

Mentalism and Epistemic Transparency

ABSTRACT: Epistemic transparency is central to the debate between factive and non-factive versions of mentalism about evidence. If evidence is transparent, then factive mentalism is false, since factive mental states are not transparent. However, Timothy Williamson has argued that epistemic transparency is a myth, since there are no transparent conditions except trivial ones. This paper responds by drawing a distinction between doxastic and epistemic notions of transparency. Williamson's argument succeeds in showing that no conditions are doxastically transparent, but it fails to show that no conditions are epistemically transparent. Moreover, this is sufficient to reinstate the argument against factive mentalism.

1. Knowledge-First Epistemology

In *Knowledge and Its Limits*, Timothy Williamson develops a distinctive approach to epistemology, which he sums up in the slogan: 'knowledge first'. Instead of explaining knowledge in terms of justification and other epistemic notions, Williamson explains justification in terms of knowledge and thereby inverts the traditional order of explanation. A central plank of Williamson's knowledge-first epistemology is his claim that knowledge is a mental state.

Mentalism about evidence is the thesis that one's evidence is determined by one's mental states.¹ Traditionally, proponents of mentalism have supposed that one's evidence is determined by one's *non-factive* mental states. However, if knowledge is a factive mental state, then there is logical space for a factive version of mentalism on which one's evidence is determined by one's *factive* mental states, rather than one's non-factive mental states.

This position in logical space is occupied by Williamson's epistemology. Williamson (2000: Ch.1) argues that knowledge is a factive mental state; indeed, it is the most general factive mental state in the sense that all factive mental states are determinate ways of knowing. Moreover, Williamson (2000: Ch.9) argues that one's knowledge determines one's evidence, since one's total evidence just is the total content of one's knowledge. This entails a factive version of mentalism on which one's evidence is determined by one's factive mental states.

An influential source of resistance to Williamson's epistemology stems from what he calls 'the myth of epistemic transparency'. On this view, one's evidence is transparent in the sense that one is always in a position to know which propositions are included in one's total evidence. However, if one's evidence is transparent, then it cannot be determined by one's knowledge, since one's knowledge is not transparent in the sense that one is always in a position to know which propositions one knows. More generally, if mentalism is true, then one's evidence is transparent only if the mental states that determine one's evidence are themselves transparent.² Thus, epistemic transparency provides the basis of the following line of argument against factive mentalism:

- (1) Evidence is transparent
- (2) Evidence is transparent only if the mental states that determine evidence are themselves transparent
- (3) Factive mental states are not transparent
- (4) Therefore, evidence is not determined by factive mental states

Williamson's response is to argue that epistemic transparency is a myth – a quaint relic of Cartesian epistemology. He argues that only trivial conditions are transparent in the sense that one is always in a position to know whether or not they obtain. If non-factive mental states are no more transparent than factive mental states, then this undermines one of the central motivations for rejecting factive versions of mentalism in favour of non-factive versions.³ Therefore, Williamson's rejection of the myth of epistemic transparency plays a central role in motivating his distinctive brand of knowledge-first epistemology.

2. The Anti-Luminosity Argument

A condition is *transparent* if and only if one is always in a position to know whether or not it obtains. A condition is *luminous* if and only if one is always in a position to know that it obtains when it does. So, a condition is transparent if and only if it is *strongly luminous* – that is, one is always in a position to know that it obtains when it does and that it does not obtain when it does not. If there are no luminous conditions, then there are no transparent conditions either.

Williamson's anti-luminosity argument is designed to establish that there are no luminous conditions except trivial ones, which hold in all cases or none.⁴ The argument exploits a tension between the assumption that there are some luminous

conditions and the assumption that knowledge requires a margin for error. These assumptions jointly entail a tolerance principle, which is falsified by any sorites series of pairwise close cases that begins with a case in which C obtains and ends with a case in which C does not obtain. So, the following assumptions yield a contradiction:

- (1) **Luminosity:** C is luminous, so if C obtains, then one is in a position to know that C obtains
- (2) **Margins:** If one is in a position to know that C obtains, then C obtains in every close case
- (3) **Tolerance:** If C obtains, then C obtains in every close case (from 1, 2)
- (4) **Gradual Change:** There is a series of close cases that begins with a case in which C obtains and ends with a case in which C does not obtain

To illustrate the problem, Williamson asks us to consider a morning on which one feels cold at dawn and then gradually warms up until one feels warm at noon. The process is so gradual that one cannot discriminate any change in one's condition from one moment to the next. By hypothesis, one feels cold at dawn. By the definition of luminosity, if feeling cold is a luminous condition, then one is in a position to know that one feels cold at dawn. By the margin for error principle, it follows that one feels cold a moment later, since this is a relevantly close case. By repeating these moves, we generate the conclusion that one feels cold at noon. However, this contradicts the initial stipulation that one feels warm, rather than cold, at noon.

Williamson urges that we should resolve the contradiction by denying that there are any luminous conditions. But why not deny instead that knowledge requires margins for error? The answer is that the margin for error principle is motivated by independently plausible assumptions – in particular, that knowledge requires safety from error and that there are limits on our powers of discrimination.

According to the safety requirement, one knows that a condition C obtains only if one does not falsely believe that C obtains in any close case. The rationale is that if one's beliefs are not safe from error, then they are not sufficiently reliable to count as knowledge. Still, there is a gap to be bridged between safety and margins. I can know that C obtains, even if C does not obtain in every close case, so long as there is no close case in which I falsely believe that C obtains. This need not impugn my knowledge so long as my powers of discrimination are sufficiently sensitive to the difference between cases in which C obtains and cases in which C does not obtain.⁵

The margin for error principle does not hold without exception, but only given the further assumption that we cannot discriminate between close cases. However, it is question-begging to assume that one is not in a position to know the conditions that make the difference between close cases. A more neutral assumption is that one's doxastic dispositions are less than perfectly sensitive to the difference between close cases. However, we cannot assume a tolerance principle on which one believes that C obtains only if one believes that C obtains in every close case. In Williamson's example, one's degree of confidence that one feels cold may gradually decrease until one falls below the threshold for believing that one feels cold. Still, if one's powers of discrimination are limited, then one's degree of confidence cannot differ too radically between close cases. This is what we need in order to derive margins from safety.

Let us focus our attention to Williamson's example in which one gradually warms up between dawn and noon. First, we may assume that throughout the process, one does everything that one is in a position to do, so if one is in a position to know that C obtains, then one knows that C obtains. Second, we may assume that one's powers of discrimination are limited, so one's degree of confidence that C obtains cannot differ too radically between close cases. Third, we may assume that knowledge requires safety from error and so one knows that C obtains only if C obtains in every close case in which one has a similarly high degree of confidence that C obtains. Given these assumptions, we may conclude that the margin for error principle is true, if not in general, then at least in this specific example:

- (1) **Position:** If one is in a position to know that C obtains, then one knows that C obtains
- (2) **Discrimination:** If one knows that C obtains, then one has a high degree of confidence that C obtains in every close case
- (3) **Safety:** If one knows that C obtains, then C obtains in every close case in which one has a high degree of confidence that C obtains
- (4) **Margins:** If one is in a position to know that C obtains, then C obtains in every close case

But if the margin for error principle is true, then there are no luminous conditions.

Broadly speaking, responses to Williamson's anti-luminosity argument can be divided into two categories: *offensive* and *defensive*. Offensive responses reject the conclusion of the argument and so take on the burden of rejecting at least one of its

premises.⁶ By contrast, defensive responses accept the conclusion of the argument, but engage in a kind of damage limitation exercise.⁷

My strategy in this paper is defensive. I will concede that Williamson's argument establishes that there are no luminous conditions. However, I will attempt to limit the damage by relocating the epistemic asymmetry between factive and non-factive mental states. Recall that what is at stake in this debate is the motivation for rejecting factive mentalism in favour of non-factive mentalism. The defensive strategy aims to show that even if there are no luminous conditions, there is nevertheless an epistemic asymmetry between factive and non-mental states. As long as there is some epistemic criterion of *quasi-transparency* that is satisfied by non-factive mental states, but not by factive mental states, we can reinstate the original form of argument against factive mentalism.

Can Williamson's anti-luminosity argument be generalized in such a way as to establish that there is no relevant epistemic asymmetry between factive and non-factive mental states? As we have seen, Williamson's knowledge-first epistemology is motivated by the claim that there is no such epistemic asymmetry. Thus, he writes: "Any genuine requirement of privileged access on mental states is met by the state of knowing *p*. Knowing is characteristically open to first-person present-tense access; like other mental states, it is not perfectly open." (2000: 25)

3. The Lustrous and the Luminous

One defensive strategy is proposed by Selim Berker (2008), who suggests that even if there are no luminous conditions, there may be some lustrous conditions. A condition is *luminous* if and only if one is always in a position to know that it obtains if it does. By contrast, a condition is *lustrous* if and only if one is always in a position to justifiably believe that it obtains if it does. So, a condition is lustrous, but not luminous, if one is always in a position to believe justifiably, if not knowledgeably, that it obtains if it does.

Can Williamson's anti-luminosity argument be extended to show that there are no lustrous conditions? Consider the following:

- (1) **Lustrousness:** C is lustrous, so if C obtains, then one is in a position to justifiably believe that C obtains
- (2) **Margins:** If one is in a position to justifiably believe that C obtains, then C obtains in every close case

(3) **Tolerance:** If C obtains, then C obtains in every close case (from 1, 2)

(4) **Gradual Change:** There is a series of close cases that begins with a case in which C obtains and ends with a case in which C does not obtain

As before, the argument relies upon a margin for error principle. However, it is more plausible that margins for error are required in the case of knowledge than justified belief. If justified belief is non-factive, then the margin for error principle is false, since one may be in a position to justifiably believe that C obtains even if C does not obtain in the actual case, which is the closest of all possible cases. As Williamson himself remarks, only factive conditions are subject to margin for error principles.⁸

This point has limited value for proponents of the defensive strategy. After all, the aim is to identify an epistemic criterion that is satisfied by non-factive mental states, but not factive mental states. Arguably, however, factive mental states satisfy the criterion of lustrousness. For instance, if one sees that *p*, then one is in a position to justifiably believe that one sees that *p*.⁹ Certainly, no factive mental state is *strongly lustrous* in the sense that one is always in a position to justifiably believe that it obtains when it does and that it does not obtain when it does not. For instance, it is not the case that if one does not see that *p*, but merely seems to see that *p*, then one is in a position to justifiably believe that one does not see that *p*, but merely seems to see that *p*. Therefore, proponents of the defensive strategy must argue that some non-factive mental states are strongly lustrous in order to motivate an epistemic asymmetry.

The problem is that if C is strongly lustrous, then justified belief that C obtains is factive. The argument is straightforward. If justified belief that C obtains is non-factive, then there are cases in which one justifiably believes that C obtains and yet C does not obtain. But if C is strongly lustrous and C does not obtain, then one is in a position to justifiably believe that C does not obtain. It follows that one is in a position to justifiably believe that C obtains and that C does not obtain. And yet one is never in a position to justifiably believe a contradiction. Therefore, if C is strongly lustrous, then one is in a position to justifiably believe that C obtains *if and only if* C obtains.

If Margins is restricted to strongly lustrous conditions, then the objection from factivity is blocked. Still, further argument is needed in order to establish that the restricted version of Margins is true. After all, justified *true* belief is factive, but it does not generally require a margin for error. Gettier cases typically involve a subject who justifiably believes that a condition C obtains, which does obtain in the actual case, but not in all or even most of the closest non-actual cases. For example, in Alvin

Goldman's (1976) fake barn case, Henry has a justified true belief that there is a barn on the road ahead, but he doesn't know this, since his belief is false in most of the closest non-actual cases. One possible reaction to Williamson's anti-luminosity argument is that it uncovers a new kind of Gettier case in which being close to the margin for error is like being in fake barn country – that is, one is in a position to form a justified true belief, which is not sufficiently reliable to count as knowledge.¹⁰

In what follows, I provide an argument for the restricted version of Margins. The argument is motivated by reflection on Ernest Sosa's (2003) version of the problem of the speckled hen. Following Sosa, we can ask: what explains the difference between the following pair of cases?

- (i) If I experience 3 speckles, then I am in a position to justifiably believe that I experience 3 speckles
- (ii) If I experience 48 speckles, then I am not in a position to justifiably believe that I experience 48 speckles

Why am I in a position to form a justified belief in the one case, but not the other?

One strategy for solving the problem appeals to facts about the determinacy of the representational content of experience. My experience in the first case represents that the hen has 3 speckles, whereas my experience in the second case does not represent that the hen has 48 speckles: it is simply not that determinate. Even if I do experience 48 speckles, it is a further question whether my experience represents that there are exactly 48 speckles; indeed, there may be no determinate number n such that my experience represents that there are exactly n speckles. If so, then my experience in the first case provides justification to believe that I experience 3 speckles, but my experience in the second case does not provide justification to believe that I experience 48 speckles. In general, my experience provides justification to believe that I experience n speckles if and only if my experience represents n speckles.

The problem with this response is that it fails to generalize to other examples. All we need to generate the problem is a case in which the determinacy of experience is more fine-grained than one's powers of discrimination in judgement. For instance, experience represents objects as having not just determinable shades, such as red, but also more determinate shades, such as red-48 and red-49. Nevertheless, I may be unable to discriminate these shades in judgement: perhaps I can tell them apart when presented simultaneously, but not when presented in sequence.¹¹ In that case, my experience might represent a shade as red-48, although I am no better than chance in

judging whether I am experiencing red-48 or red-49. Once again, we can ask: what explains the difference between the following pair of cases?

- (i) If I experience red, then I am in a position to justifiably believe that I experience red
- (ii) If I experience red-48, then I am not in a position to justifiably believe that I experience red-48

Why am I in a position to form a justifiable belief in the one case, but not the other?

Sosa's solution appeals to *safety*.¹² My belief that I experience red is safe from error, since I would not believe that I experience red unless I were to experience red. By contrast, my belief that I experience red-48 is not safe from error, since I could easily believe that I experience red-48 if I were to experience red-47 or red-49. Likewise, in Sosa's original example, my belief that I experience 3 speckles is safe from error, since I would not believe that I experience 3 speckles unless I were to experience 3 speckles. By contrast, my belief that I experience 48 speckles is not safe from error, since I could easily believe that I experience 48 speckles if I were to experience 47 or 49 speckles.

The problem with Sosa's solution is that it does not generalize to beliefs about the external world. However, the problem of the speckled hen arises for beliefs about the external world as well as the internal world. Consider the following pair of cases:

- (i) If I experience red, then I am in a position to justifiably believe that there is something red
- (ii) If I experience red-48, then I am not in a position to justifiably believe that there is something red-48

Why am I in a position to form a justifiable belief in the one case, but not the other?

In this context, safety is a red herring. If I am hallucinating a red object and I believe that there is something red, then my belief is false. Still, there is an intuitive difference between the epistemic status of my belief that there is something red and my belief that there is something red-48. And yet safety cannot explain the difference, since no false belief is safe from error. What we need in order to explain the intuitive difference is not safety from *error*, but rather safety from *lack of justification*.

Absent defeaters, my experience provides justification to believe that *p* if and only if it represents that *p*. If I believe that *p*, however, my belief is justified only if it is based in a way that is counterfactually sensitive to the representational content of my experience, which provides my justification to believe that *p*. As Sosa rightly

insists, an actual match in content is not sufficient. Given the limits on my powers of discrimination, this condition is satisfied in the first case, but not the second. My belief that there is something red is counterfactually sensitive to the content of my experience, since I would not easily believe that there is something red unless my experience represents that there is something red. By contrast, my belief that there is something red-48 is not counterfactually sensitive to the content of my experience, since I could easily believe that there is something red-48 when in fact my experience represents that there is something red-47 or red-49.

What we need to solve the problem of the speckled hen is not a *safety* requirement of counterfactual sensitivity to the facts, but rather a *basing* requirement of counterfactual sensitivity to one's evidence, which determines which propositions one has justification to believe. One justifiably believes that C obtains only if one's justifying evidence obtains in every close case in which one has a similarly high degree of confidence that C obtains. If one believes that C obtains on the basis of justifying evidence E, but there is a close case in which one's justifying evidence E does not obtain and yet one believes or has a similarly high degree of confidence that C obtains, then one's belief is unjustified.¹³

What is required for a belief to be justified is not counterfactual sensitivity to the facts, but rather counterfactual sensitivity to one's justifying evidence. In the special case of strongly lustrous conditions, however, there is no distinction between one's justifying evidence and the facts. If C is strongly lustrous, then one is in a position to form a justified belief that C obtains if and only if C obtains and moreover *because* C obtains. In that case, the condition C that justifies one's belief is one and the same as the fact that one's belief is about. Since there is no distinction in this case between one's justifying evidence and the facts, we can derive a local safety condition, which requires counterfactual sensitivity to the facts, from a more general basing condition, which requires counterfactual sensitivity to one's justifying evidence. Thus, for any strongly lustrous condition C, we can argue as follows:

- (1) **Basing:** If one justifiably believes that C obtains, then one's justifying evidence E obtains in every close case in which one has a high degree of confidence that C obtains
- (2) **Identity:** If one justifiably believes that C obtains, then C is identical to one's justifying evidence E

- (3) **Safety:** If one justifiably believes that C obtains, then C obtains in every close case in which one has a high degree of confidence that C obtains

Next we can derive Margins from Safety by the following argument:

- (1) **Position:** If one is in a position to justifiably believe that C obtains, then one justifiably believes that C obtains
- (2) **Discrimination:** If one justifiably believes that C obtains, then one has a high degree of confidence that C obtains in every close case
- (3) **Safety:** If one justifiably believes that C obtains, then C obtains in every close case in which one has a high degree of confidence that C obtains
- (4) **Margins:** If one is in a position to justifiably believe that C obtains, then C obtains in every close case

And finally we can exploit Margins in arguing that there are no strongly lustrous conditions.

4. Epistemic and Doxastic Lustrousness

The situation so far has gone from bad to worse. Williamson's argument can be extended to show not only that there are no luminous conditions, but also that there are no strongly lustrous conditions. However, all is not yet lost. In what follows, I will draw a distinction between epistemic and doxastic interpretations of lustrousness and luminosity. Moreover, I will argue that even if Williamson's argument shows that there are no doxastically luminous or lustrous conditions, it fails to show that there are no epistemically luminous or lustrous conditions. Therefore, Williamson's argument fails in its attempt to debunk the myth of epistemic transparency.

A condition is lustrous if and only if one is always in a position to justifiably believe that it obtains when it does. But what exactly does it mean to say that one is in a position to justifiably believe that a condition obtains? We can draw a distinction between epistemic and doxastic interpretations. One is in an *epistemic position* to justifiably believe that C obtains if and only if one has justification to believe that C obtains. By contrast, one is in a *doxastic position* to justifiably believe that C obtains if and only if one has the capacity to use one's justification in forming a justified belief that C obtains. There is a corresponding distinction between epistemic and doxastic interpretations of lustrousness. A condition is *epistemically lustrous* if and only if one is always in an epistemic position to justifiably believe that it obtains if it does. By contrast, a condition is *doxastically lustrous* if and only if one is always in a

doxastic position to justifiably believe that it obtains when it does. My claim is that Williamson's argument can be extended to show that no conditions are doxastically lustrous, but it fails to show that no conditions are epistemically lustrous.

The key point is that being in a doxastic position to justifiably believe a proposition is more demanding than being in an epistemic position to justifiably believe it, since it requires not only having justification to believe it, but also having the capacity to *use* one's justification in forming a justified belief. This leaves open the possibility that one has justification to believe a proposition, although one lacks the doxastic capacities required to use it in forming a justified belief.

There are various different ways in which this possibility might be realized. First, one might have justification to believe a proposition, but lack the concepts required to believe it. Second, one might have the concepts, but lack the compositional capacities required to put them together in the right way. Third, one might have all the concepts and compositional abilities required to believe the proposition in question, but lack the powers of discrimination required to *justifiably* believe it. This third possibility is best illustrated by the problem of the speckled hen.

My experience provides me with justification to believe that it represents red-48 if and only if it does in fact represent red-48. However, it is a further question whether or not I have the doxastic capacity to use my justification in forming a justified belief that my experience represents red-48. If I form the belief, then my belief is justified only if it is based in a way that is counterfactually sensitive to the representational content of my experience, which provides my justification to believe the proposition in question. However, if my powers of discrimination are limited, then I may be unable to satisfy this basing requirement, since I could easily believe that my experience represents red-48 when in fact it represents red-47 or red-49. In that case, I am unable to use my justification to believe that my experience represents red-48 because my doxastic dispositions are not counterfactually sensitive to the evidence, which determines my justification to believe the proposition in question.

The moral to be drawn from the problem of the speckled hen is that limits on one's powers of discrimination do not constrain which propositions one is in an epistemic position to justifiably believe, but rather which propositions one is in a doxastic position to justifiably believe. If my experience represents red-48, then I am in an epistemic position to justifiably believe that my experience represents red-48. However, if I cannot discriminate between a case in which my experience represents

red-48 and a close case in which my experience represents red-47 or red-49, then I am not in a doxastic position to justifiably believe that my experience represents red-48.

Margin for error principles must be interpreted in light of a distinction between epistemic and doxastic notions of closeness. Two cases are *epistemically close* only if one's evidence is the same or similar, so that one has justification to believe the same propositions to the same or a similar degree. By contrast, two cases are *doxastically close* if one's evidence is treated as the same or similar in the sense that one's justified degree of confidence in each case is the same or similar. If one's powers of discrimination are limited, then two cases may be close in the doxastic sense, but not the epistemic sense, since one's doxastic dispositions are insensitive to the evidential difference between them. This possibility is illustrated by the problem of the speckled hen.

Margin for error principles can be given either an epistemic or a doxastic interpretation. According to the doxastic interpretation, if one is in a doxastic position to justifiably believe that C obtains, then C obtains in every doxastically close case. According to the epistemic interpretation, by contrast, if one is in an epistemic position to justifiably believe that C obtains, then C obtains in every epistemically close case. Both epistemic and doxastic interpretations are plausible, but neither can be used in arguing that there are no epistemically lustrous conditions. To see this, we need to consider the extended version of Williamson's argument under both doxastic and epistemic interpretations.

First, consider the doxastic interpretation:

- (1) **Lustrousness:** C is doxastically lustrous, so if C obtains, then one is in a doxastic position to justifiably believe that C obtains
- (2) **Margins:** If one is in a doxastic position to justifiably believe that C obtains, then C obtains in every doxastically close case
- (3) **Tolerance:** If C obtains, then C obtains in every doxastically close case
- (4) **Gradual Change:** There is a series of doxastically close cases that begins with a case in which C obtains and ends with a case in which C does not obtain

This argument shows that there are no doxastically lustrous conditions, but it is entirely neutral on whether or not there are any epistemically lustrous conditions.

Next, consider the epistemic interpretation:

- (1) **Lustrousness:** C is epistemically lustrous, so if C obtains, then one is in an epistemic position to justifiably believe that C obtains
- (2) **Margins:** If one is in an epistemic position to justifiably believe that C obtains, then C obtains in every epistemically close case
- (3) **Tolerance:** If C obtains, then C obtains in every epistemically close case
- (4) **Gradual Change:** There is a series of epistemically close cases that begins with a case in which C obtains and ends with a case in which C does not obtain

This argument fails to show that there are no epistemically lustrous conditions, since the epistemic interpretation of Gradual Change is false. If C is epistemically lustrous, then there is no series of epistemically close cases that begins with a case in which C obtains and ends with a case in which C does not obtain. In Williamson's example, subsequent times in the gradual process of warming up between dawn and noon are doxastically close in the sense that one's powers of discrimination are limited, so one treats one's evidence as the same or very similar. However, there is no argument that subsequent times are epistemically close in the sense that one's evidence is in fact the same or very similar. Following Williamson, we may assume that at some point in the series there is a sharp cut-off such that one feels cold in one case, but one does not feel cold in the next case. One's evidence may differ radically between such cases even if one's powers of discrimination are limited and so one's doxastic dispositions are insensitive to the evidential difference between them.

If we consistently assume either an epistemic or a doxastic interpretation, then the argument fails to show that there are no epistemically lustrous conditions. But what if we mix and match? Consider the following mixed interpretation:

- (1) **Lustrousness:** C is epistemically lustrous, so if C obtains, then one is in an epistemic position to justifiably believe that C obtains
- (2) **Margins:** If one is in an epistemic position to justifiably believe that C obtains, then C obtains in every doxastically close case
- (3) **Tolerance:** If C obtains, then C obtains in every doxastically close case
- (4) **Gradual Change:** There is a series of doxastically close cases that begins with a case in which C obtains and ends with a case in which C does not obtain

This argument fails to show that there are no epistemically lustrous conditions, since the mixed interpretation of Margins is false. As we have seen, the moral to be drawn

from the problem of the speckled hen is that limits on one's powers of discrimination do not constrain which propositions one is in an epistemic position to justifiably believe, but which propositions one is in a doxastic position to justifiably believe. Therefore, one may be in an epistemic position to justifiably believe that some condition obtains, although it does not obtain in every doxastically close case.

It is worth considering why the argument of the previous section cannot be extended to support the mixed interpretation of Margins. Consider the following:

- (1) **Position:** If one is in an epistemic position to justifiably believe that C obtains, then one justifiably believes that C obtains
- (2) **Discrimination:** If one justifiably believes that C obtains, then one has a high degree of confidence that C obtains in every doxastically close case
- (3) **Safety:** If one justifiably believes that C obtains, then C obtains in every doxastically close case in which one has a high degree of confidence that C obtains
- (4) **Margins:** If one is in an epistemic position to justifiably believe that C obtains, then C obtains in every doxastically close case

The problem with this argument is that Position on its epistemic interpretation is in tension with Discrimination. Position is motivated by the stipulation that if one is in an epistemic position to justifiably believe a proposition, and one does what one is in an epistemic position to do, then one justifiably believes the proposition in question. Therefore, one is entitled to assume Position only if one is also entitled to assume that one always does what one is in an epistemic position to do. However, this assumption conflicts with Discrimination. If one's powers of discrimination are limited, then one is not always capable of taking advantage of one's epistemic position. That was the lesson to be learned from the problem of the speckled hen.

The conclusion so far is that Williamson's argument shows that no conditions are doxastically lustrous, but it fails to show that no conditions are epistemically lustrous. Moreover, this conclusion seems to generalize, since we can draw a related distinction between epistemic and doxastic interpretations of luminosity. A condition is *epistemically luminous* if and only if one is always in an epistemic position to know that it obtains when it does. By contrast, a condition is *doxastically luminous* if and only if one is always in a doxastic position to know that it obtains when it does.

One is in an *epistemic position* to know a proposition if and only if one satisfies all the purely epistemic conditions for knowledge, whereas one is in a

doxastic position to know a proposition if and only if one has the doxastic capacity to take advantage of one's epistemic position by satisfying doxastic as well as epistemic conditions for knowledge. Thus, being in a doxastic position to know a proposition is more demanding than being in an epistemic position to know it, since it requires not just being in an epistemic position, but also having the doxastic capacity to take advantage of it. If one does not always have the doxastic capacity to take advantage of one's epistemic position, then some conditions may be epistemically luminous even if no conditions are doxastically luminous. Williamson's argument shows that there are no doxastically luminous conditions, but for reasons analogous to those given above, it does not thereby show that there are no epistemically luminous conditions.

5. Epistemic Idealization

My defensive response to Williamson's argument is that some conditions may be epistemically lustrous or luminous even if no conditions are doxastically lustrous or luminous. The availability of this response depends on the assumption that one may be in an epistemic position but not a doxastic position to justifiably believe or to know that some condition obtains, since one lacks the doxastic capacity to take advantage of one's epistemic position. However, this assumption may be challenged. On some views, one's doxastic capacities impose corresponding limits on which propositions one is in an epistemic position to justifiably believe and know. But is there any good argument for the existence of such doxastic limits?

Some epistemologists define justification in terms of the capacity for justified belief: for instance, one has justification to believe a proposition if and only if one has a capacity for justified belief in that proposition.¹⁴ An immediate consequence of this definition is that limits on one's doxastic capacities impose corresponding limits on which propositions one has justification to believe. Arguably, however, justification is not defined in terms of the capacity for justified belief, but vice versa: thus, one has the capacity to justifiably believe a proposition if and only if one has the capacity to take advantage of one's justification to believe the proposition in question by believing it on the basis of one's justification. This definition does not imply the existence of doxastic limits on which propositions one has justification to believe.

In my view, justification can be defined in terms of idealized justified belief: roughly, one has justification to believe a proposition if and only if one has a justified belief in that proposition after appropriate idealization.¹⁵ However, this definition is

circular, since the relevant idealization is to be understood as a matter of justifiably believing all and only propositions that one has justification to believe. Moreover, there is no presumption that one has the doxastic capacities required to achieve this epistemic ideal. On this view, which propositions one has justification to believe corresponds to an epistemic ideal, which abstracts away from contingent facts about the limits on one's doxastic capacities. Therefore, the limits on one's doxastic capacities do not constrain the epistemic ideal, but only the extent to which one is capable of approximating towards the ideal.

An influential line of argument for the existence of doxastic limits appeals to a deontological conception of justification as a source of epistemic obligations, which are binding only insofar as one has the psychological capacities required to discharge them. The argument proceeds roughly as follows:

- (1) If one has justification to believe a proposition, then one ought to believe it on the basis of one's justification.
- (2) If one ought to believe a proposition on the basis of one's justification, then one can believe it on the basis of one's justification.
- (3) So, if one has justification to believe a proposition, then one can believe it on the basis of one's justification.

However, this version of the deontological conception of justification faces well known counterexamples, in which drugs, brainwashing or mental illness undermine one's doxastic capacity to take advantage of one's justification. In such cases, one's doxastic limitations excuse one from blame for one's epistemic failings. But while there are doxastic constraints on blameworthiness, there are no corresponding doxastic constraints on which propositions one has justification to believe.

What's the alternative? One option is to defend the deontological conception by arguing that 'ought' does not always imply 'can'. Feldman (2000) argues that there are so-called 'role oughts' that apply to whatever plays a certain role, regardless of how well it is capable of playing that role – for instance, chefs ought to make delicious food and jugglers ought to keep their balls in the air. Similarly, Feldman claims, there are epistemic 'oughts' that apply to us in virtue of our role as believers: "It is our plight to be believers. We ought to do it right. It doesn't matter that in some cases we are unable to do so." (2000: 676)

Another option, following Alston (1989), is to reject the deontological conception of justification in favour of an evaluative conception, according to which

justification is a source of epistemic values or ideals, rather than obligations. On this view, there are no doxastic limits on which propositions one has justification to believe because epistemic ideals – like ideals of morality, scientific understanding, and chess – may lie beyond our limited capacities. David Christensen puts the point well: “Not all evaluation need be circumscribed by the abilities of the evaluated. In epistemology, as in various other arenas, we need not grade on effort.” (2006: 162)

One might object that such an idealized conception of justification loses touch with our ordinary practices of epistemic evaluation. For instance, Frege’s belief in Axiom V was justified by ordinary standards, but not by ideal standards. Certainly, there is a distinction here, but there is also an explanatory connection. Being justified by ordinary standards is a matter of meeting some contextually determined threshold on a scale that is defined by its relation to the epistemic ideal – it is a matter of approximating the epistemic ideal to a sufficiently high degree given the limitations of one’s doxastic capacities. In this respect, Frege did much better than most of us.

In summary, the notion of an epistemic position corresponds to an epistemic ideal, which abstracts away from the contingent doxastic limits of particular subjects. Two subjects may be in the same epistemic position with respect to which propositions they have justification to believe or know, although one has the doxastic capacity to take advantage of his epistemic position, whereas the other does not. On this view, which propositions one is in an epistemic position to justifiably believe or know is constrained not by one’s doxastic capacities, but rather by an epistemic ideal, which one may be doxastically incapable of realizing.

6. The Threat of Collapse

The purpose of Williamson’s anti-luminosity argument is to debunk what he calls ‘the myth of epistemic transparency’. My aim in this paper has been to defend epistemic transparency against this attempt at debunking. Williamson’s argument establishes that there are no doxastically transparent conditions, but it fails to establish that there are no epistemically transparent conditions. Moreover, this has important ramifications for Williamson’s knowledge-first epistemology. Recall that epistemic transparency provides the basis of an argument against factive mentalism:

- (1) Evidence is transparent
- (2) Evidence is transparent only if the mental states that determine evidence are themselves transparent

(3) Factive mental states are not transparent

(4) Therefore, evidence is not determined by factive mental states

Williamson's response is to argue that there are no transparent conditions except trivial ones. If this response succeeds on the doxastic interpretation of transparency, but not the epistemic interpretation, then we can reinstate the original form of argument against Williamson's knowledge-first epistemology.

At this stage, there is an alternative strategy for blocking the argument against knowledge-first epistemology. Why not insist that factive mental states, no less than non-factive mental states, are transparent in the epistemic sense, but not the doxastic sense? If so, then my defensive strategy fails, since the distinction between epistemic and doxastic interpretations of transparency fails to capture any genuine epistemic asymmetry between factive and non-factive mental states.

If factive mental states are transparent, then one is always in a position to know whether or not one is in the so-called 'bad case' in which a skeptical scenario obtains. However, this seems counterintuitive. As Williamson himself acknowledges, "Part of the badness of the bad case is that one cannot know just how bad one's case is." (2000: 165)

In response, however, one might invoke the distinction between epistemic and doxastic notions of transparency. If one is in the bad case, then perhaps one is in an epistemic position to know that one is in the bad case, but not a doxastic position, since one lacks the doxastic capacity to take advantage of one's epistemic position. Why not say what I have already said in response to the problem of the speckled hen – namely, that one's evidence differs between the good case and the bad case in such a way that one's doxastic dispositions are insensitive to the evidential difference between them?

My objection to this proposal is that it collapses the distinction between two epistemic ideals, which should be kept apart – namely, being in an epistemic position to know a proposition and being in an epistemic position to justifiably believe it. The crux of the distinction is *factivity*: one is in a position to know a proposition only if it is true, whereas one may be in a position to justifiably believe a false proposition. The current proposal collapses this distinction, since it entails that one is in an epistemic position to justifiably believe a proposition only if one is in an epistemic position to know it. This follows from the epistemic transparency of knowledge together with a further assumption about defeaters. The argument is as follows:

- (1) **Transparency:** If one is not in an epistemic position to know that p , then one is in an epistemic position to know that one is not in an epistemic position to know that p ($\sim Kp \rightarrow K\sim Kp$)
- (2) **Defeat:** If one is in an epistemic position to know that one is not in an epistemic position to know that p , then one is not in an epistemic position to justifiably believe that p ($K\sim Kp \rightarrow \sim Jp$)
- (3) **Collapse:** If one is not in an epistemic position to know that p , then one is not in an epistemic position to justifiably believe that p ($\sim Kp \rightarrow \sim Jp$)

The conclusion of the argument is that one is in an epistemic position to justifiably believe a proposition only if one is in an epistemic position to know it. Since one is in an epistemic position to know a proposition only if it is true, it follows that one is in an epistemic position to justifiably believe only true propositions. Moreover, justifiably believing a proposition is a matter of taking advantage of one's epistemic position to justifiably believe it. So, if one is in an epistemic position to justifiably believe only true propositions, then it follows that there are no justified false beliefs.

A proponent of Transparency has two options: either to accept Collapse or to deny Defeat. In what follows, I will explain why neither option is available to a proponent of Williamson's brand of knowledge-first epistemology. Moreover, I will suggest that neither option is plausible on its own terms.

I begin by noting that Williamson is committed to accepting a close relative of Defeat by his claim that knowledge is the norm of assertion and belief, which he articulates in the form of the knowledge rule:

- **The knowledge rule:** One must: assert [or believe] p only if one knows p . (2000: 243)

According to the knowledge rule, it is permissible to believe a proposition only if one knows it. However, it does not follow that it is reasonable, rational or justified to believe a proposition only if one knows it.¹⁶ Instead, Williamson claims that we can explain when it is reasonable to believe a proposition by appealing to the knowledge rule together with the following bridge principle:

- **The bridge principle:** If one must (ϕ only if p is true), then one should (ϕ only if one has evidence that p is true). The transition from 'must' to 'should' represents the transition from what a rule forbids to what it provides a reason to do. (2000: 245)

The knowledge rule plus the bridge principle yields the following derived rule:

- **The derived rule:** One should believe p only if one has evidence that one knows p

Thus, if it is permissible to believe a proposition only if one knows it, then it is reasonable to believe a proposition only if one has evidence that justifies believing that one knows it. As Williamson writes, “the knowledge rule for assertion corresponds to the norm that one should believe p only if one knows p . Given that norm, it is not reasonable to believe p when one knows that one does not know p .” (2000: 255-6)

Williamson’s claim is further supported by examples.¹⁷ In Goldman’s (1976) example, Henry is in fake barn country, so he does not know that there is a barn ahead. But he does not know that he is in fake barn country, so it is reasonable for him to believe that there is a barn ahead. However, if Henry learns that he is in fake barn country, then it is no longer reasonable for him to believe that there is a barn ahead, since he now knows that he does not know it. Moreover, the epistemic claim of Defeat applies regardless of whatever Henry believes. If he is in an epistemic position to know that he is not in an epistemic position to know that there is a barn ahead, then he is not in an epistemic position to justifiably believe it.

If Defeat is accepted, then the only remaining option is to accept Collapse. However, this is highly counterintuitive. Indeed, I will suggest below that it is worse than counterintuitive. In philosophy, of course, no position is without its defenders. For instance, Jonathan Sutton (2005) has argued that a belief is justified if and only if it is knowledge. He therefore denies that there are any justified false beliefs or even Gettier cases in which justified true beliefs are not knowledge. By contrast, Williamson does not accept Collapse. For instance, Gettier cases play an important role in his argument that knowledge is a factive mental state.¹⁸ Moreover, he explicitly allows that one has justification to believe false propositions, which are made epistemically probable by true propositions that one knows.¹⁹

Williamson’s equation of evidence and knowledge does not imply that one has justification to believe only true propositions, since evidence is *what justifies*, rather than *what is justified*.²⁰ However, it does imply that one’s evidence in the good case is different from one’s evidence in the bad case. In each case, one’s evidence includes the proposition that it seems that one has hands, but in the good case one’s evidence also includes the proposition that one has hands. One’s evidence in the good case

entails that one has hands and so justifies believing that one has hands to the highest degree. By contrast, one's evidence in the bad case does not entail, but only makes it probable that one has hands and so justifies believing that one has hands to less than the highest degree. On Williamson's view, a belief is justified to the highest degree only if it is knowledge, but it does not follow that a belief is justified only if it is knowledge, since some beliefs are justified to less than the highest degree.²¹

As we have seen, Williamson does not accept Collapse. But since others do, we should consider what there is to be said against this view. It goes without saying that it is deeply counterintuitive to deny that justified beliefs can be false. But is there anything at stake here besides brute intuition?

I suggest that what is at stake is the nature of an epistemic ideal. What is lost on the factive conception of epistemic transparency is the idea that there is an epistemic ideal – and a corresponding dimension of epistemic evaluation – that is distinct from both omniscience and infallibility. On the non-factive conception of epistemic transparency, the ideal epistemic subject is omniscient and infallible about only very a limited domain of facts. This includes epistemic facts about which propositions comprise her evidence and which propositions she has justification to believe. It also includes non-epistemic facts about her non-factive mental states, which determine those epistemic facts. However, the ideal epistemic agent is not omniscient or infallible about facts which extend beyond her subjective point of view on the world. This is essential for making sense of an important epistemic ideal, which captures what one ought to believe given the limitations of one's subjective point of view on an objective world.

7. Conclusions

Epistemic transparency is central to the debate between factive and non-factive versions of mentalism about evidence. If evidence is transparent, then factive mentalism is false, since factive mental states are not transparent. Williamson defends factive mentalism by arguing that epistemic transparency is a myth, since there are no transparent conditions except trivial ones. This paper responds by drawing a distinction between epistemic and doxastic notions of transparency. Williamson's argument succeeds in showing that no conditions are doxastically transparent, but it fails to show that no conditions are epistemically transparent. Moreover, this is sufficient to reinstate the original form of argument against factive mentalism.

At this stage, one might attempt to defend factive mentalism by arguing that factive mental states are epistemically transparent, but not doxastically transparent. However, my objection is that this strategy collapses the distinction between epistemic ideals associated with knowing and justifiably believing. We should therefore deny that factive mental states are epistemically transparent.

In conclusion, epistemic transparency may provide the basis of a compelling argument against factive mentalism. However, it is beyond the scope of this paper to settle the debate between factive and non-factive versions of mentalism. My aim here is to defend epistemic transparency against Williamson's counterarguments. It is a further task to give positive arguments that evidence is epistemically transparent and that evidence is determined by non-factive mental states, which are also epistemically transparent. This is a task that I leave for another occasion.²²

References

- Alston, William (1989) *Epistemic Justification: Essays in the Theory of Knowledge*, Ithaca, NY: Cornell University Press.
- Berker, Selim (2008) "Luminosity Regained" *Philosophers' Imprint* 8 (2): 1-22.
- Brueckner, Anthony and M. Oreste Fiocco (2002) "Williamson's Anti-Luminosity Argument" *Philosophical Studies* 110: 285-93.
- Christensen, David (2004) *Putting Logic in its Place: Formal Constraints on Rational Belief*, Oxford University Press.
- Conee, Earl (2005) "The Comforts of Home" *Philosophy and Phenomenological Research* 70 (2): 444-451.
- Conee, Earl and Richard Feldman (2001) "Internalism Defended" *American Philosophical Quarterly* 38: 1-18.
- Feldman, Richard (2000) "The Ethics of Belief" *Philosophy and Phenomenological Research* 60 (3): 667-695.
- Goldman, Alvin (1976) "Discrimination and Perceptual Knowledge" *Journal of Philosophy* 73: 771-91.
- Goldman, Alvin (1979) "What is Justified Belief?" in (ed.) G. Pappas, *Justification and Knowledge*, Dordrecht: Reidel.
- Greenough, Patrick (ms) "Our Cognitive Homes"
- Hawthorne, John (2005) "Knowledge and Evidence" *Philosophy and Phenomenological Research* 70 (2): 452-8.
- McDowell, John (1995) "Knowledge and the Internal" *Philosophy and Phenomenological Research* 55 (4): 877-93.
- Neta, Ram and Guy Rohrbaugh (2004) "Luminosity and the Safety of Knowledge" *Pacific Philosophical Quarterly* 85: 396-406.
- Neta, Ram and Duncan Pritchard (2007) "McDowell and the New Evil Genius" *Philosophy and Phenomenological Research* 74: 381-96.
- Pollock, John and Joseph Cruz (1999) *Contemporary Theories of Knowledge*, Rowman and Littlefield.
- Raffman, Diana (1995) "On the Persistence of Phenomenology" in (ed.) T. Metzinger, *Conscious Experience*, Schoningh.
- Ramachandran, Murali (2009) "Anti-Luminosity: Four Unsuccessful Strategies" *Australasian Journal of Philosophy* 87 (4): 659-73.
- Reed, Baron (2006) "Shelter for the Cognitively Homeless" *Synthese* 148: 303-8.

- Smithies, Declan (forthcoming) “The Normative Role of Knowledge” *Nous*.
- Sosa, Ernest (2003) “Privileged Access” in (eds.) Q. Smith and A. Jokic, *Consciousness: New Philosophical Perspectives*, Oxford University Press.
- Sosa, Ernest (2009) “Timothy Williamson’s *Knowledge and Its Limits*” in (eds.) P. Greenough and D. Pritchard, *Williamson on Knowledge*, Oxford University Press.
- Sutton, Jonathan (2005) “Stick To What You Know” *Nous* 39 (3): 359-96.
- Steup, Matthias (2009) “Are Mental States Luminous?” in (eds.) P. Greenough and D. Pritchard, *Williamson on Knowledge*, Oxford University Press.
- Weatherson, Brian (2004) “Luminous Margins” *Australasian Journal of Philosophy* 83: 373-83.
- Wedgwood, Ralph (2002) “Internalism Explained” *Philosophy and Phenomenological Research* 65 (2): 349-69.
- Williamson, Timothy (1994) *Vagueness*, London: Routledge.
- Williamson, Timothy (2000) *Knowledge and Its Limits*, Oxford University Press.
- Williamson, Timothy (2005) “Replies to Commentators” *Philosophy and Phenomenological Research* 70 (2): 468-91.
- Williamson, Timothy (2008) “Why Epistemology Can’t Be Operationalized” in (ed.) Q. Smith, *Epistemology: New Philosophical Essays*, Oxford University Press.
- Williamson, Timothy (2009) “Replies to Critics” in (eds.) P. Greenough and D. Pritchard, *Williamson on Knowledge*, Oxford University Press.

¹ *Evidentialism* is the thesis that one’s evidence determines which propositions one has justification to believe. If evidentialism is true, then mentalism about *evidence* entails mentalism about *justification*, the thesis that one’s mental states determine which propositions one has justification to believe.

² The rationale for this assumption is that knowledge of epistemic truths depends on knowledge of non-epistemic truths, which determine those epistemic truths.

³ My own view is that epistemic transparency is indispensable for motivating non-factive versions of mentalism. However, Pollock and Cruz (1999), Conee and Feldman (2001) and Wedgwood (2002) argue for non-factive versions of mentalism without appealing to epistemic transparency.

⁴ Williamson (2000: Ch.4). My presentation draws on Weatherson (2004) and Berker (2008).

⁵ “One can believe that C obtains and be safe from error in doing so even if C does not safely obtain, if whether one believes is sufficiently sensitive to whether C obtains. For example, one may be safe from error in believing that the child is not falling even though she is not safe from falling, if one is in a good position to see her but not to help her.” (Williamson 2000: 127)

⁶ Offensive strategies are pursued by Brueckner and Fiocco (2002), Neta and Rohrbaugh (2004), Weatherson (2004), Berker (2008), Ramachandran (2009) and Steup (2009). I am not persuaded by these strategies, although I do not have the space to survey them here.

⁷ Defensive strategies are pursued by Hawthorne (2005), Conee (2005), Reed (2006) and Sosa (2009); see Williamson (2005) and (2009) for criticisms. These authors argue that there are some conditions C such that one is always in a position to know that C obtains if C obtains determinately, intensely, or in some other specific way. My own strategy is rather different from all of these.

⁸ “Margin for error principles...seem specific to knowledge. If one Φ s a proposition in a situation s , one leaves a margin for error only if that proposition is true in all cases similar enough to s . Since s is certainly similar enough to itself, the proposition must be true in s . Thus if Φ ing requires a margin for error, one Φ s only true propositions. Knowledge is such an attitude; reasonable belief is not.” (Williamson 1994: 224-5)

⁹ This would be accepted by some proponents of factive mentalism, if not by Williamson himself. For instance, see McDowell (1995) and the discussion in Neta and Pritchard (2007).

¹⁰ This was my own view in Smithies (2006: 84, fn.9).

¹¹ Compare Raffman’s (1995) memory constraint: given the limitations on perceptual memory, one’s ability to perceptually recognize or reidentify a colour shade over time is less accurate than one’s ability to perceptually discriminate it from other shades at the same time.

¹² “How then would one distinguish (1) an *unjustified* ‘introspective’ judgement, say that one’s image has 48 speckles, when it is a true judgement, and one issued in full view of the image with that specific character, from (2) a *justified* ‘introspective’ judgement, say that one’s image has 3 speckles? The relevant distinction is that the latter judgement is both (a) *safe* and (b) *virtuous*, or so I wish to suggest. It is ‘safe’ because in the circumstances not easily *would* one believe as one does without being right. It is ‘virtuous’ because one’s belief derives from a way of forming beliefs that is an intellectual virtue, one that in our normal situation for forming such beliefs would tend strongly enough to give us beliefs that are safe.” (Sosa 2003: 290)

¹³ Here I assume that cases in which one justifiably believes that C obtains on the basis of different evidence are not relevantly close cases.

¹⁴ Compare Goldman’s definition of *ex ante* justification in terms of *ex post* justification: “Person S is *ex ante* justified in believing p at t if and only if there is a reliable belief-forming operation available to S which is such that if S applied that operation to his total cognitive state at t, S would believe p at t-plus-delta (for a suitably small delta) and that belief would be *ex post* justified.” (1979: 21)

¹⁵ See Smithies, “Why Care About Justification?” for a more detailed discussion of epistemic idealization and its importance in the theory of justification.

¹⁶ “On this analogy between assertion and belief, the knowledge rule for assertion does not correspond to an identification of reasonable belief with knowledge. The rule makes knowledge the condition for permissible assertion, not reasonable assertion. One may reasonably do something impermissible because one reasonably but falsely believes it to be permissible.” (Williamson 2000: 256)

¹⁷ Smithies (forthcoming) uses examples like this one to motivate the JK rule, according to which one has justification to believe a proposition if and only if one has justification to believe that one is in an epistemic position to know it. For present purposes, however, we need only the $\sim K\sim K$ rule, according to which one has justification to believe a proposition if and only if one is not in an epistemic position to know that one is not in an epistemic position to know it.

¹⁸ See Williamson (2000: 2, 8, 30).

¹⁹ “True propositions can make a false proposition probable, as when someone is skilfully framed for a crime of which she is innocent. If perceptual evidence in the case of illusions consists of true propositions, what are they? The obvious answer is: the proposition that things appear to be that way.” (Williamson 2000: 198)

²⁰ “Although it has been shown that *what is justified* need not be knowledge, even when it is true, it has not been shown that *what justifies* need not be knowledge.” (Williamson 2000: 185)

²¹ “On this view, the difference is that in the non-sceptical scenario the subject’s total evidence *entails* that the child is playing, whereas in the sceptical scenario the subject’s total evidence does not entail that the child is playing, although it does (misleadingly) make it probable that the child is playing. Consequently, the belief that the child is playing is *more justified* in the non-sceptical scenario than it is in the sceptical scenario, even though in both scenarios it has considerable justification.” (Williamson 2005: xxx)

²² Many thanks for excellent and stimulating discussions with David Chalmers, Patrick Greenough, John Maier, Nicholas Silins, Daniel Stoljar and Ralph Wedgwood.