THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Interpretation and Knowledge Maximization

Citation for published version:

Mcglynn, A 2012, 'Interpretation and Knowledge Maximization', *Philosophical Studies*, vol. 160, no. 3, pp. 391-405. <https://doi.org/10.1007/s11098-011-9725-2>

Digital Object Identifier (DOI):

[10.1007/s11098-011-9725-2](https://doi.org/10.1007/s11098-011-9725-2)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Philosophical Studies

Publisher Rights Statement:

© McGlynn, A. (2012). Interpretation and Knowledge Maximization. *Philosophical Studies*, 160(3), 391-405. [10.1007/s11098-011-9725-2](https://doi.org/10.1007/s11098-011-9725-2)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Interpretation and Knowledge Maximization

1. Introduction and Preliminaries

A recurring complaint in Timothy Williamson's recent work is that contemporary epistemology and philosophy of mind is prone to attempting to let justification and belief do work that must rather be done by knowledge. Williamson has influentially championed a 'knowledge first' approach to these areas of philosophy, the core claims of which are that knowledge is unanalyzable, and that knowledge—rather than justification, warrant, reasonableness, or rationality—is the central and most important epistemic status, and knowledge—rather than belief—is the central and most important mental state with mind-to-world fit. While this knowledge first approach to epistemology and the philosophy of mind defended in Williamson's *Knowledge and its Limits* (2000) has generated a great deal of attention, his attempt to effect a similar reorientation concerning fundamental issues in the philosophy of intentionality has gone rather unnoticed. It is time to correct this.

My concern here will be with Williamson's attempt to give a knowledge first twist to a central component of David Lewis's account of intentionality. Williamson proposes that a principle of charity 'constitutive of content' (2007: 264) should be a principle of knowledge maximization; it should constrain the interpreter to favor interpretations which maximize the number of knowledgeable judgments, both verbalized and un verbalized, the speaker comes out as making. Williamson follows

Lewis in taking the plight of the radical interpreter as dramatizing, not an epistemological question concerning how we determine facts about the contents of the speech acts and attitudes of those being interpreted given prior knowledge of the relevant non-intentional facts, but rather a metaphysical question concerning how such facts about content are determined by the non-intentional facts (Lewis 1974: 333-4, Williamson 2007: 260-1). The principles constraining interpretation, including the principle of charity, will be constitutive of content in the sense that what it is for P to be the content of one's thought is for P to be the content that would be assigned by an interpreter who is omniscient with respect to the non-intentional facts, working constrained by those principles; more generally, it is in virtue of these principles that one's mental states and speech acts have the contents they do. Following now standard terminology, let us call this an *interpretationalist* account of intentional content.

According to this picture, then, it is partly in virtue of a principle of knowledge maximization that our thoughts and utterances have the contents that they do. This deserves to be called a knowledge first approach to foundational issues in the theory of content if anything does. Grappling with Williamson's proposals provides an occasion to reconsider some of these foundational issues, as well as to assess how far Williamson's knowledge first approach to philosophy can be pushed.

Here I will argue against Williamson's proposal. One strategy for doing so would be to argue against interpretationalist accounts of content as a family. A second, to be

pursued here, instead argues specifically against Williamson's proposed modification. Adopting this second strategy may suggest that I'm broadly sympathetic to interpretationalism, and want only to call into question whether it stands in need of the kind of modification Williamson proposes. This isn't so. My reason for not engaging with interpretationalism in general is not that I think it's in perfectly good standing as it is, but rather that doing so would be to engage the question of the proper solution to the problems of intentionality. I will have to keep my attention more narrowly focused on this occasion. So here I will assume that some variety of interpretationalism is correct, but I'll argue against Williamson's distinctive knowledge first version. This has repercussions for the kinds of objections to Williamson I can permit myself to raise. In particular, it won't do to push the kinds of worries about the inscrutability of reference that some (e.g. Williams 2007) have taken to cause problems for interpretationalist accounts of content in general.

There are a few other important preliminary points that I want to make before getting going, since this will give us a better picture of Williamson's project to carry forward into the critical discussion to follow. First, it is natural to wonder whether the interpretationalist story is intended to apply to both mental content and public language content (as I have tacitly assumed up until now), or only to one or the other of these. Williamson's discussion suggests that he envisages the interpretationalist story applying both to mental content and to public-language

content, and so I'll follow what I take to be Williamson's intent in not making any sharp distinction between these.

Second, unlike Lewis and some other interpretationalists, Williamson is not engaged in the task of reducing the intentional to the non-intentional (2007: 261), and so we won't concern ourselves with questions about the extent to which interpretationalism might enable one's reductionist ambitions to be realized.

Third, the approaches to interpretation and reference I will consider here are *holistic* rather than *molecular*, in the sense that their 'constraints on thinking given thoughts apply at the level of the subject's total system of thoughts, not at the level of individual constituents; they are global rather than local' (Williamson 2007: 259). Williamson advocates a holistic account (2007: 258-9), though as we'll see below, he sometimes appears to discuss particular examples in isolation in a manner one might have thought distinctive of a molecularist. The prime example of this is Williamson's treatment of the example he employs against a crude principle of truth maximization—one might have thought that a thorough-going holist would be more reticent about drawing the conclusion that a constraint on interpretation is no good just on the grounds that it fails to match with the verdict we are inclined to offer concerning particular carefully constructed cases. One response that can be made on Williamson's behalf here is that the case he discusses is just one particular example of a class of cases that can be constructed in a systematic manner (compare Grandy 1973: 444).

However, it is arguable that this response doesn't go far enough.¹ The examples to be discussed below involve how to interpret particular tokenings of the pronoun 'she' (in thought or in speech), and it is simply assumed that the interpreter is charged with the task of assigning the correct *referents* to these tokenings. But given the holistic nature of interpretation, we might suggest that the interpreter's task should rather be, in the first instance at least, that of giving the correct *rule* for the pronoun, the rule that determines how referents are assigned given a context together with the intentions of the speaker. Engaging in this task requires considering a wide range of different uses of the pronoun in question, and it's not clear on what grounds we can assume that the interpretation that best meets the constraints on interpretation overall with respect to this range will be the interpretation which best meets the constraints in a single case (or in a particular type of case), and vice versa. So there's still room for a worry that, given holism, the putative counterexamples to be discussed below may miss their mark.

This is an important line of reply to those counterexamples, but I won't consider it further here. The reason is that this kind of reply engages the example Williamson uses to motivate his proposal just as readily as it engages the examples I use to attack it. This leaves these important issues about the form holistic versions of interpretationalism should take, and the role that consideration of individual cases (or types of cases) can play in deciding between rival accounts of the constraints on

¹ As emphasized to me by Mark Sainsbury and a referee. See Sainsbury 2005: chapter 4 for relevant discussion.

interpretation, unresolved. But for the purposes of the present discussion of Williamson's view, we can rest content with adopting the same assumptions that condition Williamson's own engagement with rival views.²

Lastly, it is worth noting that Williamson's principal motivations for knowledge maximization seem to have shifted over the years. In the footnote (1996: 520fn20) in which Williamson first made this proposal and in *Knowledge and its Limits* (2000: 267) the shift from truth to knowledge in the principle of charity was presented as an upshot of his tentative suggestion that we should replace Lewis's (1969, 1975) convention of truthfulness in a language $\mathcal{L}$ with a convention of knowledgeableness in $\mathcal{L}$. Unfortunately, Williamson never really paused to spell out the argument. In more recent presentations (2004, 2007), the proposal about conventions has dropped out of sight, and a knowledge maximizing principle of charity is instead principally motivated by the desire to mount a general response to a particular kind of skeptical challenge. As is well known, Donald Davidson took his principle of charity to have anti-skeptical import (1983). For given a principle of charity that constrains an interpreter—even an omniscient interpreter—to maximize true belief, we may conclude that beliefs of their nature tend to be veridical. Error on the massive scale envisaged by the skeptic turns out to be ruled out by the nature of belief itself. Likewise, Williamson's principle of knowledge maximization is intended

² It is worth noting that Williamson's discussion hints at one response he might make to the worry. He suggests in one place that the force of single examples for a holist is that they can show that a proposed constraint on interpretation has implausible consequences concerning which interpretation is *defeasibly* favored by a given case (see Williamson 2007: 265). More needs to be said to flesh this out if it's to silence the objection, and I won't pursue the issue further here.

to have an anti-skeptical payload. On Williamson's account, beliefs of their nature tend to constitute knowledge. So knowledge maximization implies that we tend to apply concepts correctly in judgment, thereby undermining 'judgment skepticism', according to which we systematically misapply everyday concepts such as *mountain*, *person*, *belief*, or *knowledge* in judgment. I will have occasion to make reference to Williamson's goal of responding to judgment skepticism later in the discussion, and it will take center-stage in the final section.

A map of the terrain ahead will be useful. In Section 2 I introduce the example Williamson has used to motivate understanding charity as knowledge maximization, and note that the immediate significance of the example is that it reminds us of the importance of respecting causal connections between a speaker and the things he is intuitively speaking about, when such connections are present. Adopting the principle of knowledge maximization is a further step. In section 3 I argue against taking this further step on the grounds that we can construct cases of determinate reference where the principle of knowledge maximization plays no role, or even fails to be satisfied. I also argue against two fallback positions suggested by Williamson discussion. Finally, section 4 measures the import of my criticisms of knowledge maximization for our engagement with judgment skepticism. There I will argue that it is a mistake to think that judgment skepticism presents a unified target, and so a mistake to look for a unified response of the sort offered by Williamson.

2. Towards Knowledge Maximization

Williamson's argument for knowledge maximization can be naturally divided into two parts. First, he argues that a crude principle of truth maximization leads to incorrect assignments of reference since it ignores the role played by causal connections in fixing the reference of our expressions. Second, he suggests that such causal connections are important in this regard because they act as channels for knowledge. Assignments of reference that respect these causal connections win out because they maximize knowledge even when they do not maximize true belief. It is vital to distinguish these two stages in Williamson's argument. I do not doubt the importance of respecting causal connections when engaged in interpretation; Williamson is surely right to hold that a constraint on interpretation that is entirely blind to such connections is problematic. But we'll see that this is an insight that one should appreciate without making the further step to knowledge maximization.

Against a crude truth maximizing principle of charity, Williamson offers the following example. Emanuel is mistakenly confident in his ability to read off a person's character and life-history just by looking at their face. He is presently looking at a stranger, Celia, who is standing several feet in front of him. Upon seeing her face for a few moments, he judges 'She is F, G, H, ...', thereby ascribing in some detail a character and life-history. As usual, he is way off—none of it fits Celia. But it

very precisely fits Elsie, who Emanuel has never seen or even heard of. Williamson asks (2007: 263):

Does the pronoun “she” as used by Emanuel in this context refer to Celia or to Elsie? Which of them does he use it to express beliefs about? He accepts “She is standing in front of me,” which is true if “she” refers to Celia but false if it refers to Elsie. However, he also accepts “She is F,” “She is G,” “She is H, ...,” all of which are false if “she” refers to Celia but true if it refers to Elsie. We may assume that the latter group far outweighs the former. A principle of charity that crudely maximizes true belief or minimizes error therefore favors Elsie over Celia as the referent of the pronoun in that context. But that is a descriptive theory of reference gone mad. Emanuel has no beliefs about Elsie. He has many beliefs about Celia, most of them false. In virtue of what is Emanuel thinking about Celia rather Elsie?

The natural answer appeals to the fact that Emanuel’s uses of ‘she’ in this context stands in a causal relation to Celia but not Elsie. As Williamson points out, we must be careful here, since not just any causal relation will suffice for reference. But with that important caveat in mind, the idea that it’s the causal relation that secures that Emanuel’s uses of ‘she’ refers to Celia rather than Elsie is very tempting indeed.

Williamson suggests that this thought is right as far as it goes, but that it does not go far enough. He has two complaints. The first is that a purely causal story does not

enable us to say anything very informative about *which* causal relations are relevant for reference; they cannot say much more than that the speaker needs to be causally connected to the right objects ‘in the right way for reference, whatever that is’ (2007: 263). Second, we want a more general account of reference: one that can subsume the successes of the causal theory, but can also enable us to account for the cases that the causal theory fails to cover, including, perhaps, reference to abstracta (2007: 264).

What Williamson proposes is a knowledge-based account of the significance of causal relationships, in the cases where such relationships seem important for reference. In doing so he hopes to offer a more informative answer than a purely causal story can offer to questions about what fixes what it is we are thinking and talking about, and moreover an answer that generalizes to cover cases in which we seem to have reference to an object but not in virtue of any causal relationship. Williamson suggests that a ‘natural idea’ is that the significance of the perceptual link between Emanuel and Celia, and more generally the significance of any appropriate causal link, is that it is a channel for knowledge, and so leads to an assignment that maximizes knowledge. In the example we have been discussing, assigning Celia as the referent of Emanuel’s uses of ‘she’ in that context had Emanuel come out as truly accepting ‘She is standing in front of me’ and a few other claims of this sort—it was just that he came out as having more true beliefs on an interpretation which assigned Elsie as the referent. But notice that even though the latter interpretation makes each of Emanuel’s beliefs ‘She is F’, ‘She is G’, ‘She is H’,

..., true, none of those beliefs will amount to knowledge. In contrast, an interpretation that assigns Celia as the referent of Emanuel's uses of 'she' will have him come out as expressing knowledge with 'She is standing in front of me', and likewise for other perceptually-based claims of this sort. So a principle of knowledge maximization rules that we should favor the second interpretation over the first. As Williamson writes (2007: 264):

The assignment of Celia wins because it does better with respect to knowledge, even though it does worse with respect to true belief.

This is the picture I wish to resist. In the next section I will offer an argument against Williamson's proposal, and then consider two fallback positions.

3. Against Knowledge Maximization

Williamson considers whether a variant of the example he uses against a principle of true belief maximization might similarly tell against his own principle of knowledge maximization, and his discussion will provide a useful point of departure for developing my worries with his proposal. Suppose that things are just as they were in the earlier example in the following respects; Emanuel is looking at a stranger, Celia, and upon seeing her face for a few moments, he incorrectly but confidently judges 'She is F, G, H, ...', thereby ascribing a character and life-history in some detail. This time, however, these descriptions not only happen to be true of

Elsie; Emanuel is actually acquainted with Elsie, and happens to *know* each of the claims he would express with ‘She is F’, ‘She is G’, ‘She is H’, and so on, were ‘she’ interpreted as referring to Elsie. Still, in the context under consideration he judges ‘She is F’, ‘She is G’, ‘She is H’, ..., on the basis of looking at Celia, and so we still want to say that intuitively his beliefs are about Celia rather than Elsie. The problem for Williamson is that the assignment of Elsie seems to maximize knowledge, and so his principle seems to give precisely the wrong verdict here. In response, Williamson writes (2007: 267):

However, he [Emanuel] can still use “she” as a visual demonstrative to refer to Celia in judging “She is F,” “She is G,” “She is H, ...,” thereby expressing false beliefs about Celia rather than knowledge about Elsie, because those judgments are not causally based on his independent knowledge of Elsie, and therefore fail to express that knowledge.

The suggestion is that, contrary perhaps to initial appearances, knowledge maximization doesn’t get this case wrong, since Emanuel doesn’t express knowledge about Elsie even if Elsie is assigned as the referent of ‘she’. Even though Emanuel *possesses* more knowledge about Elsie than Celia, he only *expresses* knowledge if Celia is assigned as the referent, and so knowledge maximization correctly rules that we should favor interpretations that assign Celia rather than Elsie.

But it is natural to wonder whether the principle of knowledge maximization is actually doing any real work in producing the right response in this case. Take a variant of the case in which Emanuel is actually looking at a reflection of Celia. As in Williamson's example, in addition to forming the judgments 'She is F,' 'She is G,' and so on, he also forms the judgment 'She is standing in front of me', but no other judgments of this sort. This judgment is actually false; Celia is not standing in front of Emanuel, though it appears to him just as if she was. Here it is natural to think that Celia is again the referent of Emanuel's uses of 'she', even though the causal link between them is very slightly deviant. But this result isn't being delivered by knowledge maximization, since Emanuel will fail to express *any* knowledge whether we assign Elsie or Celia as the referent. It looks like the causal and other facts decide in favor of Celia all by themselves, and one might wonder why we shouldn't say the same in Williamson's original case and in his variant.

In building the variant of Williamson's examples considered in the previous paragraph, I stipulated that aside from his judgments detailing a life-history and character, Emanuel makes only one further judgment, namely 'She is standing in front of me'. That was important, since it's obviously crucial to my argument that Emanuel not form any judgments of that sort that are knowledgeable. Is this a legitimate stipulation? Well, perhaps not. But I'm inclined not to worry, since we can construct examples in which we can afford to relax such stipulations about which judgments the subject makes, since (so I'll argue) *no* such judgments will be knowledgeable. Consider the following case:

Hallucinogen: Jane is taking part in a trial of a new drug with powerful hallucinogenic properties. Sitting in her cubicle, she appears to see a number of people in front of her. In fact, there is only one other person in the cubicle, Helen, who is standing several feet away from her. Helen appears to Jane to be no more flesh and blood than the other figures that appear to her in the cubicle. Despite knowing that she will experience hallucinations while in the cubicle, Jane judges ‘She is beautiful’, ‘She is freckled’, ‘She has green eyes’, and so on, keeping her attention firmly on Helen the whole time. Each of these judgments is correct.

Whether one will agree that none of the beliefs Jane comes by in her cubicle are knowledgeable will depend on what one thinks is required for knowledge. But let me give two considerations in favor of this verdict. First, Jane has an undercutting defeater for any of the beliefs she forms on the basis of perception during the duration of the trial, and so it is plausible that her beliefs cannot even be considered justified. Second—and perhaps more to the point given Williamson’s own views on knowledge—as the case has been described, Jane could easily have focused her attention on one of the other apparent figures in the room, and had she done so, she would not have formed true beliefs. So, under some assumptions to be made explicit in a moment, Jane’s beliefs fail to meet the following plausible necessary condition on knowledge:

Safety: S knows that P only if in all the nearby possible worlds in which S forms a relevantly similar belief on a relevantly similar basis as S forms her belief that P in the actual world, that belief is true.

Now, there are two important observations we must make about *Safety*. First, this formulation differs from standard formulations of a safety condition for knowledge in allowing the content of the relevant beliefs at different worlds to vary; we are to consider beliefs with contents relevantly similar to P, formed on a basis relevantly similar to the basis S actually formed the belief that P on, rather than just P itself. It must be stressed that this is no *ad hoc* maneuver. A familiar problem for standard formulations of safety conditions for knowledge is that they are trivially met by modally-robust truths, most obviously necessary truths. Suppose one guesses that 78 and 96 make 172. Since there's no world in which that's false, there's clearly no nearby world in which one forms that belief on a basis relevantly similar to one's actual basis, and yet fails to come by a true belief. But intuitively this is precisely the sort of case a safety condition on knowledge was meant to exclude. *Safety* gets the right result here, at least if we are willing to give the hedge 'relevantly similar' a reasonable interpretation; the belief that 78 and 96 make 171 should count as similar, and there are close worlds in which one guesses and forms this belief instead of the belief that 78 and 96 make 172, thereby coming by a false belief rather than a true one. So formulating our safety condition on knowledge as

something like *Safety* provides a neat way of solving this familiar problem with modally-robust truths (Sainsbury 1997, Weatherson 2004, Pritchard 2009).³

The relevance of all this to our discussion of *Hallucinogen* is this. If ‘she’ refers to Helen, then there are no close worlds in which Jane forms the belief she would express with ‘She is beautiful’ on the basis of perception, and yet thereby fails to form a true belief. However, had Jane focused on one of the other apparent figures in the room, as she could easily have done, she would have formed beliefs on the basis of perception that are not true. So if we may assume that the beliefs Jane would form if she had shifted her focus to any of the other apparent figures in her cubicle count as sufficiently similar to the beliefs she actually forms, Jane does not meet *Safety*. On reflection, this is just what we should expect. *Safety* offers a safety condition that fails in ‘fake-barn’ cases, and *Hallucinogen* is at heart a modified barn case.

The second observation I want to make about *Safety* is that as stated it requires that the beliefs one forms in the relevant worlds be true, rather than that they just not be false. Of course, under certain assumptions these conditions are equivalent, but those are assumptions one cannot make in the present context. It is quite natural to think that *Safety* fails in *Hallucinogen*, not because Jane would form *false* beliefs if she shifted her attention, but because the pronouns she would employ would fail to refer to anything were she to do so, this reference-failure rendering her belief

³ It is worth noting that Williamson has recently offered a formulation of safety which is identical to *Safety*, except that, in line with the version of safety he defends in *Knowledge and its Limits*, it is close cases—centered worlds, in David Lewis’s terminology—rather than simply close worlds that are deemed relevant to whether one’s belief counts as safe. See Williamson 2009a: 325.

neither true nor false. This is not my own preferred take on the case,⁴ but those who do think of reference-failure as leading to truth-value gaps of this sort will have to decide if they want to construe *Safety* such that the easy possibility of gappy beliefs renders one's belief unsafe. They should (as noted in Hawthorne 2004: 56fn17); the thought behind safety conditions is that knowledge requires that there be a modal robustness to one's getting things right.

So I think it's very plausible that none of the beliefs Jane forms in this case are knowledgeable. One can resist this conclusion by maintaining that one can have knowledge even if one has an undercutting defeater and fails *Safety*, but that will be a tough pill for a lot of us to swallow. Nonetheless, although unknowledgeable, Jane's judgments are *about* Helen; the causal connection is of the right sort to act as a channel for reference. Assignments that assign Helen as the referent of Jane's uses of 'she' should be preferred, even though they do not maximize knowledge. And if knowledge maximization isn't playing any significant role in this case, why think it's any less of a spare gear in Williamson's Emanuel case and the variants we have considered?

One might suggest that Williamson doesn't really need to appeal to knowledge *maximization* to explain the significance of causal relations for reference. Recall from above that Williamson (2007: 264) makes the following proposal:

⁴ See Sainsbury 2005 for an attractive alternative.

Roughly: a causal connection to an object (property, relation, ...) is a channel for reference to it if and only if it is a channel for the acquisition of knowledge about the object (property, relation, ...).

It doesn't look like this proposal entails the principle of knowledge maximization and so it is worth detaching it from that principle and considering it as an independent claim about the significance of causal connections, when they are present. But even with the prefixed hedge, Williamson's claim is too strong. Consider again the case *Hallucinogen*. The causal connection between Jane's uses of 'she' and Helen is plausibly a channel for reference, but this connection is not a channel for the acquisition of knowledge about Helen. The argument for this is just the one offered above. This objection leaves intact the plausible thought that a causal connection is a channel for the acquisition of knowledge about an object only if it is a channel for reference to that object. But I think that the above example shows that it is not even roughly true that a causal connection is a channel for reference to an object only if it is a channel for the acquisition of knowledge about that object. So I think that a knowledge first account of the significance of causal relations for reference is not merely unnecessary to account for the cases we've considered so far, but actually mistaken.

There is, however, a second fallback worth considering. Williamson might respond to this objection in much the same way as he has responded to a recent objection to knowledge maximization due to M. G. F. Martin. Martin (2009) tries to construct a

case in which one is simultaneously attending to two qualitatively identical candidates that one's thought 'That is pink' could be about, and where one would express knowledge whichever candidate was assigned as the referent of 'that'. Martin's thought is that one could still have a thought about the object on the left, even though one would equally express knowledge if the object on the right were assigned, and even though causal facts about which objects one is visually attending to don't discriminate between the two candidates. Williamson considers a number of responses to this worry, but only one is relevant here. He writes (2009b: 471):

An appealing picture is that when one starts using 'that' with reference to lefty, one opens some sort of mental file, if only a very temporary one, with a predominant causal connection to lefty that enables it to act as a channel for perceptual information about lefty more directly than it can act as a channel for perceptual information about righty: potentially, a channel for *knowledge* of lefty rather than righty. It does not matter whether that knowledge includes the particular item that it is pink, for the thought 'That is pink' still has a compositional semantics. Even if no knowledge actually happens to be gained through that channel, the naturalness of the reference relation may still keep the reference constant between the actual case and counterfactual cases in which knowledge of lefty is gained through the channel.

We might adapt this response to the case I offered above; Helen is the referent of Jane's uses of 'she' in virtue of a causal connection that *would* be a channel to

knowledge in counterfactual cases. I confess that I don't see how to construct a counterexample to this weaker claim. But the proposed connection between reference and knowledge seems so diluted that it is unclear what significance we should attach to it, even if it holds; it is unclear why one couldn't happily admit the connection without holding that there's a particularly intimate connection between reference and knowledge, or conceding any major victory to knowledge first philosophy. We can make this a little more pointed. In section one I noted that in Williamson's recent discussions of knowledge maximization he is concerned to undermine judgment skepticism by showing that beliefs of their nature tend to constitute knowledge, and thereby tend to be true. We'll examine this line of thought in more detail in section 4 below. But the point for just now is that the revised principle suggested by Williamson's reply to Martin doesn't have this anti-skeptical consequence. For the revised principle only makes being a *potential* channel for knowledge a necessary condition on being a channel for reference, and so it does not generate a presumption that our beliefs are *actually* knowledgeable or true. There's a danger, then, that revised proposal fails to engage the judgment skeptic. Of course, by itself this doesn't have any tendency to show that the revised proposal is incorrect. But it does serve to reinforce the worry that the revised proposal isn't really articulating a point of contention between Williamson and his opponents.⁵

⁵ A referee has suggested an interesting variant of this second fallback position. Recall Williamson's claim, discussed above: 'Roughly: a causal connection to an object (property, relation, ...) is a channel for reference to it only if it is a channel for the acquisition of knowledge about the object (property, relation, ...)' (2007: 264). We might accept that *Hallucinogen* provides a counterexample to this claim, but advance a related claim about causal connection *types* in its place: roughly, a causal connection type to some objects (properties, relations, ...) is a channel for reference to those objects (properties, relations, ...) only if it is a channel for the acquisition of knowledge about those objects (properties, relations, ...). If I've understood it, the suggestion is now that a causal connection type can count as a

There may of course be other ways of developing the thesis that there is some kind of constitutive link between reference and knowledge, but none of the three suggested by Williamson's discussion have proved promising, and I don't see grounds to be optimistic that any intimate connection of the sort Williamson wants will emerge. This leaves defenders of interpretationalist accounts of content with the related challenges of explaining the significance of causal connections when they are present, and of explaining how reference might be effected when they are absent. Rather than broach those rather large issues here, however, I want to turn at last to Williamson's discussion of judgment skepticism.

4. Judgment Skepticism

Judgment skepticism has been mentioned at several junctures throughout our discussion, and I want to close by considering the import of the conclusions of the last section for our engagement with such a skeptic. The judgment skeptic attacks

channel for the acquisition of knowledge so long knowledge is acquired via its tokens in *normal* cases. That's compatible with *Hallucinogen*, if we grant both that it's an abnormal case, and that the token causal connection in the example is a token of a type that is a channel for knowledge, in the proposed sense. It's not at all clear that my objection in the text to Williamson's own version of the fallback applies with the same degree of force here; perhaps a presumption is generated that our beliefs are normally knowledgeable, and perhaps we should allow that this would be a significant victory against the judgment sceptic (though see the next section for doubts about Williamson's whole strategy here). But I have some misgivings about the proposal. First, as the referee points out, we might well wonder how we are to determine which causal connection type is relevant in a given case, and in particular whether we can do so in a way that let's us avoid getting snagged in analogues of precisely the difficulties that Williamson takes to scupper an unadorned causal theory of reference, or whether we have enabled ourselves to say something more informative about which causal connection tokens are reference relations only by assuming we have an independent grasp on which causal relation types are relevant to reference. Second, we need some control on the notion of normality in play here, since if the normal cases are characterized with reference to those in which knowledge is acquired, the proposal seems rather empty and ad hoc.

‘our practices of applying concepts in judgment’ (Williamson 2007: 220).

Williamson (2007: 220-1) offers the following (putative) examples:

Call “judgment skeptics” those skeptical in the way just sketched about some contextually relevant judgments. For example, in a context that concerns folk psychological ascriptions of belief and desire, Paul Churchland and other eliminativists about such mental states are judgment skeptics. In a context that concerns ordinary geographical judgments, Terry Horgan and other eliminativists about mountains are judgment skeptics. Such skeptics question our standards for applying ordinary concepts both in experience and thought: the concept of a mountain, the concept of belief, the concept of knowledge, the concept of possibility, the concept of the counterfactual conditional, and so on.

Williamson claims that, like more familiar skeptics, judgment skeptics argue for their position by presenting skeptical scenarios—scenarios in which most of our judgments about some contextually relevant subject matter are false but in which there is some story about how we understandably come to make those judgments despite their falsity—and challenge us to find grounds for continuing to regard those judgments as true (Williamson 2007: 221-2). But why think this challenge is any more worrying than that presented by the skeptic about, say, perception? Williamson suggests (2007: 250-1) that the reason that judgment skepticism merits special consideration is that a case can be made for thinking that the kinds of

scenarios the judgment skeptic plays with may well be instantiated by other cultures in the actual world (or at least in nearby worlds), whereas skeptics about perception invoke modally remote scenarios involving brains in vats or extended and coherent dreams. So we cannot dismiss the judgment skeptic's arguments on the grounds that the scenarios appealed to are too 'fanciful or far-fetched'. Williamson concludes (2007: 251):

If we are to refuse in good conscience to take seriously the radical scenarios for judgment skepticism, we must do so from a perspective on which there is a quite general tendency for beliefs to be true. Anything else will look like special pleading on our own behalf.

We can now see the role that Williamson's principle of knowledge maximization is to play in his response to judgment skepticism. Given the constitutive reading of the principle of knowledge maximization, we may conclude that beliefs of their nature tend to be knowledgeable. And since knowledge is factive, this implies that beliefs tend to be true, just as Davidson held. This applies not only to our own beliefs, but also those of geographically and modally proximate believers. So a successful defense of knowledge maximization would put us in a position to dismiss the judgment skeptic's scenarios once and for all. Not only would we have assurance that we are not prone to error on the scale envisaged by some judgment skeptics, we would also have grounds for holding, contrary perhaps to initial appearances, that the possibility of error on this scale is remote.

That's the line. But I'm doubtful that judgment skepticism as Williamson characterizes it represents a unified target that might be addressed with appeal to the kind of general argument Williamson wants to offer. There seem to be at least two rather different kinds of philosophical positions being run together under the label 'judgment skepticism'. The first is a kind of skepticism about a class of judgments that contends that since the judgments of members of other cultures or our past selves are or were wrong, we have no right to think that our judgments about the relevant subject matter get things right (and perhaps even have inductive reasons to think that they too are false—this kind of pessimistic induction is very familiar from ethics, the philosophy of religion, and the philosophy of science).

The second is the repudiation of a certain class of objects, typically leading to the corollary that judgments affirming or presupposing the existence of those objects are false or meaningless.⁶ The grounds offered for repudiating the ontology can be various. Consider the examples Williamson appeals to in the passage quoted above. Churchland (1981) argues that folk psychology, which commits to the existence of beliefs, desires, and mental states, is a moribund theory; it is explanatorily inadequate even when one favorably chooses its domain of applicability and it coheres remarkably poorly with our best scientific theories. From this he concludes that we should be eliminativists about mental states. According to Horgan (1995)

⁶ The qualification 'typically' is to leave room for eliminativist materialism—there is a manifest awkwardness in taking this to be a form of judgment skepticism in Williamson's sense, since this kind of skepticism presupposes rather than denies the existence of judgments.

vagueness is a form of incoherence, so since the world cannot be incoherent, there can (strictly speaking, at least) be no vague objects, mountains included. Plainly neither of these arguments invokes the kinds of skeptical scenarios Williamson takes to be characteristic of judgment skepticism, nor the kind of familiar skeptical argumentation we were led to expect.

Once we tease apart these two broad kinds of challenges to common-sense, it becomes very hard to see how appealing to knowledge maximization is supposed to help speak to either. Presumably we want to be able to say that nothing in the nature of intentional content prevents our past selves or other cultures from having mostly false beliefs about (say) morality, religion or science. If knowledge maximization lets one purchase that much security against widespread error, that is a serious problem for the proposal rather than a victory for common-sense.⁷ Now, this kind of skepticism does rely on arguments that bear considerable resemblance to traditional skeptical arguments, and Williamson argues at length that this gives us reason to be very suspicious; it is very difficult to see how such arguments could apply selectively against particular targets, and we are confident that these

⁷ Williamson recognizes a role for naturalness, of the sort familiar from Lewis 1974 and elsewhere, in determining the reference relation, and he suggests that this 'holds the anti-skeptical effect of knowledge maximization within reasonable limits' (2007: 268). But the worry Williamson is responding to here is that knowledge maximization might force us to interpret Stone Age people as talking and thinking about quantum mechanics, as this interpretation might have them come out as expressing more knowledge than one that has them believing and asserting the (presumably false) theories of the world in currency at that time. The appeal to naturalness isn't enough to escape the present worry, however. Naturalness may ensure that we *don't* interpret Stone Age people as believing aspects of contemporary quantum theory, but it can't help ensure that we *do*, despite being constrained to maximize knowledge, interpret them as believing the largely false theories of the day.

arguments cannot be sound in general. This point is perhaps well taken, but whether or not it is, it seems completely independent of knowledge maximization.

On the other hand, arguments like those presented by Churchland and Horgan surely deserve and need to be addressed head on. Neither Churchland nor Horgan relies on a style of argument that bears any significant resemblance to traditional arguments for skepticism, and so we cannot simply dismiss their arguments on the grounds that we can be confident there is a mistake somewhere. Williamson does stress that we need to judge arguments that do not turn on 'general skeptical fallacies' on their own merits (2007: 241). But it doesn't seem like Williamson has Horgan and Churchland's arguments in mind here. Rather, he seems to take these philosophers to be offering skeptical arguments. He cites Horgan and van Inwagen's claim that there are no mountains as an example of the kind of 'skeptical scenario' employed by judgment skeptics (2007: 221-2), though he notes that unlike more traditional skeptics, the judgment skeptic 'often' argues that we're actually in the skeptical scenario (2007: 222). And later he writes: 'A judgment skeptic argues that our evidence is neutral between the ordinary hypothesis that there are mountains and the skeptical hypothesis that there are no mountains, but instead only complex microphysical events the human brain usefully but untruthfully classifies as mountains, and concludes that we cannot know and are not justified in believing that there are mountains' (2007: 225-6). There's no sharp distinction drawn in Williamson's discussion between skeptical arguments that turn on the contention that our evidence doesn't favor the common sense hypothesis over some skeptical

hypothesis, leading to the epistemological conclusion that ‘we cannot know and are not justified in believing’ that things are as common sense says they are, and the kind of metaphysical or methodological arguments offered by Horgan and Churchland for the metaphysical conclusions that there are *in fact* no mountains or states of the sort posited by folk psychology.

I suspect that it is only by failing to sharply distinguish the kind of arguments offered by Churchland and Horgan from the skeptical arguments one finds directed against the claim that our present religious, ethical and scientific beliefs are justified or knowledgeable, that we might find ourselves tempted into thinking that the former philosophers might be satisfactorily addressed by the sort of sweeping anti-skeptical argument supposedly supplied by knowledge maximization. Once we get clearer on the very different challenges presented to common-sense by the various philosophers Williamson wants to classify together as judgment skeptics, it becomes doubtful that a unified response is possible or desirable.

5. Conclusion

We have no assurance of the sort Williamson wishes to provide that ignorance on a massive scale is ruled out. It remains an open question how damaging a concession this need be. On the knowledge first picture of philosophy, it is potentially devastating. If knowledge provides the normative standard for belief, assertion, and action, we face the unnerving possibility that as a believer, asserter, and agent, one

may frequently hold out hostages that cannot be redeemed.⁸ Only a careful examination of the most powerful arguments for skepticism will reveal whether this is a possibility that need cause proponents of knowledge first philosophy genuine concern; consideration of Williamson's attempt to provide independent assurance that it need not has proved to reinforce rather than to undermine this point.⁹

⁸ As Unger argued in *Ignorance* (1975).

⁹ My thanks to audiences at the University of Texas at Austin, the Northern Institute of Philosophy, and the 2010 *Episteme* conference at the University of Edinburgh, as well as to Ray Buchanan, Brian Cutter, Josh Dever, Richard Heck, Bryan Pickel, and Crispin Wright. Particular thanks are owed to Mark Sainsbury, David Sosa, and a very helpful referee for this journal.

References

- Churchland, P. M. 1981. Eliminative Materialism and the Propositional Attitudes. *Journal of Philosophy* 78: 67-90.
- Davidson, D. 1983. A Coherence Theory of Truth and Knowledge. In Henrich, D., ed. 1983. *Kant oder Hegel?*. Stuttgart, Klett-Cotta: 423-38.
- Grandy, R. 1973. Reference, Meaning, and Belief. *Journal of Philosophy* 70: 439-52.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford, Oxford.
- Horgan, T. 1995. Transvaluationism: a Dionysian Approach to Vagueness. *Southern Journal of Philosophy* 33 supplement: 97-126.
- Lewis, D. 1969. *Convention: A Philosophical Study*. Oxford, Blackwell.
- 1974. Radical Interpretation. *Synthese* 23: 331-44.
- 1975. Languages and Language. *Minnesota Studies in the Philosophy of Science* 12: 3-35.
- Martin, M. G. F. 2009. Reupholstering a Discipline: Comments on Williamson. *Philosophical Studies* 145: 445-53.
- Pritchard, D. 2009. Safety-Based Epistemology: Whither Now? *Journal of Philosophical Research* 34: 33-45.
- Sainsbury, R. M. 1997. Easy Possibilities. *Philosophy and Phenomenological Research* 57: 907-19.
- 2005. *Reference Without Referents*. Oxford, Oxford.
- Unger, P. 1975. *Ignorance*. Oxford, Oxford.

Weatherson, B. 2004. Luminous Margins. *Australasian Journal of Philosophy* 82: 373-83.

Williamson, T. 1996. Knowing and Asserting. *Philosophical Review* 105: 489-523.

— 2000. *Knowledge and its Limits*. Oxford, Oxford.

— 2004. Philosophical 'Intuition' and Skepticism about Judgment. *Dialectica* 58: 109: 153.

— 2007. *The Philosophy of Philosophy*. Oxford, Blackwell.

— 2009a. Replies to Critics. In Greenough, P, and Pritchard, D., eds. 2009. *Williamson on Knowledge*. Oxford, Oxford: 282-384.

— 2009b. Replies to Ichikawa, Martin and Weinberg. *Philosophical Studies* 145: 465-76.