

No State *A Priori* Known to Be Factive Is Mental*

May 10, 2016

1 Mind, World and Knowledge

Following *content externalism*, let's assume that, by and large, how the *mind* is depends on how the (*external*) *world* is. For instance, given that snow is H₂O in a certain state, let's assume that, necessarily, one can only believe that snow is white if hydrogen exists. Still, even on this conception of the mind, it would seem to make sense to draw a distinction between *properties of the mind* and *relations that the mind bears to the world* (in the metaphysically interesting sense of 'property' and 'relation' in which "*relational*

*I first gave the core argument of this paper in 2007 during an epistemology course co-taught with Dylan Dodd at the University of St Andrews. Earlier versions of the material in the paper have then been presented in 2013 at the session *Scepticism and Epistemic Circularity* of the 23rd World Congress of Philosophy *Philosophy as Inquiry and Way of Life* (University of Athens), at the 6th NIP Basic Knowledge Workshop *Themes in Basic Knowledge* (University of Aberdeen) and at the LOGOS Seminar (University of Barcelona); in 2015, at the 1st LanCog Day (University of Lisbon) and at a research seminar at Yonsei University; in 2016, at the IIF-SADAF Seminar (University of Buenos Aires). I'd like to thank all these audiences for very stimulating comments and discussions. Special thanks go to Lauren Ashwell, Eduardo Barrio, Colin Caret, Roy Cook, Aurélien Darbellay, Laura Delgado, José Díez, Dylan Dodd, Philip Ebert, Miguel Ángel Fernández, Nakul Krishna, Dan López de Sa, Ofra Magidor, Aidan McGlynn, Alberto Moretti, Eleonora Orlando, Nikolaj Pedersen, Sònia Roca, Sven Rosenkranz, Ernie Sosa, Andreas Stokke, Célia Teixeira, Crispin Wright, David Yates, José Zalabardo. I'm also grateful to the editors for inviting me to contribute to this issue and for their patience throughout the process. At different stages during the writing of the paper, I've benefitted from an AHRC Postdoctoral Research Fellowship, from the FP7 Marie Curie Intra-European Research Fellowship 301493 on *A Non-Contractive Theory of Naive Semantic Properties: Logical Developments and Metaphysical Foundations* (NTNSP) and from the FCT Research Fellowship IF/01202/2013 on *Tolerance and Instability: The Substructure of Cognitions, Transitions and Collections* (TI), as well as from partial funds from the project CONSOLIDER-INGENIO 2010 CSD2009-00056 of the Spanish Ministry of Science and Innovation on *Philosophy of Perspectival Thoughts and Facts* (PERSP), from the FP7 Marie Curie Initial Training Network 238128 on *Perspectival Thoughts and Facts* (PETAF), from the project FFI2011-25626 of the Spanish Ministry of Science and Innovation on *Reference, Self-Reference and Empirical Data* and from the project FFI2012-35026 of the Spanish Ministry of Economy and Competition on *The Makings of Truth: Nature, Extent, and Applications of Truthmaking*.

properties” such as e.g. being a sibling and “*modified properties*” such as e.g. walking with a stick are not properties but relations). Properties of the mind are presumably things such as e.g. believing that snow is white, seeing white, wondering what colour snow has; relations that the mind bears to the world are presumably things such as e.g. tracking that snow is white, measuring snow, testing what colour snow has.¹

The distinction is naturally introduced along such lines, but, ultimately, these are somewhat misleading. For they evidently presuppose that there is such an object as “the mind”, a reifying presupposition which, while harmless for many other purposes, for those of this paper has worse consequences than merely that of an animistic ringing. To wit, there are properties of the mind that, for most purposes, we shouldn’t count as “*mental*”: mind-body substance monists can observe that, if the body exemplifies the property of, say, weighing 72 kg, it presumably follows that the mind also exemplifies that property; the rest of us can still observe that the mind exemplifies the property of being an object. Not every property of the mind is a mental property.² Fortunately, just as one can distinguish the *chemical* properties and relations (for example, being a molecule of H₂O) of an object from its other properties and relations (for example, being a molecule of H₂O surrounded by dark matter) without presupposing that there is such an object as “the chemical object”, and do so on the basis of a conception of a range of properties and relations that are so nomologically networked among themselves as to constitute a distinct level of reality—the chemical level—so can one distinguish the *mental* properties and relations (for example, believing that snow is white) of an object from its other properties and relations (for example, tracking that snow is white) without presupposing that there is such an object as “the mind”, and do so on the basis of a conception of a range of properties and relations that are so nomologically networked among themselves as to constitute a distinct level of reality—the (*personal*) *mental* level. On this conception of the mental, it would seem to make sense to draw a distinction between one’s *mental properties and relations* and one’s *non-mental properties and relations*, and in particular between one’s mental properties and one’s non-mental relations to the world which nevertheless involve to a substantial extent some of one’s mental properties. Having clarified this, I’ll henceforth use the pair ‘property of the mind’/‘relation that the mind bears to the world’ and its relatives simply to mark in a pictorial way the latter distinction, which is really the one of interest for the purposes of this paper; similarly, I’ll henceforth use ‘the mind’ and its relatives simply to point in a pictorial way to the range of mental properties, which is really what is of interest for the purposes of this paper.

¹Notice that our conceptual repertoire also arguably envisages as central quite a few relations that the mind bears to objects that I’m officially excluding from “the world” (since I’m equating it with the external world). One prominent kind of such relations is formed by relations that the mind bears to *minds*, presumably things such as e.g. revealing that snow is white, discovering snow, asking what colour snow has.

²Don’t say that mental properties are those properties of the mind *that only minds exemplify*. On the one hand, only minds might exemplify the property of being realised either by physical system s_0 or by physical system s_1 or by physical system $s_2 \dots$, but that is a property that, for most purposes, we shouldn’t count as “mental”. On the other hand, for most purposes we should count as “mental” the property of calculating, even if also objects that are not minds (such as e.g. calculators) exemplify it.

With this distinction in place, in theorising about *knowledge* it then becomes a natural question to ask on which side of the distinction knowledge falls.³ Is knowledge a property of the mind—akin in the relevant respect to believing that snow is white—or rather a relation that the mind bears to the world—akin in the relevant respect to tracking that snow is white?⁴ The answer would seem clear, and is one of the few things that the allegedly traditional definition of knowledge as justified true belief arguably gets right (if brutally so). However, Williamson [2000] for one has influentially disagreed, and so it’ll be worth not to be too dogmatic about this issue and investigate it in some detail.

2 Mind, Introspection and Knowledge

Williamson [2000], pp. 23–27 himself discusses a little one natural objection to the mentality of knowledge. By way of ostensive definition, let ‘*introspection*’ cover whatever method (or methods) we use when we know about our mental properties in the typical “introspective” way we’re all aware of. Introspection is arguably *immediate* in the sense that, roughly, if one knows something by introspection one’s knowledge (or, if you prefer, justification) does not epistemically depend on one’s knowledge of any proposition. Introspection is also arguably *non-perceptual* in the sense that, roughly, if one knows something by introspection one’s knowledge (or, if you prefer, justification) does not epistemically depend on perception.^{5,6} Unsurprisingly, it is not wholly uncontroversial that introspection is either immediate or non-perceptual (Ryle [1949] can be seen as contesting both immediacy and non-perceptuality),⁷ but these are nevertheless extremely plausible

³There arguably are different kinds of knowledge: *propositional* knowledge (such as e.g. knowing that snow is white), *objectual* knowledge (such as e.g. knowing snow), *practical* knowledge (such as e.g. knowing how to ski on snow) etc. (Stanley and Williamson [2001] would notoriously disagree about the last item. For what it’s worth, I myself have given reason against the reduction of practical knowledge to propositional knowledge in Sgaravatti and Zardini [2008]; Zardini [2013a].) Throughout, unless otherwise specified ‘know’ and its relatives are understood to express propositional knowledge.

⁴Obviously, one has at least some knowledge of, say, one’s mental properties (for example, one knows that one believes that snow is white) and such pieces of knowledge are not relations that the mind bears to the *world*. Throughout, overtly unrestricted talk of knowledge as a “relation that the mind bears to the world” and its relatives are understood as covertly restricted to such pieces of knowledge as knowledge that snow is white.

⁵Throughout, by ‘*perception*’ and its relatives I mean, roughly, the *system (or systems) by which one’s sensory modalities represent one’s environment*. Emphatically, this does not include the variety of properties traditionally grouped under the label ‘*inner sense*’ (for example, the property of feeling pain). Inner sensations are also mental properties, and it is very plausible that one’s knowledge of one’s inner sensations does epistemically depend on them (for example, it is very plausible that one’s knowledge that one feels pain does epistemically depend on one’s feeling pain).

⁶Notice that the *immediate/non-immediate* distinction is *orthogonal* to the *non-perceptual/perceptual* distinction (where knowledge is *non-immediate* iff it is not immediate, and knowledge is *perceptual* iff it is not non-perceptual): immediate knowledge can be perceptual (as in the case of e.g. visual information), non-perceptual knowledge can be non-immediate (as in the case of e.g. mathematical theorems) and non-immediate knowledge can be perceptual (as in the case of e.g. physical theories).

⁷One might want to allow for local systematic non-immediate or perceptual methods. For example, one might hold that one method for knowing about one’s beliefs is provided by the so-called “*Evans-*

assumptions⁸ that I'll henceforth make.

Now, a quick survey of paradigmatic mental properties (belief, hope, love etc.) suggests the principle of *introspection of mind*:

style method" (Evans [1982], p. 225), interpreting such method as relying on a straightforward inference from ' P ' to 'I believe that P ' (so that, for instance, one's knowledge that one believes that the theory of general relativity is true would epistemically depend on one's knowledge that the theory of general relativity is true, thus being neither immediate nor non-perceptual). For what it's worth, I find the idea (defended by e.g. Byrne [2005]) that the Evans-style method so interpreted *is as a matter of fact systematically employed* to yield knowledge that one believes something utterly implausible. Firstly, presumably systematic employment of the method is supposed to cover also cases in which one's belief that P is *false and completely unjustified*; but inferences from a false and completely unjustified premise very plausibly cannot yield *knowledge* of their conclusion. Secondly, the method is extremely unreliable when employed in contexts of *suppositional reasoning*, so that one presumably needs to have independent knowledge that one is not in such contexts in order for employment of the method to yield knowledge; but that boils down to the requirement that one have independent knowledge that one does *believe*—rather than merely *suppose*—the premises in question, which makes employment of the method useless. Thirdly, contrary to what systematic employment of the method implies, generally, being able to assess whether one believes something is independent of being able to assess whether anything as *esoteric* as the theory of general relativity is true. Fourthly, contrary to what systematic employment of the method implies, typically, one is not in a position to be significantly *more certain*, say, that one believes that there are infinitely many prime numbers than that one believes that the theory of general relativity is true, even though one has nothing less than a *deductive proof* that there are infinitely many prime numbers but only an *abductive argument* that the theory of general relativity is true. Fifthly, contrary to what systematic employment of the method implies, ordinarily, one does not know that one believes something *by such a complex argument* as those used to justify that the theory of general relativity is true. Sixthly, contrary to what systematic employment of the method implies, ordinarily, one subject's knowledge that she *believes* that the theory of general relativity is true does not involve a significantly different *kind of epistemic achievement* than another subject's knowledge, say, that she *does not believe* that the theory of general relativity is true. (Notice that the first four points, if correct, actually establish the stronger conclusion that the method *cannot be systematically employed* to yield knowledge that one believes something. My suspicion is that the idea that it can be so employed derives its spurious plausibility from the much more attractive idea that the method can be systematically employed to yield knowledge that one *knows* something, see fn 36.) Henceforth setting these issues aside, it is fairly uncontroversial that the Evans-style method so understood cannot even be systematically employed in order to yield knowledge either that one *does not believe* that P or that one *is more confident than not* that P or that one *believed* that P *two seconds ago* etc. But then the relevant method (or methods) systematically employed to yield those kinds of knowledge will *also* presumably be systematically employed to yield knowledge that one believes that P , and we can let 'introspection' cover such method (or methods) but not the Evans-style method. Anyways, even if the far-fetched hypothesis that the Evans-style method so understood is the *only* method systematically employed to yield knowledge that one believes or knows something were true, the method's *extremely local* character would make the hypothesis' truth quite irrelevant for the purposes of this paper (sketch of reason: at the relevant places, substitute knowledge that one knew that P two seconds ago for knowledge that one knows that P). (*Mutatis mutandis*, the points made apply also to other interpretations of the Evans-style method, e.g. the one defended by Fernández [2013].) Thanks to Lauren Ashwell and an anonymous referee of another paper of mine for discussion of some of these issues.

⁸Both assumptions are made extremely plausible by reflection on *paradigm cases* of introspection (for example, knowing by introspection that one is in pain). Moreover, it is extremely plausible that *introspection is immediate if it is non-perceptual*—on what other kind of knowledge would introspective knowledge epistemically depend if not on perceptual knowledge? (Knowledge of one's *practical and theoretical traits* (such as e.g. being courageous, being intelligent and being a philosopher) might actually epistemically depend on introspective knowledge, and much of what I say about introspection does not

(IM) Every mental property M is such that:

(IM^P) Metaphysically necessarily, if one is M one is in a position to know by introspection that one is M ;

(IM^N) Metaphysically necessarily, if one is not M one is in a position to know by introspection that one is not M

(see Alston [1971] for a seminal list and discussion of possible epistemic privileges of introspection), where one is in a position to know that P iff [P and, in order for one to (come to) know that P , one only needs to (come to) believe that P on grounds one has already available]⁹ (see Zardini [2013b], pp. 766–767 for further explications of the notion of being in a position to know). (IM) is not only suggested by a piecemeal *examination of cases*. The general claims it consists in are also plausible given our *conception of what one’s mind is and how one is related to it*: namely, as what is *most intimate to one*, and so as *always open to one’s view*.¹⁰ Such conception also inspires more detailed argumentative routes to (IM).

Firstly, while the world is given to one only *from a perspective* (for the purposes of this paper, especially in the respect in which some few things are open to one’s view but most are hidden) such perspective *itself*—i.e. one’s mind—is not in turn given to one only from

apply to such knowledge. For the purposes of this paper, we should only count as candidate for being “mental”, roughly, *intentional* or *phenomenal* properties.) And there are further reasons for thinking that introspection is non-perceptual. Firstly, if introspection were perceptual its epistemic structure would presumably be essentially the same as that of knowledge of *others’* mental properties by *observation* rather than by *testimony* (presumably, perception plus inference). Yet, *natural language* registers a stark *contrast* between the two cases which would then be difficult to explain: a speaker would describe herself with ‘I believe that Ann and Beth are having an affair’ but would describe another subject she has only observed with something like ‘She seems to believe that Ann and Beth are having an affair’. (Such contrast is in some respects even more acute in Japanese: for example, it seems to persist even in the case of knowledge by testimony (see e.g. Kamio [1995], pp. 251–261). The contrast is also in some respects even more patent in Tibetan (a language in which *evidentiality* is grammatically encoded): for example, barring conclusive observational evidence the direct-evidence marker ‘*dug* seems to be used in attributions of certain mental properties to oneself but not to others (see e.g. Garrett [2001], pp. 77–82). I should emphasise that I’m not in the least inclined to read epistemic structure directly off natural language, but I do think that the contrasts pointed to are suggestive nevertheless.) Notice that, if, on the contrary, the speaker has heard the other subject assert something to the effect that Ann and Beth are having an affair, she would indeed describe her with ‘She believes that Ann and Beth are having an affair’. Secondly, the *very high degree of certainty* with which we typically know about our mental properties would be hard to explain if such knowledge epistemically depended on perception, given the apparent *poverty of perceptual evidence* about one’s mental properties which is available to one. Thirdly, it would seem perfectly possible that there are subjects who *lack a perceptual system* but who *are still able to know about their mental properties by introspection*. (Notice that these three further reasons in themselves speak even more conclusively against the idea that *introspection can be reduced to perception and inference*; but, presumably, if introspection were perceptual it would be such precisely because it would be so reducible.)

⁹Throughout, I use square brackets to disambiguate constituent structure.

¹⁰*Maximal intimacy* corresponds not only to *openness to oneself*, but also to *hiddenness to others*. If *privileged self-knowledge* is rooted in such intimacy, then it has the same root as the *problem of knowledge of other minds*—the latter is the flip side of the former.

a *further perspective*. More accurately, surely there is no infinite regress of perspectives (which perspective if any would one then really occupy?), and, at whichever perspective the regress stops, that is surely the perspective that one really occupies—i.e. one’s mind. Therefore, one’s mind is given to one not *perspectivally*, but *absolutely*. Restricted to our own minds, we do occupy a “God’s-eye point of view”. Everything must be open to such point of view, and that presumably implies (IM).

Secondly, presumably the version of (IM) in which ‘*is in a position to know*’ is replaced by the *much weaker* ‘*can in principle know*’ is true (presumably, *there are no facts about one’s mental properties that are in principle unknowable by introspection*). Now, in general the gap between being in principle able to know and being in a position to know would only seem to concern further *inferential* or *empirical* work that one who can in principle know might need to do in order to get in a position to know.¹¹ But, as argued in fn 8, introspection is both *immediate* and *non-perceptual*, and so being in a position to know by introspection requires neither inferential nor empirical work. Therefore, being in principle able to know *by introspection* would seem to collapse on being in a position to know *by introspection*, and so (IM) follows.

Thirdly, a distinctive feature of the mind as the *locus of subjectivity* is that there is a guarantee that each state of *consciousness* (whose content is at least partially encoded, for example, by ‘Snow is white’) can be the object of a state of *self-consciousness* (whose content is at least partially encoded, for example, by ‘In this very state, I believe that snow is white’)—broadly as *per* Kant [1787], pp. 108–109, without such guarantee of being able to recognise oneself as the subject of a state of consciousness in what relevant sense could any such state be a state that *belongs to the subject*? The argument then splits. For the less adventurous, it suffices to note that, if the possibility of such *recognition* is indeed *guaranteed*, it must then be *independent of any further epistemic work*, so that the possibility in question is “*positional*” (i.e. the same as the one involved in the notion of being in a position to know), thus yielding (IM^P), but not (IM^N) (low risks, low returns). For the more adventurous, the argument continues by observing that the object of a state of self-consciousness must involve that state itself (if it is to count as a state of genuine *self-consciousness*), and so every state of consciousness must already involve an at least implicit state of self-consciousness. Consciousness is self-consciousness. However, self-consciousness being in this sense *omniscient* also implies that self-consciousness is *infallible* in the sense of always being correct. For an incorrect state of self-consciousness would be a state whose content is at least partially encoded, for example, by ‘In this very state, I believe that snow is white’ without its being at least partially encoded by ‘Snow is white’. But, by the argument just given, a state with a content with those two features would be

¹¹This probably requires an important qualification. Many psychological approaches postulate the existence of unconscious mental states that can only be made conscious by some sort of *broadly therapeutic* work going beyond what can be delivered by simple *reflection* (see the fifth next paragraph about the latter). The arguments in favour of (IM) do not apply to such states, which are in effect such that it is not metaphysically necessary that, if one is in them, one is in a position to know that one is in them. For the purposes of this paper, just as we don’t count subpersonal-level states as “mental” we shouldn’t count this clearly demarcated kind of states as “mental” either (obviously, for many other purposes the opposite policy would be much more appropriate).

as such the object of a state of self-consciousness, and one could not rationally be in the latter state while also being in the former state (given the former state's first feature). Self-consciousness is consciousness. Putting omniscience and infallibility together (and identifying introspection with self-consciousness), we get a version of (IM^P) and of its converse in which '*is in a position to know*' is replaced by something like '*believes*', which in turn constitutes a very good reason in favour of (IM^P) itself—indeed in favour of its stronger version in which '*is in a position to know*' is replaced by '*knows*'. And, from omniscience and infallibility, the analogous version of (IM^N) itself would similarly follow given the further assumption that, at some level, one's beliefs about one's mental properties are *consistent* and *maximal*, in the sense that one does not believe [that one is *M*] iff one believes that one is not *M* (high risks, high returns).¹²

Yet, plausible as it may seem (IM) creates an important problem for the mentality of knowledge, given that *it systematically fails for knowledge*. One might think that this failure can be *immediately inferred from common-or-garden examples*, but a lot of care should be exercised here. For example, if one lacks the concept of knowledge, or if one is on an LSD trip, one might not know, for instance, that one is on the moon, but one is not in a position to know that one does not know that one is on the moon, and so (IM^N) for knowledge fails. But that only shows that (IM) is more interestingly understood as restricted to subjects who *possess the relevant conceptual resources* and who *are alert* (and, indeed, (IM) will henceforth be so understood).

Even so, one might still think that (IM)'s failure for knowledge can be immediately inferred from common-or-garden examples. Williamson himself would seem to do so by putting forth the following example. Before Lincoln's assassination, one knows on the strength of ordinary evidence that Lincoln is the President of the US. Just after Lincoln's assassination, one is not yet aware of anything amiss. Then, at that time one does not know that Lincoln is the President of the US, but one is not in a position to know that one does not know that Lincoln is the President of the US. Such example clearly relies on one's *having undefeated misleading evidence* (the evidence one has just after Lincoln's assassination, which includes the fact that everything seems to be business as usual). But then examples that similarly rely on one's having undefeated misleading evidence can be used to show that (IM) fails *even for paradigmatic mental properties* like belief (Conee [2005], pp. 448–450). One is in the grip of a seemingly flawless eliminativist argument seemingly establishing that there are no such things as belief states. Eliminativism about mental properties does not have implications for the colour of snow, it is well argued for and it is false, and so it is perfectly rational for one to believe it while also believing that snow is white. Then, one believes that snow is white, but one is not in a position to know that one believes that snow is white. *If* (IM) is understood as including subjects having undefeated misleading evidence, it fails for paradigmatic mental properties just as well as for knowledge.¹³ But that only shows that (IM) is more interestingly understood as

¹²There are also arguments in favour of one or both components of (IM) as restricted to some *specific* mental properties: for example, Zardini [2013b], pp. 777–778 offers an argument in favour of (IM^P) as restricted to the property of being in pain.

¹³Williamson implies that the *kind of case* instantiated by the Lincoln example only applies to (IM^N)

restricted to subjects who *do not have undefeated misleading evidence* (and, indeed, (IM) will henceforth be so understood; see Zardini [2012], pp. 376–377 for further, and more general, discussion). So understood, (IM) for knowledge would seem to be no less immune to immediate refutation from common-or-garden examples than (IM) for paradigmatic mental properties is.

However, in the grip of the undefeated misleading evidence provided by the Lincoln example Williamson does not think so, and thus he sees himself faced with the challenge of an apparent difference between paradigmatic mental properties and knowledge. He addresses the challenge *via* a two-fold strategy. On the one hand, he tries to *play down* the extent to which paradigmatic mental properties satisfy (IM). This in turn he does *via* a two-fold strategy. On the one hand (of the first “main hand”), he gives (sundry) *common-or-garden examples* where one does not seem to be in a position to know about one’s paradigmatic mental properties. Let’s take these in turn.

There is first the example (slightly expanded for concreteness) of one only discovering that one hopes for Juventus’ victory after experiencing disappointment at Juventus’ drawing, before which one only felt indifference. Williamson claims that, before the draw, one hopes for Juventus’ victory, but one is not in a position to know that one hopes for Juventus’ victory. Now, it is not perfectly clear to me how, in this discussion, Williamson understands the rather vague phrase ‘be in a position to know’ and its relatives, and if the positional modality implicit in that phrase can be understood in a strict enough way—so as to include as possible ground for knowledge only, say, one’s superficial feeling of indifference—surely it can then be claimed that one “is not in a position to know” that one hopes for Juventus’ victory. That would be at best a Draconian understanding, as witnessed by the pronounced oddity of the claim under usual conditions. Still, to fix ideas I myself have officially adopted an understanding of this kind in the seventh last paragraph (and mean to stick to it), since the example is naturally understood so that the only ground one has already available is one’s superficial feeling of indifference. Even so, such example clearly relies on one’s *not reflecting enough on the relevant issue*, since, as far as the example has been described, I see no reason to doubt that, if one had reflected enough on the issue, one would have seen through one’s superficial feeling of indifference to one’s deep inclinations. But then examples that similarly rely on one’s not reflecting enough on the relevant issue can be used to show that a principle similar to (IM) fails *even for paradigmatically transparent facts* like arithmetical ones. One goes too quickly over the definition of a prime number assuming that it rules out all even numbers as these are always divisible by 2 too, overlooking the fact that, in the peculiar case of 2 itself, such divisibility is compatible with primeness. One still needs to see such compatibility, and so this is not a ground one has already available. Then, 2 is prime, but one is not in a position to know that 2 is prime. *If* (IM) and similar principles are understood as

for knowledge, and that a very different kind of case (and considerations) is needed to establish failure of (IM^P) for knowledge (see the tenth next paragraph). But, as the kind of case instantiated by the Lincoln example is at bottom the kind of case in which one has undefeated misleading evidence, that is exactly the kind of case which the text has just used to establish failure of (IM^P) for belief, and which, clearly, can similarly be used to establish failure of (IM^P) for knowledge.

including subjects not reflecting enough on the relevant issue, they fail for paradigmatically transparent facts just as well as for paradigmatic mental properties. But that only shows that (IM) and similar principles are more interestingly understood as restricted to subjects who *reflect enough on the relevant issue* (and, indeed, (IM) will henceforth be so understood).¹⁴

There is then the example of one's being too self-pitying, and so mistaking an itch for a pain. Presumably, the idea here is that one has a itch, but one is not in a position to know that one has a itch. Such example clearly relies on one's *having psychological biases*. But then examples that similarly rely on one's having psychological biases can be used to show that a principle similar to (IM) fails *even for paradigmatically transparent facts* like arithmetical ones. One might be stuck with the belief that a number can only be divided by another number, and so believe that 2 is not divisible by 2, and hence mistake it for odd. Then, 2 is even, but one is not in a position to know that 2 is even. *If* (IM) and similar principles are understood as including subjects having psychological biases, they fail for paradigmatically transparent facts just as well as for paradigmatic mental properties. But that only shows that (IM) and similar principles are more interestingly understood as restricted to subjects who *do not have psychological biases* (and, indeed, (IM) will henceforth be so understood).¹⁵

There are thus excellent reasons for restricting to subjects who, in addition to possessing the relevant conceptual resources and being alert, also do not have undefeated misleading evidence, also reflect enough on the relevant issue and also do not have psychological biases. Indeed, there are further *epistemic flaws* by which a subject might be affected in a situation, and the restriction should rule out all such cases. I take it to be clear that our intuitive conception of an epistemic flaw, *vague* as it may be, is *substantial* enough as *not* to turn (IM) (and similar principles) *trivially true* once such restriction is imposed.¹⁶

We can push this a bit further. For, once it is realised that a restriction to certain subjects is needed anyways in (IM), *any modalisation of knowledge becomes dispensable in it*. For example, given my official understanding of positional modality in the tenth

¹⁴Williamson also offers an example in which it is not straightforward whether one believes something or merely fancies it. I take it that such example does not present issues different from those addressed in the text.

¹⁵Williamson also offers the example of one's being too little self-pitying, and so mistaking a pain for an itch. I take it that such example does not present issues different from those addressed in the text. I should mention that this style of argument against knowability principles of some sort or other is uncharacteristic of Williamson's, who has elsewhere offered very challenging arguments that do not rely on such cheap possibilities (I'll mention those arguments in the sixth and seventh next paragraphs).

¹⁶Does the restriction make facts about, say, *one's surroundings* just as directly accessible as facts about *one's mental properties*? Not at all, for at least two reasons. Firstly, it might be the case that there is a cube in one's surroundings but one is not in a position to know that as one is looking somewhere else (and so has *no relevant grounds already available in the first place*). Secondly, it might be the case that there is a cube in one's surroundings but one is not in a position to know that as there are also a lot of cube holograms around (and so the relevant grounds one does have already available are *not sufficient for knowledge*; cf Burge [1988], pp. 657–658 on the possibility of “*brute error*” in perception and its impossibility in introspection).

last paragraph in my particular case the modalisation does the only job of allowing one to form the relevant beliefs, but we can now straightforwardly further restrict instead to subjects who also *form the relevant beliefs*. Making the restrictions explicit and dropping (IM)'s modalisation, we thus get the principle of *restricted introspection of mind*:

(RIM) Every mental property M is such that:

(RIM^P) Metaphysically necessarily, if an epistemically flawless subject is M , and forms a belief as to whether she is M on introspective grounds she has already available, she knows by introspection that she is M ;

(RIM^N) Metaphysically necessarily, if an epistemically flawless subject is not M , and forms a belief as to whether she is M on introspective grounds she has already available, she knows by introspection that she is not M .

It is natural also to try to pursue the *opposite* strategy, and, instead of *letting the restriction do also the work that the positional modality does, let the modality do also the work that the restriction to epistemically flawless subjects does*, so that we get the principle of *possible introspection of mind*:

(PIM) Every mental property M is such that:

(PIM^P) Metaphysically necessarily, if one is M one can know by introspection that one is M ;

(PIM^N) Metaphysically necessarily, if one is not M one can know by introspection that one is not M

(where 'can' expresses the relevant "*flawlessifying*"-cum-positional modality, and where 'one' is understood as not restricted to epistemically flawless subjects). Are (RIM) and (PIM) for all intents and purposes equivalent, or are there reasons for preferring one over the other? I'd like to present here two reasons for preferring (RIM) over (PIM).

Firstly, notice that, in some degenerated cases, it would actually not be straightforward to understand the flawlessness modality in (PIM) so that it allows one to overcome every kind of epistemic flaw, as there could be *wacky enough subjects* for which we might not have any *clear conception* of how they could overcome all the epistemic flaws affecting them. (PIM) gets bogged down in all sorts of obscure *de re* flawlessness-cum-positionally modal claims about wacky subjects, whereas (RIM) wisely steers clear of these by only committing to *de dicto* metaphysically modal claims about whoever happens to be a subject of the relevant kind. Here as elsewhere, prevention is better than cure. And, in any event, the *amount of modalisation* required would be such as to make one wonder in exactly what sense the mind is more directly *accessible* than the world: by all natural measures, it might be easier for me to move around the world and discover

all sorts of things about it than for an extremely sophisticated, one-sided and committed eliminativist to recant and come to accept that she does exemplify mental properties.¹⁷

Secondly, notice that it would seem that it is not only “*atomic*” mental properties (like the properties of believing that snow is white, of hoping that grass is green, of loving water etc.) that are introspectable, but also “*compound*” mental properties (like the properties of believing that snow is white and hoping that grass is green, of believing that snow is white or hoping that grass is green, of believing that snow is white if hoping that grass is green etc.). Indeed, (PIM) already explicitly admits that mental properties formed out of mental properties by negation are introspectable, and, if this is so with respect to the Boolean operation of *negation*, it is hard to see how it can fail to be so also with respect to the Boolean operation of *conjunction*: if negating does not take us beyond the introspectable, then surely neither does conjoining. But then (PIM) (and, for that matter, (IM)) falls prey to a version of the *Church-Fitch paradox of knowability* (Fitch [1963]). Take the compound mental property *C* of hoping that grass is green and not believing that one hopes that grass is green. Suppose for *reductio* that one knows that one is *C*. By *distribution of knowledge over conjunction*, one knows [that one hopes that grass is green] and one knows that one does not believe that one hopes that grass is green. From the former piece of knowledge it follows, by the principle that *knowledge requires belief*, that one believes that one hopes that grass is green; from the latter piece of knowledge it follows, by the principle of *factivity of knowledge*:

(F^K) If one knows that *P*, then *P*,

that one does not believe that one hopes that grass is green. Contradiction. Therefore, by *reductio* it is not the case that one knows that one is *C*, and so, by *metaphysical necessitation*, it is metaphysically impossible that one knows that one is *C*. *Contraposing* on (PIM), it follows that it is not the case that one is *C*. But that means that, if one hopes that grass is green, one believes that one hopes that grass is green, which is too strong (at least barring something like the adventurous argument in the tenth last paragraph). Therefore, if we want to stick to (PIM), we’re faced with the exciting prospect of developing an account of exactly *which* compound mental properties are covered by (PIM) (ditto for (IM)). Fortunately, (RIM) does not fall prey to this version of the Church-Fitch paradox of knowability: *contraposing* on (RIM), the fact that it is metaphysically impossible that one knows that one is *C* entails that an epistemically flawless subject *who forms a belief as to whether she is C on introspective grounds she has already available* is such that, if she hopes that grass is green, she believes that she hopes that grass is green, which, by the principle that knowledge requires belief (which is needed anyways for the

¹⁷This reason complements the three reasons I’ve given in Zardini [2015], pp. 384–385 for preferring a certain principle related to (RIM) ((ICTK)) over a certain principle related to (PIM) ((TFPK)): *mutatis mutandis*, the reason given in the text applies also to the choice between (ICTK) and (TFPK) and, conversely, the three reasons given in that paper apply also to the choice between (RIM) and (PIM). Notice that, in Zardini [2015], I’ve introduced (ICTK) as a principle that, contrary to (TFPK), does not fall prey to a certain problem, and the second reason I’m about to offer in the text develops the same theme with respect to the choice between (RIM) and (PIM).

version in question of the paradox of knowability), is just an obvious and unproblematic consequence of the instance of (RIM) for the property of hoping that grass is green (assuming, extremely plausibly, that forming a belief as to whether one is C on introspective grounds one has already available implies forming a belief as to whether one hopes that grass is green on introspective grounds one has already available). Having noted all this about (RIM) and (PIM) (and (IM)), for simplicity's sake I'll mostly conduct the rest of my discussion in terms of (IM) understood as restricted to epistemically flawless subjects.

On the other hand (of the first “main hand”), Williamson offers the (infamous) *systematic argument* against *luminosity* to the effect that, for virtually every fact whatsoever, and so *a fortiori* for virtually every fact involving one's mental properties, it is possible that [that fact obtains, but one is not in a position to know that that fact obtains] (Williamson [1995]; Williamson [1996]; Williamson [2000], pp. 93–113). I've already explained in detail in Zardini [2012] why I think, more negatively, that the anti-luminosity argument is irremediably *fallacious* and, more positively, that a domain of facts *can* be luminous, and that it can be so *without violating the plausible principle that knowledge requires reliability* which also constitutes the starting point of the anti-luminosity argument. Without giving you any spoiler, let me only note here that the shift to (RIM) is also useful in that it makes it easier to see where the weak point of the anti-luminosity argument lies on which luminosity can thrive (see especially Zardini [2012], pp. 404–408).¹⁸

Loosely related to this, it might be worth mentioning that Williamson also puts

¹⁸Srinivasan [2015], essentially followed by McGlynn [2014], pp. 155–160, argues that there is no domain of facts that is luminous for “us” relying (partly) on the allegedly compelling *psychological* claim that, for example, for a being like us who, when slightly above the threshold for feeling cold, believes that she feels cold, there is an easy possibility in which she does not feel cold and yet still believes that she feels cold. I don't see good reasons in favour of such an implausible claim (Srinivasan only offers as justification for it an extremely general and vague principle, which could just as naturally be used to derive, for example, that, for a being like us who, when considering 49, believes that she's considering a number smaller than 50, there is an easy possibility in which she's considering 50 and yet still believes that she's considering a number smaller than 50). Be that as it may, what is most relevant for the purposes of this paper is that, because of her psychological assumption, Srinivasan in effect focusses exclusively on *subjects with certain psychological biases*. The original anti-luminosity argument thus turns in her hands into an argument to the conclusion that *epistemically flawed subjects*—in particular, *subjects who sometimes believe [that they feel cold] whether or not they feel cold*—are not always in a position to know [that they feel cold] when they feel cold. That is indeed a very plausible conclusion, although, as Srinivasan pp. 305–306 herself basically brings out, not one for whose establishment *anything as elaborated* as the original anti-luminosity argument is needed (in fact, at this point the conclusion is much better established otherwise, for the original anti-luminosity argument only establishes its conclusion for a weird scenario in which one is constantly thinking from dawn to noon whether one feels cold, a scenario which is not very relevant *once the discussion is constrained by psychologically realistic assumptions*). But it is neither a conclusion that threatens (IM) nor the conclusion of the original anti-luminosity argument, whose interest largely resided (for me at least, see Zardini [2012], pp. 376–377) in the fact that *it made no assumptions of psychological biases or of other epistemic flaws*. There might be *psychological obstacles* (i.e. obstacles stemming from *one's dispositions to believe and the like*) for beings like us to know even certain obvious facts (for example, the fact that one's baby is cute), but this falls very much short of there being *epistemological obstacles* (i.e. obstacles relating to *the quality of reasons available to one*, independently of what psychological dispositions to believe and the like one happens to have) to such knowledge. In epistemology as elsewhere, psychologism should be resisted.

forth a systematic argument against *KK*, and so *a fortiori* against (IM^P) for knowledge (Williamson [1992]; Williamson [2000], pp. 114–134). I’ve already explained in detail in Zardini [2016b] why I think, more negatively, that the anti-*KK* argument is irremediably *fallacious* and, more positively, that *KK* *can* hold, and that it can do so *without violating the plausible principle that knowledge requires reliability* which also constitutes the starting point of the anti-*KK* argument. Without giving you any spoiler, let me only note here that the shift to (RIM) is also useful in that it makes it easier to see where the weak point of the anti-*KK* argument lies on which *KK* can thrive. Therefore, as in the case of the considerations from common-or-garden examples also in the case of considerations from systematic arguments Williamson’s attack against (IM) for paradigmatic mental properties *not only fails, but is also anyways uncalled for* in that it is motivated by what I’ve argued is the unwarranted assumption that, in some *relevant* situations (i.e. situations of *epistemic flawlessness*), one is not in a position to know whether one knows.¹⁹

On the other (“main”) hand, Williamson tries to *play up* the extent to which knowledge satisfies (IM), claiming that it is (henceforth adopting his own wording) “*characteristically*” if not “*perfectly*” accessible to introspection, so that (IM) for knowledge holds in characteristic situations (Williamson [2000], p. 25). Strangely, he claims this without providing any supporting reason, apparently taking the claim to be too obvious to need any defence. And, in fact, for what it’s worth I can report that, on my first reading of that passage, the claim did strike me as obvious enough (especially given the background belief I’ve defended in this section that, in every relevant situation, one is in a position to know whether one knows). Yet, on reflection I soon discovered what seems to me a compelling argument to the surprising effect that, in a sense that I’ll specify in the next period, *knowledge is systematically inaccessible to introspection*, and it is essentially this argument that I’d like to develop in the rest of this paper. Therefore, while I’ve argued that Williamson has failed to provide reasons for thinking that there is any relevant situation in which one is not in a position to know *by some method or other* whether one knows I’ll attack the broadly related but, in at least one respect, importantly stronger claim that, in characteristic (and relevant) situations, one is in a position to know, or at least can in principle know, *by introspection* whether one knows, contending that, quite on the contrary, in characteristic (and relevant) situations one is *not* in a position to know, and *cannot* even in principle know, *by introspection* whether one knows. While knowledge might be systematically accessible to some method or other, it is in this sense systematically inaccessible to introspection.

¹⁹I should add that, in my own view as developed in Zardini [2016b], the phenomenon of *inexact knowledge* on which the argument mentioned in this paragraph focusses does provide reasons for thinking that, in some relevant situations, (not so much [one *knows*, but one is not in a position to know that one knows] but) [one *does not know*, but one is not in a position to know that one does not know]. However, given that the point is both moot and irrelevant for the main purposes of this paper, having noted that I’d ultimately take it as a telling difference between paradigmatic mental properties and knowledge I’ll henceforth ignore it.

3 Knowledge Is Systematically Inaccessible to Introspection

If (IM) for knowledge holds in characteristic situations, let's suppose without loss of generality that it holds for one's knowledge that snow is white. Let's suppose then that one does know that snow is white. By the relevant instance of (IM^P), one is in a position to know by introspection that one knows that snow is white. Let's suppose without loss of generality that one in effect does form in the appropriate way the belief that one knows that snow is white, so that one knows by introspection that one knows that snow is white. Moreover, by way of ostensive definition let 'a priori *reflection*' cover whatever method (or methods) we use when we know about logic, mathematics, philosophy etc. in the typical "a priori reflective" way we're all aware of. Then, by a priori reflection one also knows (F^K), and so in particular one knows that, if one knows that snow is white, snow is white. Introspection and a priori reflection are clearly different kinds of methods, and yet they arguably *share the important feature of being non-perceptual* (in the sense explained in section 2). Just as with introspection, unsurprisingly it is not wholly uncontroversial that a priori reflection is non-perceptual (Williamson himself can be seen as contesting non-perceptuality, see e.g. Williamson [2013]), but it is nevertheless an extremely plausible assumption that I'll henceforth make.

Now, it is widely thought that a principle of *closure of knowledge under logical consequence* along the lines of:

- (C) If one knows that P_0 , that P_1 , that $P_2 \dots$, and one knows that the proposition that Q logically follows from the propositions that P_0 , that P_1 , that $P_2 \dots$, one is in a position to know that Q

must be essentially correct: it is indeed subject to a lot of counterexamples, but these are typically felt to concern boring issues of fine-tuning that do not detract from (C)'s *essential correctness*, let alone from the idea behind it that *competent deduction transmits knowledge*.²⁰ Let's assume that this is indeed so, so that, for the purposes of this paper, (C) holds. However, if it is granted that (C) holds it's hard to see what objection there could be to a principle of *closure of non-perceptual knowledge under logical consequence* along the lines of:

- (C^{NP}) If one non-perceptually knows that P_0 , that P_1 , that $P_2 \dots$, and one non-perceptually knows that the proposition that Q logically follows from the propositions that P_0 , that P_1 , that $P_2 \dots$, one is in a position to know non-perceptually that Q .

²⁰I myself would make an exception here at least for the kind of counterexample exemplified e.g. by the *preface paradox* (Makinson [1965]), which would seem to show that (C) *embodies an essential mistake about the relations between knowledge and logical consequence*. However, notice that the features generating that kind of counterexample would seem to be *absent* in the situation I'm describing in the text, so that this otherwise important issue would seem to be apt to be legitimately ignored for the purposes of this paper. Notice also that the argument I'm constructing in the text concerns a *single* entailment *with only two premises*, and that hence the corresponding instances of (C) fall under the defence I've given in Zardini [2016a] of closure principles.

It is true that principles broadly related to (C^{NP}) obviously fail: for example, the principle got from (C^{NP}) by substituting ‘immediately’ for ‘non-perceptually’ obviously fails. Still, it is also obvious why such variations fail: the idea behind (C) does not allow to rule out—indeed, it strongly suggests to accept—that *knowledge of the conclusion epistemically depends on knowledge of the premises and on competent deduction from the premises to the conclusion*, so that no variation on (C) is likely to hold if it implies that knowledge of the conclusion does not epistemically depend on those. But in this respect (C^{NP}) itself is presumably safe: its consequent only rules out that knowledge of the conclusion epistemically depends on perception, and its antecedent precisely rules out that knowledge of the premises and knowledge of the entailment, and so presumably competent deduction from the premises to the conclusion, epistemically depend on perception.^{21,22}

Let’s suppose then that, by *a priori* reflection, one also knows (*pace* Williamson’s beloved McGee [1985] reference) that *modus ponens* is valid, and so in particular that one knows that the proposition that snow is white logically follows from the propositions [that one knows that snow is white] and [that, if one knows that snow is white, snow is white]. Since both knowledge by introspection and knowledge by *a priori* reflection are non-perceptual, it follows by (C^{NP}) that one is in a position to know non-perceptually that snow is white. But this conclusion is blatantly absurd: it just ain’t possible to know non-perceptually the colour of a substance.^{23,24} Therefore, keeping fixed the rest of the very plausible principles appealed to in this argument it follows that the relevant instance of (IM^P) is false.²⁵ Given that an analogous argument goes through for the vast majority

²¹After conceiving the main argument of this section, I came across Neta and Pritchard [2007], who consider a problem broadly similar to the one I’m developing in this section which affects what they take to be John McDowell’s position, and offer McDowell a resolution effectively involving a heroic rejection of (C^{NP}). Neta and Pritchard argue against (C^{NP}) by in effect assuming the principle that, if one competently deduces that *Q* from (among other things) a premise to the effect that one has perceptual justification for believing that *P*, then, assuming everything goes well, one thereby knows *only perceptually* that *Q*. But that principle smacks of a sort of *level confusion* and is anyways *independently very implausible*, as it implies, for example, that, if one competently deduces that one has justification of *some kind or other* for believing that things are thus-and-so from (among other things) a premise to the effect that one has *perceptual* justification for believing that things are thus-and-so, then, assuming everything goes well, one thereby knows *only perceptually* [that one has justification of some kind or other for believing that things are thus-and-so], whereas, on the view pushed by Neta and Pritchard, one knows on the contrary *non-perceptually* that one has perceptual justification for believing that things are thus-and-so!

²²Thanks to Roy Cook and Crispin Wright for pushing me to be clear about this issue.

²³And it is if possible even more absurd that one non-perceptually knows the colour of a substance *with the high degree of certainty* which would be conferred by (the certainty of (F^K) and *modus ponens* plus) the very high degree of certainty with which we typically know about our mental properties—indeed, *with the authority* which would be conferred by (the certainty of (F^K) and *modus ponens* plus) the authority with which we typically know about our mental properties.

²⁴The argument focusses on the *factivity* consequence of knowledge only for simplicity’s sake: clearly, analogous arguments can be run which focus on *stronger a priori* consequences of knowledge, like, taking as example knowledge that snow is white, the consequence that one is not related in a Gettier-like way to the fact that snow is white. Thanks to Sven Rosenkranz for comments that provoked this point.

²⁵The argument bears an obvious relationship with McKinsey [1991]’s seminal argument concerning the *apparent conflict between (IM) for mental properties and content externalism* (see McKinsey [2002],

of pieces of perceptual knowledge,²⁶ it follows that *(IM^P) and so (IM) for knowledge systematically fail.*²⁷ Knowledge is systematically inaccessible to introspection.

One might object that the claim in the last paragraph to the effect that it is absurd that one is in a position to know non-perceptually the colour of a substance is *false*, since, to take two standard examples, one can *remember* that snow is white *while sitting in the armchair* or *be told* that snow is white *while sitting in the armchair*.²⁸ In reply, I stress that

pp. 217–218, n. 1 about the history of his argument). However, *none of the approaches to that apparent conflict lends much solace to the claim that (IM) for knowledge characteristically holds*. On one approach (see e.g. McLaughlin and Tye [1998]), although one can know by introspection that one believes that snow is white one cannot know non-perceptually that belief that snow is white requires that the world exemplify certain non-trivial properties. But the analogue of that approach in our context would involve the very improbable claim that one cannot know non-perceptually (F^K). On another approach (see e.g. Wright [2000]), one cannot know non-perceptually *via* the McKinsey argument that the world exemplifies certain non-trivial properties because one actually non-perceptually warrantedly presupposes that at least in order for one both to know by introspection that one believes that snow is white and to know by *a priori* reflection that content externalism holds (so that one cannot use the McKinsey argument to transmit one’s knowledge of its premises to its conclusion). But the analogue of that approach in our context would involve the very improbable claim that one can non-perceptually warrantedly presuppose that snow is white. On yet another approach (see e.g. Warfield [1998]), which I favour, one can know non-perceptually *via* the McKinsey argument that the world exemplifies certain very abstract non-trivial properties (like, on my view, the property of containing some sort of substance or other that is snow) on the basis of one’s knowledge by introspection that one believes that snow is white and of one’s knowledge by *a priori* reflection that content externalism holds. But the analogue of that approach in our context would involve the improbable claim that one can know non-perceptually that the world exemplifies certain very concrete non-trivial properties (like the property of containing a white substance that is snow). On one final approach (see e.g. McKinsey [2002]), although one can know by *a priori* reflection that belief that snow is white requires that the world exemplify certain non-trivial properties one cannot know by introspection that one believes that snow is white. But the analogue of that approach in our context would involve precisely the claim that (IM) for knowledge fails. In sum, contrary to the impression one might get from Williamson [2000] there are substantial differences between the *world involvement of properties subject to content externalism* and the *world involvement of properties subject to factivity*, differences that might actually be *crucial for mentality*. Thanks to Aurélien Darbellay and Crispin Wright for discussions of these issues.

²⁶Similar arguments actually go through also for many pieces of *non-perceptual* knowledge. For example, a variation on (C^{NP}) that, essentially for the reason given in the last paragraph, is equally safe says that, if one knows by introspection that P_0 and knows by *a priori* philosophical reflection that P_1 , and one knows by *a priori* logical reflection that the proposition that Q logically follows from the propositions that P_0 and that P_1 , one is in a position to know by introspection, *a priori* philosophical and *a priori* logical reflection alone that Q . But then we can run an argument similar to the one in the text to the absurd conclusion that one is in a position to know by introspection, *a priori* philosophical and *a priori* logical reflection alone that Fermat’s Last Theorem is true, where such knowledge does not epistemically depend on any mathematical consideration.

²⁷Indeed, to avoid distracting lines of defence against the main argument of this section notice that the version of (IM) mentioned in section 2 in which ‘*is in a position to know*’ is replaced by the *much weaker* ‘*can in principle know*’ presumably holds for mental properties whereas, *mutatis mutandis*, the main argument of this section also establishes, as I’ve already flagged later in section 2, that *even such version of (IM) systematically fails for knowledge*.

²⁸One could also point to the *broadly Cartesian project of grounding knowledge of the world in non-perceptual knowledge (in particular, in knowledge by introspection and knowledge by a priori reflection)*. As far as I can tell, philosophers have *in fact not* yet been able to carry out that project successfully, so I’d actually take that case as further evidence for the absurdity claim in the last paragraph. Moreover,

non-perceptuality is an *epistemic* property concerning epistemic dependence and that the relevant notion of epistemic dependence is suitably *inter-temporal* and *inter-subjective*, so that knowledge based on (episodic) memory epistemically depends on the past perception the memory is linked to and so that knowledge based on testimony epistemically depends on the other subject’s assertion the testimony is linked to.^{29,30}

One might object that the absurdity claim in the second last paragraph is (maybe not false but anyways) *unwarranted* in that it *relies on the general presupposition* that, if the “canonical” method for knowing a proposition is of a certain kind, the proposition *cannot be known by any other method*,³¹ and then object that *that presupposition is false*, since, for example, one can know that either snow is white or it is not the case that snow is white by seeing that snow is white and then inferring by addition that either snow is white or it is not the case that snow is white (rather than by the canonical method of *a priori* reflection establishing it as a logical truth). In reply, I deny that the absurdity claim relies on *such a general* presupposition: the claim that one cannot know non-perceptually the colour of a substance is perfectly consistent with the claim that one can know perceptually things that one would canonically know by *a priori* reflection—the fact that sometimes one needs to go to the field is perfectly consistent with the fact some other times one needn’t or even shouldn’t.³²

One might object that the absurdity claim in the third last paragraph is (maybe neither false nor unwarranted but anyways) *ambiguous* and that, under one reading (compatible with (IM) for knowledge characteristically holding), it is true and, under another reading (incompatible with (IM) for knowledge characteristically holding), it is false. To wit, one could first suggest that, under one (true) reading, the absurdity claim means that it is absurd that one is in a position to know the colour of a substance *without ever leaving the armchair* and, under another (false) reading, the absurdity claim means that it is absurd that one is in a position to know non-perceptually the colour of a substance *full stop*. One could then observe that the main argument of this section only shows that, if (IM) for

that project is characterised by the search of *non-easy mind-to-world bridge principles*, which belies the project’s underlying assumption that *no easy mind-to-world bridge principle is available*, which is in turn yet further evidence for the claim that it is absurd that one is in a position to know non-perceptually the colour of a substance *in the easy way envisaged by the (F^K)-based argument*. For what it’s worth, I myself would reject that justification of easy mind-to-world bridge principles (such as e.g. a principle going from its perceptually looking to one as though *P* to its being the case that *P*) is non-perceptual (Zardini [2014]), and I’d accept that knowledge of some very abstract non-easy mind-to-world bridge principles is non-perceptual (fn 25).

²⁹Cases in which one is still justified but there is actually no such past perception or no such other subject’s assertion are accommodated by the fact that the relevant notion of epistemic dependence is suitably *intensional* (as already intimated by its inter-temporal character): a justification can epistemically depend on a perception, or on an assertion, even if there are in fact no relevant perceptions, or relevant assertions (just like someone can desire a perfect partner even if there are in fact no relevant perfect partners).

³⁰Thanks to Célia Teixeira for discussion of this objection.

³¹One could mention instead the related but different presupposition that, if one can know something perceptually, one cannot know it non-perceptually. *Mutatis mutandis*, everything to follow in the text would apply to this case too.

³²Thanks to Sònia Roca for discussion of this objection.

knowledge characteristically holds, one is indeed in a position to know non-perceptually that snow is white, but that is *compatible with such knowledge always presupposing an antecedently existing piece of perceptual knowledge* that snow is white (for introspection does seem to presuppose the antecedent existence of the state which it introspects), so that the claim that (IM) for knowledge characteristically holds does not imply that one is in a position to know that snow is white without ever leaving the armchair. And one could finally object that the apparent force of the main argument of this section vanishes once these two readings are distinguished. In reply, I observe that, in the relevant respects, the absurdity claim is obviously *unambiguous*: ‘It is absurd that one is in a position to know non-perceptually the colour of a substance’ unambiguously means that it is absurd that, in some way or other, one is in a position to know non-perceptually the colour of a substance, and that is at worst only ever so slightly less compelling than the claim that it is absurd that one is in a position to know non-perceptually the colour of a substance without ever leaving the armchair. I submit that the objection might seem to some appealing because of its *confusion between epistemic dependence and causal dependence* (in a broad sense of ‘causal’ encompassing the relation between knowledge and what is known). But, once the two kinds of dependence are distinguished, the absurdity of the claim that some piece of knowledge that snow is white does not *epistemically* depend on perception would not seem to be to almost any extent lessened by the observation that it is still the case that every such piece of knowledge *causally* depends on perception (and indeed on perceptual knowledge of the same proposition). Compare: it might be that one needs to have had perceptions, say, as of certain objects being white (and indeed have had perceptual knowledge, say, that snow is not white all over and black all over) in order for one to possess the concept of white,³³ and so in order for one then standardly to entertain the proposition that snow is not white all over and black all over. If so, one’s subsequent knowledge by *a priori* reflection that snow is not white all over and black all over causally depends on such perceptions (and indeed on the antecedently existing piece of perceptual knowledge that snow is not white all over and black all over), but such causal dependence would not seem to detract in the least from the idea that, with the subsequent piece of knowledge, in virtually every epistemologically interesting sense one knows from the armchair that snow is not white all over and black all over. Just so, the fact that one’s non-perceptual knowledge that snow is white causally depends on the relevant perceptions (and indeed on the antecedently existing piece of perceptual knowledge that snow is white) would not seem to detract in the least from the idea that, with the subsequent piece of knowledge, in virtually every epistemologically interesting sense one knows from the armchair that snow is white. Alas, in virtually every epistemologically interesting sense one cannot know from the armchair the colour of a substance.³⁴

One might object that the absurdity claim in the fourth last paragraph is (maybe neither false nor unwarranted nor ambiguous but anyways) *harmless* since, even granting that the main argument of this section is sound and hence that (IM) for knowledge

³³I certainly wouldn’t recommend this specific view, but its sheer coherence is already enough to make the conceptual point I want to make.

³⁴Thanks to Colin Caret, Célia Teixeira and David Yates for discussion of this objection.

systematically fails, it might still be the case that, metaphysically necessarily, if one knows that P by method m one is in a position to know by introspection *plus* m that one knows that P . In reply, I don't see that such *Ersatz* principle vindicates any interesting sense in which knowledge patterns in terms of accessibility with paradigmatic mental properties *as opposed to paradigmatic non-mental relations*. For example, since one knows by perception that snow is white, according to the *Ersatz* principle, one is in a position to know by introspection plus perception that one knows that snow is white. The problem is that one would also seem to be in a position to know by introspection plus perception that one truly believes that snow is white, but truly believing that snow is white is a paradigmatic non-mental relation.³⁵

4 No Property Non-Perceptually Known to Be Factive Is Mental, but Seeing Is Not One of Them

The conclusion of the main argument of section 3 is that *knowledge is systematically inaccessible to introspection*.³⁶ Given the considerations in section 2, I'll henceforth assume (with Williamson) that mental properties are not systematically inaccessible to introspection, so that the conclusion of the main argument of section 3 implies that *knowledge is not a mental property*. It should now be made clear that this argument actually has *much wider implications*. For notice that the argument only relies, in the relevant respect, on the assumption that knowledge is *non-perceptually known to be factive*. Therefore, if good, the argument actually establishes that *no property non-perceptually known to be factive is mental*.³⁷

Now, that implication will strike many epistemologists—especially those adhering to the recent “*factivity turn*”—as false. To take a trite example, it is commonplace both to

³⁵Deploying a point emerged in fn 24, I should also mention that, on plausible assumptions, the *Ersatz* principle actually succumbs just as well to the main argument of this section. For example, one plausibly knows that there is a barn in front simply by perception as of a barn in front, and one plausibly knows by *a priori* reflection that, if one knows that there is a barn in front, there are not many fake barns around. Assuming that it is in the spirit of the *Ersatz* principle that in such a case the method can be taken to be the relatively fine-grained one of perception as of a barn in front, the principle would then absurdly deliver that one is in a position to know by introspection plus perception as of a barn in front that there are not many fake barns around. Thanks to Sven Rosenkranz for discussion of this objection.

³⁶That conclusion in itself still leaves open a variety of *non-introspection-based epistemologies of knowledge*. For what it's worth, my independently motivated approach has it, very roughly, that, *under normal conditions, justification for believing that P is also justification for believing that one knows that P* (Zardini [2016b]); if that is the method we use when we know that we know, it makes *knowledge of knowledge systematically perceptual rather than introspective*, as the main argument of section 3 shows it'd better be.

³⁷In fact, the other generalisation of the argument adumbrated in fn 26 indicates that its really fundamental assumption is simply, a bit roughly, that *the relevant propositions can only be known in a way which need not be involved in introspection and in knowing that the relevant property is factive*. Of course, generalisations of the argument in other directions might also be worth investigating (for example, substituting other epistemic properties for knowledge in the “outermost” occurrences of knowledge in the argument, as adumbrated in fn 23), but in this paper I'll leave it at that.

maintain the principle of *factivity of seeing*:

(F^S) If one sees that P , then P

and at the same time hold that (propositional) seeing (such as seeing that snow is white) is a mental property. (Ditto for another central property like (propositional) remembering (such as remembering that snow is white); to fix ideas, I'll henceforth focus on seeing, but I think that everything I say applies to remembering as well.)³⁸ Many epistemologists will then probably conclude that there must be something wrong with the main argument that has so far been developed. For the purposes of this paper, that is not however an option of much interest.

For the purposes of this paper, of some more interest is the option of taking the main argument that has so far been developed to apply also to seeing, yielding the conclusion that also seeing is not a mental property. Presumably, one could then sweeten the pill by adding that, although *seeing that P* is not a mental property, something along the lines of *its visually looking as though P* is. To me, these two properties sound however pretty much equivalent, and so I think that, for the purposes of this paper, of even (much) more interest is the option of taking the main argument that has so far been developed also to speak in favour of the idea that (F^S) *fails*.

Indeed, it seems to me fairly clear that, on reflection, (F^S) is more like a *presupposition typically in place in context* rather than an *unrestrictedly true principle* (see Rimoldi [2014], pp. 181–191 for a broadly congenial discussion). This is to me evidenced by the fact that the presumption of (F^S) seems to me in a few cases *suspended*, especially (and tellingly, as far its presupposition-like status is concerned) when the relevant seeing occurs in *other subjects, times or situations than those of utterance*, as witnessed by:

- (1) Jim saw that the ball was moving, whereas Jon saw that it was still;
- (2) Although Jim saw that the ball was moving, it actually was still;
- (3) Jim saw that the ball was moving even if it actually was still;
- (4) Jim continued to see that the ball was moving even after it had stopped;
- (5) It might happen that Jim sees that the ball is moving whereas it is actually still;
- (6) If Jim sees that the ball is moving, he is wrong;
- (7) If the ball were still but Jim looked at it under unfavourable conditions, he would see that it is moving;

³⁸Williamson [2000], pp. 37–39 follows suit. Such combination of theses notoriously lands him into trouble with his other thesis that *knowledge is the most general factive mental state*, for the three theses together entail that, *if one sees that P , one knows that P* , which, given a plausible *reliability constraint on knowledge* that Williamson would also seem to accept, in turn implies (in my view, absurdly) that someone who is looking at a barn might fail to see that there is a barn in front only because there is an unseen fake barn behind.

- (8) The ball is still. Does Jim see that it is moving?
- (9) Jim saw that the ball was moving, but I'm not sure whether it was;
- (10) Jim hallucinated and thereby saw that the ball was moving;
- (11) When Jim saw that the ball was moving he was under drugs;
- (12) What Jim saw was that the ball was moving

(if it helps, add to the background that Jim and Jon have just smoked some pot). Notice that the *consistency* of (1)–(5), the *consistency* of (6)–(7) *with their antecedents* and the *felicity* of (8) more or less directly prove the point; the consistency of (9)–(12) does it more indirectly in the sense that it would seem to extend to their *consistency with its not being the case that the ball was moving*.³⁹ Notice also that neither of (1)–(5) is consistent, neither of (6)–(7) is consistent with its antecedent and (8) is not felicitous *if 'know' and its relatives are substituted for 'see' and its relatives*, and that, under the same substitution, none of (9)–(12) is consistent with its not being the case that the ball was moving.⁴⁰ Notice finally that, in these respects, 'can see' would *prima facie* seem to pattern with 'know' rather than 'see' (see the discussion in Vendler [1957], pp. 148–149, 156–157). But, arguably, not even in this case do we really have a seeing-related factive property. For, arguably, *what should be seen* (which is what is seen when it “can be seen”) *is determined by context* and, while many, “*investigative*” contexts do determine that what should be seen is *something that is the case*, some other, “*recreational*” contexts determine that what should be seen is *something else*: for example, after enough pot smoking at the party, and after a few smokers reporting their surprisingly converging experiences as of the local mountain moving, ‘Jim too can now see that the mountain is moving’ can truthfully describe Jim’s own hallucinatory state.^{41,42}

³⁹My judgements about (1)–(12) do not seem to be universally accepted among English speakers. I regard this as an issue about English ripe for further investigation. My own sense is that the presumption of (F^S) might indeed be usually stronger in English than it is in quite a few other languages (some sundry examples: Arabic, Hindi and Italian) in which it is easily suspended. In fact, I wouldn't wish to rule out that, in some language *L*, (F^S) is indeed an unrestrictedly true principle for a verb standardly used in *L* for seeing. But, letting *schsee* be such verb, I'd insist that the situation should be described as one in which *schsee* still expresses the property of *seeing*—rather than an extensionally different property of “schseeing”—and *L* simply somehow places a *further necessary condition for the relevant sentences* in which *schsee* occurs to be *true*.

⁴⁰More generally, I'm not aware of convincing counterexamples to (F^K). Hazlett [2010] offers some candidates, but they all strike me as unconvincing. For example, ‘Everyone knew that stress caused ulcer before it was proved that ulcers are actually caused by bacterial infection’ does sound like an OK thing to say, but so does ‘Tobacco was a harmless treat before its negative health effects were proved’.

⁴¹The presumption of (F^S) can be suspended not only in *ordinary* discourse, but also in *philosophical* discourse: “[...] the subject can see that the stone (for example) is moving and that it is not moving” (Crane [1988], p. 145; in the same debate, making the issue salient Mellor [1988], p. 149 candidly observes that “[i]t isn't absurd to say ‘I see that *a* is *F* but it isn't’ [...]”).

⁴²The presumption of (F^S) can be suspended not only in *ordinary* discourse, but also in *scientific* discourse: “[...] you may also see that these imaginary dots “link up” to form a series of circles that radiate out from the middle of the illusion” (Joyce et al. [2000], p. 81). More interestingly, to the best of

A slightly more indirect argument against (F^S) starts with the extremely plausible principles:

(ANNIA) If one sees a *ADJECTIVE NOUN*, one sees that a *NOUN* is *ADJECTIVE*;

(DNTDN) If one sees *DETERMINER*^{WEAK} *NOUN*s, one sees that there are *DETERMINER*^{WEAK} *NOUN*s;

(NVNIV) If one sees a *NOUN VERB*ing, one sees that a *NOUN* is *VERB*ing,

and continues with the unassailable observations that, suitably deceived, one sees a pink elephant; that, suitably deceived, one sees two unicorns; that, suitably deceived, one sees an elephant flying. From the first observation and (ANNIA), it follows that one sees that an elephant is pink; from the second observation and (DNTDN), it follows that one sees that there are two unicorns; from the third observation and (NVNIV), it follows that one sees that an elephant is flying. Since however no elephant is pink, there are not two unicorns and no elephant is flying, (F^S) fails. Notice that objectual seeing (such as seeing snow) seems to allow at least two different readings: on a “transparent” reading, if one sees a pink elephant there is a pink elephant *x* and one sees *x* (possibly without having a visual experience as of a pink elephant); on an “opaque” reading, if one sees a pink elephant one has a visual experience as of a pink elephant (possibly without there being a pink elephant). It is the opaque reading that is understood as operative in (ANNIA) and (DNTDN).⁴³ Notice also that there is *no full analogue of the opaque reading for objectual knowledge*: on every possible reading, if one knows a pink elephant there must be a pink elephant.⁴⁴

A predictable reply to arguments against (F^S) is to claim that there must be something wrong with them, given that, when one says ‘I see that the ball is moving’ but it is then shown that the ball was not moving, one would naturally react by saying something along the lines of ‘OK, I thought I saw that the ball was moving, but I was wrong’. But that can hardly be taken as a ground for believing that one did not see that the ball was moving. Compare: when one says ‘I have good reason for thinking that the butler is the murderer’ but it is then shown that the butler is not the murderer, one would naturally react by saying something along the lines of ‘OK, I thought I had good reasons for thinking that

my knowledge current *psychology of perception* does not assume, and does not provide any good reasons to assume, that, beyond the kind of psychological state constituted by usual *fallible visual representation* of objects, properties and states-of-affairs, there is a further kind of psychological state constituted by some sort of *infallible visual representation* of objects, properties or states-of-affairs (the contrast with the situation for content externalism is sharp, see e.g. Burge [2005]).

⁴³Even so, one might still want to place some systematic restrictions on *ADJECTIVE*, so as to rule out e.g. ‘fake’ as an admissible value.

⁴⁴A final argument against (F^S) is from authority: “And the woman saw that the tree was good for food” (*Genesis* 3, 6; “*wattēre’ hā’iššāh ki tōv hā’ēš l’ma’ākhol*”; *Septuagint*: “*kai eiden hē gynē hoti kalon to xylon eis brōsin*”; *Vulgate*: “*Vidit igitur mulier quod bonum esset lignum ad vescendum*”; unsurprisingly, the interpretation of the passage is not completely uncontroversial, but it would seem rather dubious to infer from it that the tree was indeed good for food).

the butler is the murderer, but I was wrong’, but that clearly is no ground for believing that one did not have good reasons for thinking that the butler is the murderer. Indeed, turning the tables on the reply under discussion, if the respect in which, in the former situation, one was wrong is that one was wrong in thinking that one saw that the ball was moving,⁴⁵ we’d expect ‘OK, I thought I saw that the ball was moving, but actually I didn’t see that’ to be a similarly natural reaction, but it isn’t (to me, it implies that one didn’t have a visual experience as of the ball moving in the first place).⁴⁶ And we’d also expect a remark like ‘I didn’t see that the ball was moving, as I was hallucinating that the ball was moving’ to sound natural, but it doesn’t. Notice that, on the contrary, in the case of knowledge, when one says ‘I know that the ball is moving’ but it is then shown that the ball was not moving, the reactions ‘OK, I thought I knew that the ball was moving, but I was wrong’, ‘OK, I thought I knew that the ball was moving, but actually I didn’t know that’ and the remark ‘I didn’t know that the ball was moving, as I was hallucinating that the ball was moving’ are equally natural.

5 Mind and Cognition

The investigation in this paper has confirmed the initial answer with which the paper began: knowledge is not a property of the mind, but a relation that the mind bears to the world—in the relevant respect, knowledge that snow is white is more akin to tracking that snow is white than to believing that snow is white. Yet, this by no means implies that, in the relevant respect, knowledge that snow is white is more akin to a radically non-mental relation like causing that snow is white than to a mental property like believing that snow is white. For, while non-mental, knowledge is nevertheless what may be called a “*cognitive*” relation: a relation that requires from its subject possession of the *representational resources* involved in its content and that applies to contents exhibiting the marks of *intensionality* (understanding this as the broad opposite of *extensionality*, not e.g. as the specific dimension between extensionality and *hyperintensionality*).⁴⁷

Let’s first note that knowledge requires from its subject possession of the representa-

⁴⁵In passing, why on earth should one invariably have had back then the weird second-order attitude of thinking that one sees? This is further evidence that the reactions in question must be handled very carefully.

⁴⁶Reactions that on the surface might seem pretty much equivalent are in fact crucially different: ‘OK, I thought I saw that the ball was moving, but I didn’t REALLY see that’ (‘OK, I thought I saw that the ball was moving, but I didn’t really SEE that’) is indeed a similarly natural reaction. But then ‘If you came back at midnight, you didn’t REALLY go out last night’ (‘If you came back at midnight, you didn’t really GO OUT last night’) is a similarly natural thing to say.

⁴⁷These two components of a cognitive relation are arguably by and large independent: the relation of *intentionally causing* arguably requires from its subject possession of the representational resources involved in its content, but applies to contents that do not exhibit many of the marks of intensionality (for example, none but the first one of those mentioned in the next paragraph); the relation *between the contents of the antecedent and the consequent of a conditional* arguably applies to contents exhibiting the marks of intensionality (for example, all those mentioned in the next paragraph), but does not involve a subject in the first place.

tional resources involved in its content: if Jim has never heard of gold, ‘Jim knows that the gold melted’ is false. By contrast, even if Jim has never heard of gold ‘Jim caused that the gold melted’ can be true. Let’s then go through a few indicative marks of intensionality to check that the contents of knowledge exhibit them. (Warning: we’ll do so by making some not completely uncontroversial assumptions.) Firstly, the contents of knowledge *do not license intersubstitutability of identicals*: ‘Jim knows that Tully shaved’ can be true without ‘Jim knows that Cicero shaved’ being true. By contrast, ‘Jim caused that Tully shaved’ can only be true if ‘Jim caused that Cicero shaved’ is true. Secondly, the contents of knowledge *allow for non-specific interpretations*: ‘Jim knows that a friend is dying’ can be true even if Jim knows of no particular friend that she is dying. By contrast, ‘Jim caused that a friend died’ can only be true if Jim caused a particular friend to die. Thirdly, the contents of knowledge *allow for reference failure*: ‘Jim knows that Bucephalus rather than Pegasus won the race’ can be true even if there is no such an object as Pegasus. By contrast, ‘Jim caused that Bucephalus rather than Pegasus won the race’ can only be true if there is such an object as Pegasus. Fourthly, the contents of knowledge *can be logical truths in a discriminating way*: ‘Jim knows that either snow is white or it is not the case that snow is white’ can be true or false depending on Jim’s logical acumen. By contrast, ‘Jim caused that either snow is white or it is not the case that snow is white’, if it has a truth value at all, is presumably necessarily false. Notice that virtually all these points about the cognitive character of knowledge also apply *mutatis mutandis* to our examples in section 1 of tracking, measuring (which is understood as a state and for which however something like the fourth mark would not seem to be applicable) and testing (and to many more relations that the mind bears to the world).

Moreover, it is plausible to assume—although it lies beyond the scope of this paper to argue—that knowledge *does not reduce to a mental state having such-and-such features* (be such features added to the state *extensionally*—for example, “*s* knows that *P* reduces to the fact that *s* believes that *P* and...”—or *intensionally*—for example, “*s* knows that *P* reduces to the fact that *s* believes that *P* in virtue of...”⁴⁸). If so, knowledge offers a prime example of a *non-mental, cognitive relation whose cognitive character is not derivative from that of a mental property* (as opposed, for example, to believing truly), and so a prime example of how, in some cases, cognitions just ain’t in the mind.⁴⁹

⁴⁸Many have argued for this assumption by way of *induction from the post-Gettier literature*. Williamson [2000], pp. 65–92 himself also offers a more detailed argument restricted however to a subcase of the extensional case. On my own views, over and above the specific induction from the post-Gettier literature there stands the much more general fact that *most properties and relations expressed by atomic concepts are not in any way reducible* (arguably, even the property of being a bachelor is not reducible to the property of being an unmarried man). Moreover, on my own views *knowledge does not require belief nor any other mental state* (with Radford [1966], the student who ventures ‘1492’ only after being forced to guess when Columbus first landed in America arguably knows that Columbus first landed in America in 1492 even if she does not believe it; against Lackey [1999], the professor converted to nominalism who continues to teach that there are prime numbers arguably still knows that there are prime numbers even if she no longer believes it; beyond Steup [2006], the husband who has a psychological block on catching his wife *in flagrante* arguably already knows that his wife betrays him even if he does not yet believe it), from which the non-reducibility in question follows even more straightforwardly.

⁴⁹Such cognitions will have to fail the relevant versions of (IM) and so will have to fail the relevant

References

- William Alston. Varieties of privileged access. *American Philosophical Quarterly*, 8: 223–241, 1971.
- Tyler Burge. Individualism and self-knowledge. *The Journal of Philosophy*, 85:649–663, 1988.
- Tyler Burge. Disjunctivism and perceptual psychology. *Philosophical Topics*, 33:1–78, 2005.
- Alex Byrne. Introspection. *Philosophical Topics*, 33:79–104, 2005.
- Earl Conee. The comforts of home. *Philosophy and Phenomenological Research*, 70: 444–451, 2005.
- Tim Crane. The waterfall illusion. *Analysis*, 48:142–147, 1988.
- Gareth Evans. *The Varieties of Reference*. Oxford University Press, Oxford, 1982.
- Jordi Fernández. *Transparent Minds: A Study of Self-Knowledge*. Oxford University Press, Oxford, 2013.
- Frederic Fitch. A logical analysis of some value concepts. *The Journal of Symbolic Logic*, 28:135–142, 1963.
- Edward Garrett. *Evidentiality and Assertion in Tibetan*. PhD thesis, University of California, Los Angeles, Los Angeles, 2001.
- Allan Hazlett. The myth of factive verbs. *Philosophy and Phenomenological Research*, 80:497–522, 2010.
- Katherine Joyce, Michael DiSpezio, and Charles Paraquin. *Optical Illusions*. Sterling Publishing, New York, 2000.
- Akio Kamio. Territory of information in English and Japanese psychological utterances. *Journal of Pragmatics*, 24:235–264, 1995.
- Immanuel Kant. *Kritik der reinen Vernunft*. Hartknoch, Riga, 2nd edition, 1787.
- Jennifer Lackey. Testimonial knowledge and transmission. *The Philosophical Quarterly*, 49:471–490, 1999.

versions of the arguments mustered in section 2. In effect, they are cognitions that are not most intimate to a subject (they are not part of the perspective that a subject really occupies; they do not in the relevant sense belong to a subject) and they are cognitions that are only known perceptually (natural language registers no contrast in attributions of them to oneself and to merely observed others; their occurrence requires of a subject perceptual activity to such an extent that she will typically also have available enough perceptual evidence to know with a high degree of certainty that that they occur).

- David Makinson. The paradox of the preface. *Analysis*, 25:205–207, 1965.
- Vann McGee. A counterexample to modus ponens. *The Journal of Philosophy*, 82:462–471, 1985.
- Aidan McGlynn. *Knowledge First?* Palgrave MacMillan, Basingstoke, 2014.
- Michael McKinsey. Anti-individualism and privileged access. *Analysis*, 51:9–16, 1991.
- Michael McKinsey. Forms of externalism and privileged access. *Philosophical Perspectives*, 16:199–224, 2002.
- Brian McLaughlin and Michael Tye. Externalism, twin earth, and self-knowledge. In Crispin Wright, Barry Smith, and Cynthia Macdonald, editors, *Knowing Our Own Minds*, pages 285–320. Clarendon Press, Oxford, 1998.
- Hugh Mellor. Crane’s waterfall illusion. *Analysis*, 48:147–150, 1988.
- Ram Neta and Duncan Pritchard. McDowell and the new evil genius. *Philosophy and Phenomenological Research*, 74:381–396, 2007.
- Colin Radford. Knowledge: By examples. *Analysis*, 27:1–11, 1966.
- Florencia Rimoldi. ¿Puede el conocimiento ser un estado mental? *Análisis Filosófico*, 34: 171–201, 2014.
- Gilbert Ryle. *The Concept of Mind*. Hutchinson, London, 1949.
- Daniele Sgaravatti and Elia Zardini. Knowing how to establish intellectualism. *Grazer Philosophische Studien*, 77:217–261, 2008.
- Amia Srinivasan. Are we luminous? *Philosophy and Phenomenological Research*, 90: 294–319, 2015.
- Jason Stanley and Timothy Williamson. Knowing how. *The Journal of Philosophy*, 98: 411–444, 2001.
- Matthias Steup. The analysis of knowledge, 2006. URL <http://plato.stanford.edu/entries/knowledge-analysis/>.
- Zeno Vendler. Verbs and times. *The Philosophical Review*, 66:143–160, 1957.
- Ted Warfield. A priori knowledge of the world: Knowing the world by knowing our minds. *Philosophical Studies*, 92:127–147, 1998.
- Timothy Williamson. Inexact knowledge. *Mind*, 101:217–242, 1992.
- Timothy Williamson. Does assertibility satisfy the S4 axiom? *Crítica*, 27:3–22, 1995.
- Timothy Williamson. Cognitive homelessness. *The Journal of Philosophy*, 93:554–573, 1996.

- Timothy Williamson. *Knowledge and Its Limits*. Oxford University Press, Oxford, 2000.
- Timothy Williamson. How deep is the distinction between a priori and a posteriori knowledge? In Albert Casullo and Joshua Thurow, editors, *The A Priori in Philosophy*, pages 291–312. Oxford University Press, Oxford, 2013.
- Crispin Wright. Cogency and question-begging: Some reflections on McKinsey’s paradox and Putnam’s proof. *Philosophical Issues*, 10:140–163, 2000.
- Elia Zardini. Luminosity and vagueness. *Dialectica*, 66:375–410, 2012.
- Elia Zardini. Knowledge-how, true indexical belief and action. *Philosophical Studies*, 164: 341–355, 2013a.
- Elia Zardini. Luminosity and determinacy. *Philosophical Studies*, 165:765–786, 2013b.
- Elia Zardini. Confirming the less likely, discovering the unknown. Dogmatism: surd and doubly surd, natural, flat and sharp. In Dylan Dodd and Elia Zardini, editors, *Scepticism and Perceptual Justification*, pages 33–70. Oxford University Press, Oxford, 2014.
- Elia Zardini. Truth, demonstration and knowledge: A classical solution to the paradox of knowability. *Theoria: An International Journal for Theory, History and Foundations of Science*, 30:365–392, 2015.
- Elia Zardini. Closed without boundaries. ms, 2016a.
- Elia Zardini. Open knowledge of one’s inexact knowledge. ms, 2016b.