's a cup on the table?

Internalist¹ answer: Percy and Hallie have the same evidence for their beliefs; that evidence is *phenomenal*: it's determined by how things sensorily seem

¹ SS seems to have *mentalist* internalism in mind

Externalist² answer: Percy has *introspective* evidence about how things appear to her, and she also has *factive perceptual* evidence. Hallie has the same introspective evidence as Percy, but Hallie has no *factive perceptual* evidence.

² SS takes Williamson to be a paradigm externalist for the purposes of this article

SS's proposal (first pass): Percy has both phenomenal evidence and perceptual *factive* evidence; Hallie has phenomenal evidence but not perceptual *factive* evidence.³

Both phenomenal evidence and perceptual *factive* evidence 'have the same rational source': the perceptual capacities employed in experience; this common source unifies the 'two types of perceptual evidence' theory.

³ Both Percy and Hallie have introspective evidence, but that's not the focus of this paper.

1 – Perceptual Evidence and Introspective Evidence

Purpose of this section is simply to distinguish perceptual evidence from introspective evidence.

Perceptual evidence:

- is evidence about what the environment is like
- stems from perception
- is obtained by attending to the environment
- is *direct* evidence: we don't need to attend to our experiences to have that evidence

Introspective evidence:

- is evidence of what our mental states are like
- stems from introspection
- is obtained by attending to one's mental states
- is *indirect* evidence: we must attend to our experience to have that evidence⁴

⁴ Isn't it then natural to say that that introspective evidence is direct evidence about our mental states?

2 – *The Phenomenal Evidence Argument*

SS's 'basic argument' for perceptual phenomenal evidence:

- (1) If a subject S is perceptually directed at her environment (while not suffering from blindsight or any other form of unconscious perception), then it sensorily seems to S as if her environment is a certain way (regardless of how it in fact is).
- (2) If it sensorily seems to S as if her environment is a certain way (regardless of how it in fact is), then S is in a sensory state that provides phenomenal evidence.
- (3) So, if S is perceptually directed at her environment (while not suffering from blindsight or any other form of unconscious perception), then S is in a sensory state that provides phenomenal evidence.

Premise (1) merely describes what happens when we direct our attention to the environment.

Premise (2) is defended below.

2.1 – *Premise 2: Sensory states and phenomenal evidence*

Point of this section is to argue that sensory states provide phenomenal evidence.

Consequence of that claim (with other plausible premises): phenomenal evidence supervenes on mental states.

Question: why is it rational to heed the testimony of the senses, even in the case of (non-apparent) hallucination?⁵ If perceptual evidence is the same whether you're hallucinating or perceiving accurately, then how can experience support belief in the one rather than the other?

⁵ note that if phenomenal evidence supervenes on mental states, then the only reason to think that the hallucinator and the non-hallucinator have the same evidence is that they have the same mental states. TW would reject that, since one knows (in the 'seeing-way') while the other doesn't. Hence SS seems to presuppose that the mental states that phenomenal evidence supervenes upon are not factive.

Important: not everybody thinks that rationality, and hence evidence, requires a reason to believe that experience tracks truth; those internalists will not require an explanation. SS isn't really arguing that point, but she is specifically looking for an explanation for why experience is truth-conducive, a question that even the internalist can ask, even if they think that rational belief on the basis of experience doesn't require an answer.

SS's provisional answer:

I will give support to premiss 2 by arguing that sensory states provide us with phenomenal evidence, since sensory states are systematically linked (in ways to be explained) to the particulars that they single out in the case of an accurate perception. Due to the existence of this systematic link it is rational to heed the testimony of our senses.

The thing that sensory states 'systematically link' us to is the environment, so (2) can be stated more helpfully as:

2* If it sensorily seems to S as if her environment is a certain way (regardless of how it in fact is), then S is in a sensory state that provides phenomenal evidence for the presence of particulars of the type that the sensory state is of in the good case.

Argument for (2*):

- (2a) If it sensorily seems to a subject S as if her environment is a certain way (regardless of how it in fact is), then S is in a sensory state that is systematically linked to external, mind-independent particulars of the type that the sensory state is of in the good case.
- (2b) If S is in a sensory state that is systematically linked to external, mind-independent particulars of the type that the sensory state is of in the good case, then S is in a sensory state that provides phenomenal evidence for the presence of particulars of the type that the sensory state is of in the good case.
- (2*) So If it sensorily seems to S as if her environment is a certain way (regardless of how it in fact is), then S is in a sensory state that provides phenomenal evidence for the presence of particulars of the type that the sensory state is of in the good case.

2.2 – Premise (2a): *Sensory states and perceptual capacities*

The goal of this section is to motivate (2a):

- (2a) If it sensorily seems to a subject S as if her environment is a certain way (regardless of how it in fact is), then S is in a sensory state that is systematically linked to external, mind-independent particulars **of the type that the sensory state is of in the good case.**

The tricky part is to motivate the part in bold. Since phenomenal seemings are consistent with both the good case and the bad case, then why think that a given sensory state supports beliefs that are true in the good case (e.g. *there's a cup on the table*) instead of beliefs that true in the bad case (e.g. *I'm hallucinating a cup on the table, but there probably isn't one*)?

The classic move is to identify something that's common to the two cases (e.g. sense data), assert that phenomenal seemings are 'systematically linked' to these internal, mind-dependent particulars, and hence are good evidence for facts about those particulars. But there's no reason to believe that those mind-dependent particulars are systematically related to things in the world, so phenomenal seemings provide no reason to believe that you're in the good case rather than the bad case. Hence, skepticism.

SS's move is to provide a metaphysical account of phenomenal states on which those states are the same in both the good case and the bad case (contra the disjunctivist) but there's a systematic link between a state and the external, mind-independent particulars that it's a state that exists whether you're in the good case or the bad case. So the state is shared between the good and bad cases, but the state is not neutral about which one you're in, so it justifies good-case beliefs over bad-case beliefs.

Metaphysical account of phenomenal states

The basic idea of this view is that when we perceive, we employ perceptual capacities by means of which we differentiate and single out particulars in our environment. The relevant particulars are external and mind-independent objects, events, property-instances, and instances of relations. Sensory states are understood as determined by employing perceptual capacities in a sensory mode, that is, modes such as seeing, hearing, touching, smelling, or tasting. I will argue that if a subject S's environment sensorily seems to contain F particulars to her (regardless of how it in fact is), then S is in a sensory state that is determined by employing perceptual capacities that function to single out F particulars. (709)

In the good case in which Percy perceives the white cup, he employs capacities to discriminate white things from non-white things, to differentiate cup-shaped things from non-cup-shaped things, to differ-

entiate cups from non-cups (maybe), etc. It is in virtue of employing these capacities that Percy is in a sensory state that is of a white cup, and he thereby becomes perceptually related to the cup.

In the bad case in which Hallie hallucinates the white cup she deploys exactly the same capacities, so she has exactly the same sensory state, but this time that state fails to be perceptually related to the cup (because there isn't one).

How to understand the capacities in question?

- the science supports thinking of them as 'discriminatory, selective capacities [which] are the cognitively most low-level mental capacities employed in perception... that function[] to differentiate, single out, and in some cases type the kind of particulars that the capacity is of.' (710-11)
- They're non-conceptual: animals and small kids can single out properties and objects in their environments. (711)
- they're 'necessarily determined by functional connections between perceivers and their environment' such as 'global patterns of the organism's response to its environment'. (712)
- but these patterns/ connections needn't be perfect: it's possible for a capacity to be employed by something other than the thing in the environment that it's functionally connected to, like LSD
- sensory states are determined by employing perceptual capacities, rather than the capacity itself. Since employing a capacity to select and discriminate doesn't require that one successfully select and discriminate an actual object, employing that capacity doesn't require that the environment be what's triggering the employment of the capacity, so the 'baselessness' of a token sensory state is not revealed in that state (i.e. you can't tell from the inside whether your sensory states are hallucinatory)

Provisional conclusion:

If it is right that two experiences in which *ceteris paribus* all the same perceptual capacities are employed in the same sensory mode have the same sensory character, then subjectively indistinguishable perceptions, hallucinations, and illusions will share a **metaphysically substantial common factor**. The common factor is determined by the perceptual capacities that the subject employs in a sensory mode. But... the fact that there is such a common factor does not imply that we are aware of a common factor, nor does it imply that the good case is analyzed as a conjunction of a common factor and some additional element, such as a causal perceptual relation. (712-13)

Epistemological upshot of the metaphysical account

How does SS's metaphysical account of sensory states help to explain why it's rational to heed the testimony of the senses?

Basic idea:

[S]ensory states are systematically linked to what they are of in the good case in the sense that the perceptual capacities employed in the bad case are explanatorily and metaphysically parasitic on their employment in the good case. (714)

The good case is *explanatorily prior* to the bad case because the capacities employed in the bad case are always explained in terms of what would have happened if those capacities had been employed in the good case: Hallie the hallucinatory employs capacities that would have picked out a white cup on the table had she been in Percy's position.

The good case is *metaphysically prior* to the bad case because you can⁶ possess the capacities employed in the bad case only 'in virtue of being the kind of being that could employ those very capacities in the good case'. This is SS's **Metaphysical Primacy Thesis**.

⁶ metaphysical possibility intended here

- Since the capacities are defined functionally, '[i]t would be unclear what it would mean to possess a discriminatory, selective capacity, the very function of which is to single out a kind of particular, without being in a position to single out such a particular when perceptually related to one'. (715)
- '...while discriminatory, selective capacities can be employed in hallucination, they are necessarily determined by relations between perceivers and their environment in so far as the function of the capacity is to differentiate and single out, say, instances of red in perception. In this sense, there is a metaphysical priority of the good over the bad case'. (715)

Objection: suppose that's right, and the *mere employment* of the capacity in the bad case is parasitic on the employment of the capacity in the good case. Nonetheless, the Metaphysical Primacy Thesis entails a symmetry when it comes to the *successful employment* of that capacity:

- if one possesses the capacity to single out the white cup, then you'll (typically) succeed in singling out the white cup in good case.

But,

- if one possesses the capacity to single out the white cup, then you'll fail in singling out the white cup in the bad case.

The Metaphysical Primacy Thesis is intended to establish an asymmetry between the good case and the bad case that makes it rational to trust the testimony of the senses, but when it comes to the successful employment of those senses there's no asymmetry.

SS's response:

The asymmetry buttressing the thesis is an asymmetry of function. Perceptual capacities function to single out particulars. They do not function to fail to single out particulars. It is compatible with this that they may be employed in hallucination thereby failing to single out particulars... The heart has the function to pump blood. It does not have the function to fail to pump blood – though in the bad case it will fail... 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. Likewise, perceptual capacities have the function to single out particulars in the environment. They do not have the function to fail to single out particulars.' (716)

BTM: Ursula was genetically engineered to have a capacity for detecting unicorns; that's its function. It does not function to fail to detect unicorns. So the good case (in which there are unicorns) is more explanatorily and metaphysically basic than the bad case (no unicorns). So it's rational for Ursula to trust her unicorn-dar.

Is that case possible? SS waffles a bit. On p. 717 she seems sympathetic to the idea that I couldn't possess a capacity to single out a particular unless someone, somewhere had to have perceived a particular of that kind, but she doesn't commit to it. If that's right, then since no one has ever seen a unicorn, a capacity to single out unicorns is impossible.

Problem: How does this process of capacity-acquisition get going?

Example: no one anywhere has ever seen a particular that's the super-missing shade of blue. Presumably it's possible for a first person to see it, meaning that they must have the capacity to single out such a particular. But on this account that's impossible. Rinse and repeat for every kind of particular: at some point in time, no one had ever seen it, so no one would have the capacity to single it out, so the chain never gets going. So no one ever has the capacity to pick out any particulars. (Compare: Transcendental Aesthetic) [End BTM]

New case (718): Swampman, who came into existence when lightning struck the swamp, so has no causal history before that. Does Swampman have perceptual capacities? Here SS's concern seems to be about how such capacities could have acquired their *functions*.

She's sympathetic to an evolutionary account, but Swampman didn't evolve. Capacities could be innate, so it's at least possible. SS's official position: neutrality.

So in what sense does a capacity that functions to single out X depend upon X?

- S's possession of a capacity to pick out X does not depend on S's ever been perceptually related to an X. It could be innate, acquired through testimony or imagination
- It's possible (at least, consistent with the argument in the paper) that S could possess a capacity pick out X even though there are no actual X's.
- S can employ a capacity to pick out X even if there is no X present; that's exactly what's happening in cases of hallucination.

2.3 – *Premise (2b): Phenomenal evidence and systematic linkage*

The purpose of this section is to defend premise (2b):

- (2b) If S is in a sensory state that is systematically linked to external, mind-independent particulars of the type that the sensory state is of in the good case, then S is in a sensory state that provides phenomenal evidence for the presence of particulars of the type that the sensory state is of in the good case.

(2b) depends on two principles:

'The first principle is that if sensory states are systematically linked to what they are of in the good case in the sense specified, then it is epistemically rational to heed the testimony of these sensory states.'
(719)

This principle was defended in §2.2, the basic idea being that the systematic link is established by the metaphysical and epistemic dependence of the employment of the capacity in the bad case on the employment of the capacity in the good case, which is in turn established by the function of the capacity being determined by what happens when you employ that capacity in the good case.

NB: this view does not presuppose any form of reliabilism: seemings are evidence, no matter how often they lead us astray.⁷

'The second principle is that if it is epistemically rational to heed the testimony of sensory states, then they provide evidence.'
(719)

⁷ though presumably if you *know* that they frequently lead us astray then that evidence is defeated

Follows from an uncontroversial premise: that ‘if it is epistemically rational to heed x in the absence of defeaters, then x provides evidence.’

2.4 – Coda

NB: this view is *internalist* in the sense that the environment doesn’t have to be some way or other in order to obtain perceptual evidence: the mental state needn’t be factive or even the input to a reliable process. But the view is *externalist* in the sense that the thing that determines what evidence we have — the capacities employed — is not accessible to the agent.

So the account is consistent with *mentalist* internalism (assuming that the employment of a capacity is a type of mental state) but it’s not consistent with *access* internalism.

The upshot:

This externalist notion of phenomenal evidence makes room for the idea that having evidence is a matter of being in an epistemic position that is a guide to how the world is,⁸ while allowing that we can have evidence even if we happen to have been led astray and so are in a state that is not accurate with regard to our environment.⁹ As a consequence, the suggested capacity view shows how experience provides us with phenomenal evidence even in the bad case without retreating to introspective evidence.

⁸ presumably what Goldman is after when he demands a ‘systematic link’ between how things seem and how they are if the seemings are to justify beliefs.

⁹ this is what internalists want, as illustrated in NED scenarios.

Big-picture criticism (super-rough)

Way back on p. 706-7, SS described the puzzle motivating her project as this:

If our conscious mental states can be the very same regardless of our environment and if these conscious mental states determine our perceptual evidence, then our evidence will be the very same in the good and the bad case – that is, our evidence will be the very same regardless of whether we are accurately perceiving or suffering a hallucination. But if perceptual evidence is the very same in the good and the bad case, then it is mysterious why it would be rational to heed the testimony of our senses (see Goldman 1999 for this line of criticism). It is plausible that the reason for why it is valuable to take how our environment seems to us at face value is because doing so constitutes a useful way of pursuing an accurate view of the world. Evidence can play that role, however, only if there is a systematic link between our

sensory seemings and the way our environment actually is. In so far as evidential internalists do not account for such a link, they fail to account for the role of evidence as being a guide to how the world is. (706-7)

It's true that Goldman and other externalists think that it's only rational to take experience at face value if there's a systematic link between how things seem and how things are.

The whole point of SS's proposal re: phenomenal evidence is to provide such a systematic link and thereby to make plausible the claim that sensory seemings provide evidence for what the world is like.

For SS, then link comes in the form of a this metaphysical and explanatory dependence of the employment of perceptual capacities in the bad case upon what those capacities function to single out when employed in the good case.

This is a kind of systematic link, so in a sense SS seems to have met Goldman's challenge. But there's reason to worry that it's not the kind of systematic link that Goldman wants.

Arguably, the 'systematic link' that Goldman really wants is a positive correlation between how things seem and how they are. This is particularly clear for process reliabilists: if the reliability of a process is simply a matter of its outputs being true more often than not (70% of the time? 90%?), and if believing the content of a seeming is a process, then seemings provide evidence iff there's a positive correlation between how things seem and how they are.

One way to ensure a stronger positive correlation between how things seem and how they are is to restrict the class of seemings to *factive* seemings, as TW does. In that case the correlation very strong: if it seems_{factive} that p then p . This approach is equivalent to a kind of hyper-reliabilism: a process is hyper-reliable only if its outputs are true 100% of the time (as in the case with the process of forming beliefs on the basis of *factive* seemings). Forming beliefs on the basis of hyper-reliable evidence is a hyper-reliable process: it cannot lead you astray.¹⁰

We can agree that SS's view establishes *some sort* of systematic link between how things seem and how they are. But still: is it the sort of link that Goldman wants? Does it ensure a positive correlation between how things seem and how they are?

This will really depend on some of the details that SS waffles upon.

Example: if the function of one's capacity to single out particulars is determined by an evolutionary process, then we can reason as follows:

¹⁰ I don't mean to suggest that this is TW's actual position – the point is simply that views of this sort imply a positive correlation between evidence for p and the truth of p , and the existence of such a positive correlation is plausibly what motivates the view.

1. evolution selects for traits that facilitate surviving long enough to pass on your genes
2. the capacity to single out food and predators and potential mates would facilitate surviving long enough to pass on your genes only if those things really are in the environment
3. so the mere fact that I have these capacities means that tokens of these types of particulars are (probably) present in my environment
4. so I can rule out all of those possible worlds in which tokens of those types *don't* exist
5. All those world that I just ruled out are ones in which my seemings fail to track the truth: they're bad-case worlds. Assuming that logical space is roughly 50% good-case worlds and 50% bad-case worlds, when I rule out a bunch of good-case worlds the odds shift in the favor of good-case worlds, i.e. a positive correlation between how things seem and how they are is established.

Instead of offering an evolutionary explanation for how capacities get their functions we could assert a necessary condition for having a capacity to single out X: that someone, somewhere must have experienced an X. Assuming the further requirement that we be in some sort of causal contact with that someone, the somewhere must be in *this* possible world, meaning that at some point X's must be in *this* possible world. This again rules out a bunch of bad-case worlds, establishing the positive correlation between how things seem and how they are (plausibly).

So it's at least plausible that SS could achieve the sort of systematic link that Goldman is after by helping herself to these add-ons to her theory. But here's the point: in that case it looks like it's the add-ons doing all the work of establishing the systematic link rather than SS's capacities-based approach.

Example: Putnam¹¹ argued from content-externalism to the claim that experiences are positively correlated with truth (sort of – I'm taking some liberties here). I'm able to think *the water is wet* in part because I'm able to have thoughts about water, and a necessary condition for my having thoughts about water is that at some point I've had some causal contact with water. So I'm in a world in which I can come into contact with water, so I'm not a BIV. But if I can rule out BIV scenarios then I've ruled out lots of bad-case worlds, so I've got my positive correlation.¹²

¹¹ Chapter 1 of *Reasons, Truth, and History*

¹² OK, that required me to take a *lot* of liberties with the original Putnam, but you get the point - SS's account is inessential to the overall argument, as I'm just using one of the add-ons that SS could embrace.

3 – *The factive evidence argument*

SS has argued that Percy the perceiver and Hallie the hallucinator each have the same phenomenal evidence that there's a white cup on the table. She also thinks that overall, Percy has more evidence than Hallie because *accurate* perceptions yield phenomenal evidence *and factive evidence*. Argument:

- (4) If a subject S accurately perceives her environment, then S accurately represents her environment on the basis of her environment.
- (5) If S accurately represents her environment on the basis of her environment, then S has factive evidence determined by her environment.
- (6) So, if S accurately perceives her environment, then S has factive evidence determined by her environment.

3.1 – *Premise (4): Perceptual content*

The purpose of this section is to defend (4):

- (4) If a subject S accurately perceives her environment, then S accurately represents her environment on the basis of her environment.

(4) follows from two theses:

First: perception is representational

The sense of representation here is very weak: it doesn't require that representations be propositional, or conceptual, or that it involves propositional attitudes between agents and contents. 'In so far as employing perceptual capacities determines sensory states and in so far as sensory states are correlated with perceptual content, we can say that employing perceptual capacities yields perceptual content.'
(724)

Second: 'perceivers accurately represent their environment on the basis of their environment'.

SS doesn't really defend this thesis,¹³ but she does explain that 'on the basis of' is intended to be equivalent to 'is caused in a non-deviant way by'.

¹³ perhaps she means only to spell out what she means by 'perceive'?

3.2 – *Premise 5: Perceptual content and the factivity of perception*

The purpose of this section is to defend (5):

(5) If S accurately represents her environment on the basis of her environment, then S has factive evidence determined by her environment.

(5) follows from two theses:

- (i) we have evidence if we accurately represent our environment on the basis of our environment
- (ii) accurate representations of the environment on the basis of the environment yields factive evidence

Re: (i)

(i) is consistent with lots of ways of understanding what evidence is. SS doesn't really defend this thesis, she thinks it's something that will follow from 'any reasonable conception of what having evidence requires'.

Re: (ii)

SS motivates (ii) in two ways.

Simple way: 'perception is factive and it is reasonable that the evidence provided by perception inherits the factivity of perception.'
(726)

The not so simple way depends upon SS's theory of perceptual content. Overview of the theory:

- there are two levels of perceptual content: representational content *types* and relational content *tokens*.
- type-level contents are determined by the perceptual capacities employed. Any set of subjectively indistinguishable experiences share a content type, so an experience and a subjectively indistinguishable hallucination share the same content. So token-contents are not factive.
- SS has already argued that employing perceptual faculties results in phenomenal experiences that themselves provide non-factive evidence
- token contents are partially determined by the environment: I'm having an experience as of Pam because Pam is in my environment; if it had be her identical sister Sara in the environment instead then I'd be having an experience as of Sara. Importantly, those two experiences might be the result of employing the very

same capacities, and hence they might be subjectively indistinguishable, but because those capacities are employed in different environments they single out different particulars. Hence although they might have the same type-level content, they are of different things, so they have different token-level content.

- SS thinks of token-contents as proposition-like.¹⁴ If I have a phenomenal experience as of Pam’s tallness, the token-content of my experience is composed of a particular object, Pam, and an instance of the property of being tall. Importantly, both of those components are essential in order for this token-content to have truth conditions; to represent the world as being some way or other. If I’m hallucinating and Pam isn’t really there in my environment then the token content is missing, so the content is gappy. If Pam really is there in the environment but she isn’t really tall — that part is an illusion — then there’s a gap in my token-content where the property instance ought to be.
- Hence hallucination and illusion both produce only gappy-contents, which fail to represent the environment as being some way or other. Only non-gappy contents represent the world as being some way or other, and only veridical experiences produce non-gappy token contents, so token contents are factive. That means that they’re strongly correlated with the environment,¹⁵ and hence they provide factive evidence.

¹⁴ at least the way that she describes them is suggestive of proposition-like token contents

¹⁵ perfectly correlated, at least in the content to environment direction

SS doesn’t want her theory of evidence to depend upon the details of her theory of content, so she provides a more general account. Here the assumption is only that the *accuracy conditions* of experience are sensitive to the particulars in the environment:

If the content of experience lays down a condition under which it is accurate in a way that is sensitive to which particulars (if any) are perceived, then the way the experiencing subject’s environment is will make a difference to the content of her experience... [W]hich particular object in a subject’s environment is represented [is] determine[d by] whether the subject’s environment is as it is represented to be. If this is right, then for an experience with the content ‘that coffee cup is white’ to be accurate it is not sufficient that ‘that’ refers to some coffee cup instantiating the right properties. It is necessary that ‘that’ refer to the particular object perceived. If the content of experience lays down the conditions under which the experience is accurate and the accuracy of an experience depends on the environment, then the particulars to which the subject is perceptually related will make a constitutive difference to the token content of her experience. (731-2)

BTM: this picture is much more plausible when we take as particulars the *objects* that are in our environment and that inform our

token-contents. But on this we're supposed to treat their properties the same way: we're supposed to treat them as property=*instances*. That's much less plausible. When I see a white cup, and then I have a subjectively identical experience but there's a different white cup in the environment, then it's plausible to claim that my experience has a different content. But what happens when I see the white cup at one moment, but then I see the very same cup that's exactly the same shade of white, but this time it's a distinct instance of that shade of white? Are the contents of those experiences really distinct? The answer isn't clear to me, but I find myself grimacing involuntarily as I read this bit, which is clear evidence that something is amiss...

If my skepticism is borne out, then SS's account of what's going on in cases of illusion fails: illusions don't result in *gappy* token-contents, they result in *inaccurate* token-contents. The whole reason that token-contents are supposed to produce factive evidence is that experiences with those contents are *factive* — it's impossible to have an experience with a given token-content unless the accuracy conditions of that content are satisfied — but if I'm right then they aren't factive after all. So the mischief that this observation makes (if it's borne out) cuts deeply into SS's account. [End BTM]

4 – *The common rational source of phenomenal and factive evidence*

The point of this section is to sum up and to discuss how the capacities view relates to alternatives.

Three Strengths of the Capacities View:

1. it is an externalist view of evidence that makes room for hallucinations providing us with evidence without retreating to introspective evidence, a general content, or an existentially quantified content.
2. it implies that we can have perceptual evidence only if we are in a sensory state, i.e. only if we're having a phenomenal experience: since we have factive perceptual evidence in virtue of employing perceptual capacities and since employing such capacities yields a sensory state, we cannot have factive perceptual evidence without being in a sensory state.
3. provides for a way of explaining why it is that a perceiver is in a better epistemic position than a hallucinator.¹⁶

¹⁶ is that really something that needs to be explained? Is it true?

Capacities view and Williamson

Williamson offers an alternative way to understand (3). TW thinks all evidence is factive, so contra SS there is no phenomenal evidence. But Hallie the Hallucinator still has evidence through introspection on her own experiences: she doesn't have Percy's factive perceptual evidence *there's a cup*, but she does have factive introspective introspective evidence *it appears that there's a cup*.

Three criticisms:

- (i) seems false that Hallie doesn't have any *perceptual* evidence
- (ii) TW's view over-intellectualizes evidence-acquisition in hallucination: I can't know *it appears that there's a cup* unless I can conceptualize appearances, but even cognitively simple creatures who can't do that obtain evidence from hallucinations.¹⁷
- (iii) on TW's view, the sources of evidence is different in the good and the bad case, which is weird

¹⁷ recall that TW worries about this too, considers whether the introspectively known propositions are demonstrative in nature: *that's a white cup*