

WHAT DO PHILOSOPHERS KNOW?¹

Andrew MELNYK
University of Missouri

In his new book, Timothy Williamson makes a frontal assault on the questions of what philosophy is, what sort of knowledge it can attain, and what methods it can and should use. He follows no familiar party line, and indeed has something to offend everyone². In the book's first part, comprising chapters one through four, he works toward the conclusion that philosophy, unlike, say, astronomy, has no special subject matter, and that, in particular, "few philosophical questions are conceptual questions in any distinctive sense, except when philosophers choose to ask questions about concepts" (Williamson 2007, 3). In the book's second part, comprising chapters five through eight, he makes a start on a new epistemology of philosophy, one intended to vindicate the view that the knowledge achieved and methods used in philosophy are no different in kind from the knowledge achieved and methods used in everyday inquiry and in science. But he does not proceed, as perhaps one might have expected, by defending a thorough-going naturalism, according to which philosophy, being continuous with science, uses versions of scientific methods; he never has a good word for naturalism, and calls it "crude empiricism" (1-2). In an afterword, entitled "Must Do Better", he urges philosophers to hold themselves to higher methodological standards and thus realize what he takes to be philosophy's hitherto unrealized potential. (In the UK, "Must do better" used to be a stock phrase from end-of-term reports on the academic performance of schoolchildren.) The book ends with two formal appendices.

Williamson does not intend his book to be a systematic treatment of the issues that it addresses, and indeed it is not. But it is full of challenging, penetrating, and highly creative philosophy, and to read the book is to keep company with an outstanding intellect.³ It sets a new standard for philosophical inquiry

1. Review essay of Timothy Williamson, *The Philosophy of Philosophy*. Malden, MA: Blackwell Publishing, 2007. xiv+332 pp. ISBN: 978-1-4051-3397-5.

2. I mean, of course, that $\forall x (x \text{ is a person} \rightarrow \exists y (y \text{ is in Williamson's book} \ \& \ y \text{ offends } x))$!

3. *Caveat emptor!* The book "is based on a series of articles in which earlier versions of the ideas were formulated, although hardly any pages have survived completely unchanged" (xiii).

into philosophy itself, and ought to elevate such inquiry from its present low place to that of a specialty in its own right. It also contains material of general philosophical interest, e.g., the third and fourth chapters in which the idea of analytic or conceptual truths is subjected to an unusually searching and imaginative examination.

So I recommend the book highly. But I do have a reservation. I found reading the book to be a work-out, and no doubt this was due in part to the originality of Williamson's thought and his fondness for formalism. But it was also in part because the book too often fails to be sufficiently clear. In this regard the book is a strange mixture. On the one hand, Williamson understands the importance of clarity in philosophy, aims to achieve it, and often succeeds admirably, both with and without the use of formalism. On the other hand, there is much that is needlessly puzzling, often, though not always, because of what Williamson has *omitted* to say. Chapter eight, discussed at some length below, provides an extended example. For a briefer example, consider the following paragraph, which I quote in its entirety:

In general, our capacity to evaluate counterfactuals recruits *all* our cognitive capacities to evaluate sentences. A quick argument for this uses the assumption that a counterfactual with a true antecedent has the same truth-value as its consequent, for then any sentence **A** is logically equivalent to **T** $\square \rightarrow$ **A**, where **T** is a trivial tautology; so any non-logical cognitive work needed to evaluate **A** is also needed to evaluate the counterfactual **T** $\square \rightarrow$ **A**. [Irrelevant footnote omitted.] For if we could evaluate that counterfactual without doing the non-logical work, we could also evaluate **A** without doing it, by first evaluating the counterfactual, then deriving its equivalence to **A** and finally extending the evaluation of the former to the latter. Any logical work needed to evaluate **A** will also be needed to evaluate **T** $\square \rightarrow$ **A** when **T** is chosen to be irrelevant to **A**. (152)

The immediate conclusion of the "quick argument", I think, is that, for any cognitive capacity to evaluate sentences that we possess, there's a counterfactual that *requires* for its evaluation the exercise of that very capacity; this counterfactual's antecedent is a trivial tautology, its consequent a sentence that requires for its evaluation the capacity in question. But what next? How are we meant to move from here to the official conclusion, stated in the first sentence, that all our sentence-evaluative sub-capacities are actually recruited by our capacity to evaluate counterfactuals? We are not told. I assume the missing premise is that we can in fact evaluate any counterfactual whose antecedent is a trivial tautology and whose consequent is a sentence that we can already evaluate (e.g., "If it were the

case that Gordon Brown is doomed or Gordon Brown is not doomed, then I would be rich”). But I don’t find this premise obvious. There are other passages in the book where the formal material is clear, but the relation of the formal to the informal material is not (e.g. 232–234).

*

Let me turn now to the substance of Williamson’s book. Since I must be highly selective, I will confine my discussion to the book’s two main positive proposals, one from chapter eight and one from chapter five.

A Davidsonian principle of charity in interpretation implies that an interpretation of—an assignment of content to—the mental states of a subject must result in the subject’s turning out to have mostly true beliefs. In chapter eight, entitled “Knowledge Maximization”, Williamson begins by arguing, persuasively, that such a principle often yields implausible interpretations. So he proposes that we should “replace true belief by *knowledge* in a principle of charity constitutive of content” (264; my emphasis). Thus, “The right charitable injunction for an assignment of reference is to *maximize knowledge*” (265; my emphasis). He also mentions two constraints on putting this injunction into practice:

Objective limits on what subjects are in a position to know appropriately constrain the maximization of knowledge by the assignment of reference. (266)

The compositional structure of . . . thoughts further constrains the ascription of knowledge, because the inferential processes in which subjects engage are sensitive to that structure: to interpret those processes as yielding knowledge, one must interpret them as valid inferences. (266)

Now Williamson calls his proposal a “picture of the mind . . . sketched, with the broadest strokes” (273). It is, indeed, an elegant and intriguing idea, but it is so underdeveloped as to be very hard to assess. Here are some questions that are left unanswered. Is the idea meant to be that we know which states of a subject are mental states *prior* to our interpreting them? If so, as suggested by Williamson’s talk of the “assignment of reference [*sc.* to mental states]”, then how do we know this—especially while suspending judgment on the reference of mental states? And if not, as suggested by Williamson’s calling his proposal a “picture of the mind”, then what is the relationship between being *interpretable*

as being in mental states, on the one hand, and being *in* mental states, on the other? If Williamson's proposal is a variant of the view that being in mental states just *is* being interpretable as being in mental states, then how are standard objections to such a view to be met?

Whatever the precise nature of Williamson's principle of knowledge maximization in interpretation, what does it have to do with the rest of his book? The answer is clear in outline, but that's all. In the immediately preceding chapter seven, Williamson works to undermine what he calls "judgment skepticism"—roughly, skepticism about the Sellarsian manifest image (220–241). And he certainly means to continue this work in chapter eight:

Even if such scenarios [i.e., radical scenarios for judgment skepticism] are rare or absent in the actual universe, but only by good luck, it remains uncomfortable for opponents of judgment skepticism. 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*. (251; my emphasis)

He returns to this thought, after having laid out his principle of knowledge maximization in interpretation:

Although maximizing knowledge is not equivalent to maximizing true belief, the nature of reference grounds a general, highly defeasible tendency for beliefs to constitute knowledge, and therefore to be true. (270)

However, I don't see how the principle of knowledge maximization in interpretation is meant to ground a "general...tendency for beliefs to constitute knowledge". The principle of knowledge maximization merely requires that as many of a subject's beliefs as possible constitute knowledge. But there's a general tendency for a subject's beliefs to constitute knowledge, I take it, only if more than 50% of the subject's beliefs constitute knowledge. Since the maximum number of the subject's beliefs that can be interpreted as constituting knowledge might still amount to fewer than 50% of the subject's beliefs, the principle of knowledge maximization doesn't require—hence doesn't ground—a general tendency for a subject's beliefs to constitute knowledge. It appears, then, that Williamson has three options:

- replace the principle of knowledge maximization with an explicitly majoritarian principle of interpretation, e.g., "Interpret a subject so that most of her beliefs constitute knowledge";

- add a premise enabling him to argue that the principle of knowledge maximization is for some reason bound to yield an interpretation in which a majority of a subject's beliefs constitute knowledge;
- adopt a weaker construal of what constitutes a general tendency for a subject's beliefs to constitute knowledge.

I won't comment further except to note that the plausibility of the third option obviously turns in part on what sort of work Williamson has in mind for the putative general tendency for a subject's beliefs to constitute knowledge. Clearly he means the tendency somehow to undermine or discredit judgment skepticism; but how exactly?

You might think that a tendency for a subject's beliefs to constitute knowledge entails that the mere having of a given belief is in itself *prima facie* evidence that the belief is knowledge, and hence true. But Williamson is opposed to taking this tack, which, he says, "depends on the fallacy...of psychologizing evidence" (274).⁴ His tack is apparently different: the nature of reference, he says,

... helps put the burden of proof on judgment skeptics to argue that their radical scenarios deserve to be taken more seriously than do the radical scenarios for skepticism about perception ... [M]uch more than [metaphysical possibility] is needed to justify serious doubt. (273)

Two pages later he writes that

Knowledge maximization is a factor, typically unnoticed by judgment skeptics, that makes their scenarios more far-fetched than they realize. (275)

These remarks hint at the following line of reasoning. Skeptical scenarios that are sufficiently far-fetched do not deserve to be taken seriously; and "radical scenarios for skepticism about perception" fall into this category. The scenarios of judgment skeptics can also be seen to fall into this category, once (but presumably only once) one appreciates the general tendency for a subject's beliefs to constitute knowledge that follows from the (correct) principle of knowledge maximization.

If this is indeed the line of reasoning that Williamson intends, then it needs a lot more work. Is being far-fetched simply a matter of being improbable on current beliefs? If not, then what is it? Is there a rationale for the convenient doctrine that sufficiently far-fetched skeptical scenarios can be dismissed? How far-fetched is "sufficiently" far-fetched? And why should we think that the general

4. I don't claim to understand what exactly this fallacy is meant to be.

tendency for a subject's beliefs to constitute knowledge makes the scenarios of judgment skeptics more far-fetched *by enough* to carry them to the point where they can be dismissed in good conscience?

*

As noted above, Williamson sees philosophical knowledge and methods as no different in kind from those proper to everyday inquiry and to science. How, then, does he account for our knowledge of metaphysically necessary truths? He offers his account in chapter five, entitled "Knowledge of Metaphysical Modality". I will only discuss what he says about metaphysical necessities, though he intends his account also to cover our knowledge of metaphysical possibilities. My concern about his account of our knowledge of metaphysically necessities will not be that it's false, but rather that, for all its undoubted interest, it doesn't take us as far as one might like.

Williamson suggests that "the epistemology of metaphysically modal thinking is tantamount to a special case of the epistemology of counterfactual thinking"(158), so that the capacity we have for assessing the truth value of counterfactual conditionals can also be used to assess the truth value of claims of metaphysical necessity. How, then, do we assess the truth value of counterfactuals? Williamson summarizes his answer as follows:

We can ... schematize a typical overall process of evaluating a counterfactual conditional thus: one supposes the antecedent and develops the supposition, adding further judgments within the supposition by reasoning, offline predictive mechanisms, and other offline judgments. The imagining may but need not be perceptual imagining...Some but not all of one's background knowledge and beliefs are ... available within the scope of the supposition as a description of the counterfactual circumstances, according to complex criteria (the problem of cotenability). To a first approximation: one asserts the counterfactual conditional if and only if the development eventually leads one to add the consequent. (153)

Earlier (142–152), Williamson has stressed that the background knowledge that constrains the development of a supposition needn't take the form of an explicit representation; it can be, and often is, embodied in imaginative capacities or methods of belief formation. In similar vein, he has expressed doubt that the development of a supposition can usually be understood as a process of inference in any non-trivial sense.

Now although Williamson promises an *epistemology* of metaphysical modality, his account of the process by which in fact we come to assert a counterfactual does not specify the conditions under which such an assertion would express *knowledge*. However, Williamson does go on to say that

... the reliability of our cognitive faculties in their online applications across a wide range of possible circumstances induces reliability in their offline applications too. (155)

Since very soon thereafter he also says that “we should admit that our methods sometimes yield knowledge of counterfactuals” (155), I assume that, at least roughly, he takes the reliability of the process by which we come to assert a true counterfactual as both necessary and sufficient for the assertion to express knowledge of the counterfactual.

According to Williamson, the capacities we exercise when we come to know counterfactuals can also give us knowledge of metaphysical necessities. What makes this possible, he argues, is that any claim of metaphysical necessity is logically equivalent to a counterfactual; that is, where ‘ \perp ’ is a contradiction, $\Box A \equiv (\neg A \Box \rightarrow \perp)$. Thus, to assess the claim that $\Box A$, we need only suppose that $\neg A$, and then develop this supposition. If we reach a contradiction, then we assert that $\Box A$. If we assert that $\Box A$, and the assessment procedure by which we came to do so was reliable, then we know that $\Box A$.⁵

Note that Williamson’s assessment procedure, when all goes well, yields the knowledge that $\Box A$, rather than the knowledge that A . Thus, he is offering an account of our knowledge of the status of metaphysically necessary claims *as* metaphysically necessary. He isn’t offering an account of our knowledge of metaphysically necessary truths, which truths one might know without knowing that they’re metaphysically necessary. Nor is he offering an account of how one acquires knowledge of metaphysically necessary truths.

Now sometimes Williamson seems to argue that, given his account of our knowledge of metaphysical necessities, one cannot be skeptical about metaphysical necessity without also being skeptical about counterfactual conditionals:

Since our capacity for modal thinking cannot be isolated from our capacity for ordinary thinking about the natural world, which involves counterfactual thinking, skeptics about metaphysical modality cannot excise it from our conceptual scheme without loss to ordinary thought about the natural world, for the former is implicit in the latter. (162)

5. I omit various subtleties in the account. Williamson also has an account of denying that $\Box A$, which I won’t discuss.

On the face of it, however, there is no difficulty at all in combining skepticism about metaphysical modality with non-skepticism about counterfactual thinking. For counterfactual thinking, as Williamson understands it, is not all of a piece. Because one can assess counterfactuals about a given subject-matter only to the extent that one already has relevant knowledge (whether explicit or not) about that subject-matter, one's capacity to assess counterfactuals may vary from one subject-matter to another. One's capacity to assess counterfactuals is a domain-general capacity only if it's construed as a meta-capacity, i.e., a capacity to recruit whatever sub-capacities one happens to have in the service of assessing counterfactuals; and two people could both possess the meta-capacity while differing greatly in the relevant sub-capacities they possess. So, for all that Williamson says in the passage just quoted, one might possess knowledge sufficient for assessing a whole host of ordinary counterfactuals without possessing knowledge sufficient for assessing counterfactuals logically equivalent to claims of metaphysical necessity.

Williamson could respond, it seems, by saying more. He could specify the knowledge sufficient for assessing counterfactuals logically equivalent to claims of metaphysical necessity; and he could then contend that this knowledge is, or could be, had by everyone capable of "ordinary thinking about the natural world". Let us therefore ask: according to him, what knowledge *is* sufficient for assessing counterfactuals logically equivalent to claims of metaphysical necessity? I think his answer is implicit in the following passage about how, on his view, we come to know that it's metaphysically necessary that gold is the element with atomic number 79:

If we know enough chemistry, our counterfactual development of the supposition that gold is [*sc.* not⁶] the element with atomic number 79 will generate a contradiction. The reason [i.e., why the contradiction will emerge] is not simply that we know that gold is the element with atomic number 79, for we can and must vary some items of our knowledge under counterfactual assumptions. Rather, part of the general way we develop counterfactual suppositions is to hold such constitutive facts fixed. (164)

In this passage, Williamson envisages one's coming to know that it's metaphysically necessary that gold is the element with atomic number 79 *on the basis of one's prior knowledge that gold is the element with atomic number 79*. I infer that this latter knowledge is the knowledge he thinks is sufficient for assessing the counterfactual logically equivalent to the claim that it's metaphysically necessary

6. Williamson has kindly confirmed by email that the word "not" was erroneously omitted here.

that gold is the element with atomic number 79. If I'm right, then surely Williamson would wish to claim, quite generally, that the knowledge sufficient for assessing the counterfactual logically equivalent to a claim that it's metaphysically necessary that $x = y$ is the knowledge that $x = y$. Furthermore, he would presumably extend his approach to other kinds of metaphysical necessities, adding, for example, that the knowledge sufficient for assessing the counterfactual logically equivalent to a claim that it's metaphysically necessary that x had origin O is the knowledge that x had origin O .

Let me now support my contention that Williamson's account of our knowledge of metaphysical necessities doesn't take us as far as one might like. I will make four points. The first is that the account doesn't in the end vindicate the idea that non-skepticism about counterfactual thinking entails non-skepticism about metaphysical modality. The reason lies in the explanation that Williamson gives in the passage just quoted of how a contradiction emerges from the supposition that gold is not the element with atomic number 79: this explanation requires not merely that we already know that gold is the element with atomic number 79 but also that "part of the general way we develop counterfactual suppositions is to hold such constitutive facts fixed".⁷ The indispensable role of this disposition to hold constitutive facts fixed opens up at least two possibilities that invalidate the deduction of non-skepticism about metaphysical modality from non-skepticism about counterfactual thinking. The first possibility is that we might be perfectly competent to assess ordinary counterfactuals but not competent to assess the special ones equivalent to necessities of identity because we simply lack the disposition to hold constitutive facts fixed. The second possibility is that we might be perfectly competent to assess ordinary counterfactuals *and* possess the disposition to hold constitutive facts fixed, but that this disposition should not yield knowledge because there is just no such thing as metaphysical necessity. Either way, non-skepticism about counterfactual thinking could coexist with skepticism about metaphysical modality.

Second, Williamson's account of our knowledge of metaphysical necessity doesn't as yet explain our knowledge, if such it be, of the metaphysically necessary truth of those abstract principles debated by contemporary metaphysicians (e.g., mereological universalism), even though such (possible) philosophical knowledge is the kind most likely to provoke skepticism. As noted above, without the idea that "part of the general way we develop counterfactual suppositions is to hold...constitutive facts fixed", Williamson can't explain how a contradiction emerges from the supposition that gold is not the element with atomic number

7. Someone might want to put this by saying that, if we hold constitutive facts fixed as we develop counterfactual suppositions, then we embody the knowledge that things have their constitutions essentially. Doing so would not undermine my point.

79. Likewise, to explain the emergence of a contradiction from the supposition that an object which in fact had origin **O** lacked that origin, he must posit a new disposition to hold certain facts fixed—presumably a disposition to hold facts about an object’s actual origin fixed.⁸ But he doesn’t even hint as to what sort of disposition to hold facts fixed he must posit to explain our knowledge of the metaphysically necessary truth of abstract metaphysical principles.

Third, Williamson’s account of our knowledge of metaphysical necessity is incomplete, even if its scope is restricted to Kripke’s necessities of identity, origin, or constitution. The assessment procedure that Williamson’s account posits yields knowledge only if it’s reliable. However, since the procedure requires one to have a disposition to hold certain facts fixed in developing whatever supposition is in question, the procedure is reliable, I presume, only if this disposition to hold facts fixed contributes somehow to the procedure’s reliability. But how? Does it do it by corresponding, in some suitable sense, to a metaphysically modal reality? If so, what sort of a thing is modal reality? To what extent, or in what ways, is it mind-independent? And how did we come by a disposition to hold certain facts fixed that corresponds to this modal reality? I don’t say that answering such questions is necessary for rebutting skepticism about metaphysical necessity. But we should answer them if we want a full understanding of our knowledge of metaphysical modality.

Finally, even if Williamson’s account of our knowledge of metaphysical necessity is correct, we should not misinterpret its significance for philosophical methodology. Williamson is not guilty of this, but the risk of misinterpretation exists. Suppose that I come to know that $\Box A$ by going through the assessment procedure that Williamson’s account posits. Still, my coming to know that $\Box A$ in this way doesn’t provide any *evidence* that $\Box A$: it doesn’t provide me with any such evidence, it doesn’t put me in a position to present such evidence to a dissenting interlocutor, and it doesn’t provide someone who observes me with any such evidence either. Likewise, no evidence is provided if I come to believe, but falsely, that $\Box A$ by going through the assessment procedure that Williamson’s account posits. The reason is that, on Williamson’s account, whether my development of the supposition that $\neg A$ leads to a contradiction depends on whether I already believe⁹ that **A**: if I do believe that **A**, then, given that I have the dispositions to hold facts fixed that Williamson’s account posits, I will run into a contradiction; but if I don’t believe that **A**, then, given the same conditions, I won’t run into a contradiction. So the fact that my development of the supposition that $\neg A$ leads to a contradiction is only evidence that I already believe that

8. Here I just assume Kripke’s necessity of origin *arguendo*.

9. The attitude needn’t be belief. In fact, in the case where “**A**” expresses an identity claim, I think the relevant attitude isn’t belief.

A; the fact that my development of the supposition *doesn't* lead to a contradiction is only evidence that I *don't* believe that **A**. The emergence of a contradiction sheds no light on whether, given that I do believe that **A**, I believe *truly* that **A**; with regard to that question, our evidential position is unchanged. This matters in cases where "**A**" expresses an identity claim (e.g., that pain = such-and-such sensory representation of bodily damage, or that knowledge = reliably acquired true belief), and whether it's true that **A** is what chiefly interests us.