METAPHYSICS

Comments on *Overfitting and Heuristics*

RESEARCH

HARVEY LEDERMAN

ubiquity press

ABSTRACT

In the second half of *Overfitting and Heuristics in Philosophy*, Timothy Williamson defends intensionalism, the view that necessarily equivalent propositions are identical, against objections in the area of the 'problem of logical omniscience'. His defense relies on proposed psychological heuristics for making and interpreting attitude ascriptions and on the notion of having an attitude towards a proposition under a guise. These comments raise three questions about this apparatus. First, Williamson says almost nothing systematic about guises, even though they bear the weight of his account of cognitive achievement in mathematics and metaphysics. Second, he specifies no relationship between having an attitude under a guise and having it *simpliciter*, with the result that his theory makes no predictions about the truth of even ordinary attitude reports. Third, his objections to contextualism assume a connection between semantics and psychological processing that his own view appears to deny. In a final section, I argue that Williamson's own commitments about guises and communication supply the materials for a contextualist semantics of attitude reports, that his reasons for rejecting contextualism do not withstand scrutiny, and that there are good reasons, given his other commitments, to endorse this form of contextualism.

CORRESPONDING AUTHOR:
Harvey Lederman
University of Texas at Austin, US
harvey.lederman@utexas.edu

KEYWORDS:
intensionalism; guises; Frege's puzzle; logical omniscience; attitude reports; contextualism

TO CITE THIS ARTICLE:
Lederman, H. 2026. Comments on *Overfitting and Heuristics*. *Metaphysics* 9(2), pp. 1–15. DOI: https://doi.org/10.5334/met.302

DOI: 10.5334/met.302

1. INTRODUCTION

There's a saying, sometimes serious, sometimes tongue in cheek, that you shouldn't let some bad data ruin a good theory. The first half of Timothy Williamson's book, *Overfitting and Heuristics in Philosophy* (Williamson 2024) (hereafter, *O&H*), develops and defends a version of this idea, applied to philosophical methodology. Williamson argues that appeal to psychological heuristics can help us to explain away bad data in philosophy, allowing us to preserve good, strong theories. The second half of the book deploys this methodology in defense of *intensionalism* in metaphysics and semantics, the view that propositions, ordered by entailment, are isomorphic to sets of possible worlds, ordered by the subset relation.[1]

In this second half of the book, Williamson responds to several challenges to intensionalism. In Chapter 3, he defends intensionalism against various metaphysical arguments. In Chapters 4 and 5, which will be my focus here, he turns to a collection of issues in the vicinity of what is sometimes called 'the problem of logical omniscience', but which both Williamson and I would prefer to call something like 'the problem of metaphysical equivalence'. Broadly, the problem is that knowledge, belief, desire, hope, fear and other propositional attitudes (as well as the natural-language expressions for them) seem to be sensitive to distinctions beyond those intensionalism allows. In Chapter 4, Williamson confronts this problem as it arises in connection to proper names (assuming what he calls 'anti-Fregeanism' about names, but which seems to be roughly the same as what is commonly called 'Millianism'). Even though George Eliot was Mary Evans, it seems that someone could believe that George Eliot was a great writer, without believing that Mary Evans was a great writer. But intensionalism combined with Millianism implies that this is impossible. Chapter 5 confronts a second version of the problem of metaphysical equivalence, which arises in connection to apparent cognitive achievements. It seems an achievement to learn that the axioms of second-order Zermelo–Fraenkel set theory with choice imply Fermat's last theorem, even if one already knew that it is raining or it is not raining. But since, given intensionalism, these two are the same proposition, intensionalism makes it unclear what this achievement could be. Very roughly, Williamson's responses to these problems are, in the first case, to hold that while the sentences *seem* inequivalent, this is an illusion, and, in the second, to hold that the relevant cognitive achievement is explained not by our coming to know a new proposition, but by our coming to know an old proposition in a new way, or, as he puts it, under a new guise.

Section 2 of this response presents the views from Chapters 4 and 5 of *O&H* in more detail. Section 3, the core of the essay, presents three critical questions about those ideas: highlighting gaps in Williamson's discussion of guises (3.1); pointing to related gaps in his discussion of the relationship between having propositional attitudes and having those attitudes under guises (3.2); and raising doubts about his understanding of the function of semantics (3.3). Section 4 shifts gears, and argues that there is a good Williamsonian case for contextualism about attitude reports, along the lines of the proposal in Goodman & Lederman (2021) (which builds on Schiffer 1977, Crimmins & Perry 1989, and Dorr 2014).

The ideas in Section 3 were greatly improved by several conversations with Williamson before and around the presentation of these comments at the Eastern APA in 2026. I thank him for giving me a chance to develop my ideas in response to those discussions.

2. KEY IDEAS IN *O&H* CHAPTERS 4 AND 5

The positive part of Williamson's extended chapter on attitudes and attitude reports (*O&H,* Chapter 4) opens with a discussion of Kripke's case of Puzzling Pierre. Kripke observed that, in this case, natural principles about belief-ascription seem to lead *us*, the ascribers, into contradiction: we are inclined to say that Pierre believes that London is pretty and also that he does not believe that

1 Williamson's official formulation of the view is in terms of the broadest objective modality, which he hypothesizes is metaphysical modality (p. 235). Williamson often relies on claims about this objective modality that some would reject, e.g. that mathematical claims are necessary with respect to it. In what follows, I'll follow Williamson's verdicts, but note here that some of my worked examples of the predictions of intensionalism (though I think not the general points) will be sensitive to the details of one's views about this broadest necessity.

DOI: 10.5334/met.302

London is pretty. Williamson proposes that we are led into this contradiction because of our reliance on a *heuristic*, related to Kripke's 'Strengthened Disquotational Principle':

> **Strengthened Disquotational Principle (SDP):** a normal English speaker who is not reticent will be disposed to sincerely assent to 'P' if and only if he believes that P. (Kripke 1979, 113)

To recall, Kripke understands this principle to apply also to languages other than English, and crucially relies on it to derive the above contradiction. On the one hand, he observes that Pierre is *not* disposed to sincerely assent to 'London is pretty', so SDP tells us that Pierre does *not* believe that London is pretty. On the other hand, he observes that Pierre *is* disposed to sincerely assent to 'Londres est jolie', so the analogous principle for normal French speakers tells us that: Pierre croit que Londres est jolie. Since this is a truth in French, its translation into English (appealing to what Kripke calls the 'Principle of Translation') should be true, so we can conclude that Pierre *does* believe that London is pretty. But now we have derived a contradiction: Pierre both does not believe that London is pretty, and he also does believe that London is pretty.

As Williamson notes, SDP has some intuitive appeal. But it is extremely difficult, perhaps impossible, to spell out the idea behind it without inviting obvious counterexamples. Williamson proposes that this fact, together with the way that SDP leads us into contradiction, motivates taking something in the area of SDP to be a heuristic. He suggests the following simplified version:

> **BDP:** English speakers assent to 'P' if and only if they believe that P.

Understood as a *general principle*, BDP, like SDP, is subject to all sorts of counterexamples. But this is no barrier to its being a *heuristic*. Understood as a heuristic in Williamson's sense, the qualifications we would need to add to avoid counterexamples can be reconfigured as 'exception-barring clauses inserted to guard against defeaters' (p. 185). This is typical of heuristics, in Williamson's view: heuristics are shortcuts used by default, and often applied defeasibly.

Williamson's disquotational heuristic is not on its own general enough to explain all apparent puzzle cases similar to Kripke's. We can raise such puzzles for non-linguistic animals, for monoglottal speakers of other languages, and also for speakers of our own language who may not know the words we are using. To deal with these cases Williamson appeals to the notion of a *guise*. He introduces this notion for the first time as part of his own theoretical apparatus in the following passage:

> To explain why a cock robin in his own territory who sees himself in a mirror flies aggressively at the mirror, we assume that he believes (falsely) that he is a rival. To keep track of the robin's mental states, we may distinguish guises: he believes that he [under a visual demonstrative guise] is a rival, but does not believe that he [under a self-relating guise] is a rival. In effect, we treat the robin as having a belief state aptly manifested by saying 'He is my rival' but no belief state aptly manifested by saying 'I am my rival', although of course we know that the robin himself cannot manifest his belief in words. … To grasp such cases firmly, we may need to conceive them in terms of such as-if saying, and make implicit notional use of BDP. (p. 187)

In other words, although the heuristic cannot directly be applied to the robin, who does not speak a language, we nevertheless apply it to him 'notionally', implicitly treating him as if he did.

A different limitation of BDP is that it does not apply to variants of Kripke's puzzle which feature attitudes other than belief. We can concoct versions of Kripke's case where we feel a temptation on the one hand to say that Pierre *hopes* that he can live in London, and, on the other, that he does *not* hope that he can live in London (and similarly for other attitudes). To generalize his view to these other attitudes, Williamson suggests the 'Just Ask' heuristic, that is, the heuristic of (perhaps imaginatively) asking the ascribee 'Do you X that P?' (p. 188). (In the case of factive attitudes, Williamson generalizes this further to 'Do you X whether P?' (p. 189)) He does not explicitly discuss how this heuristic would extend to non-linguistic animals. But presumably the generalization would involve distinguishing guises and using a notion of 'as-if answering'. To determine whether the

DOI: 10.5334/met.302

cock-robin wants to scare off his rival, we consider whether the robin has a desire-state that would manifest in an affirmative or negative answer to our question.[2]

In Chapter 5, Williamson turns to the question of how to make sense of the cognitive achievement of coming to know new metaphysical or mathematical truths, given intensionalism. Intensionalism implies that any sentences which express necessarily equivalent propositions express the same proposition. Since most people already know that 2+2=4, it is unclear, given intensionalism (and the necessity of mathematical truths), what they would learn if they learned a 'new' mathematical theorem, for instance the truth or falsity of the Riemann hypothesis. As Williamson helpfully emphasizes here and elsewhere, this problem is not restricted to necessary truths or falsehoods; it arises for any two sentences that express necessarily equivalent propositions. (Hence my suggestion above that it would be better called 'the problem of metaphysical equivalence', than the more usual 'problem of logical omniscience'.)

In the closing substantive section of the book (5.5), Williamson proposes that the achievement in question is to come to know the relevant proposition *under a new guise*. When we learn a mathematical truth, the achievement is not primarily learning a new proposition, but instead learning an old proposition in a new way. Neo-Russellians typically say that, when Lois Lane learns that Superman is Clark, she does not learn a new proposition, but instead an old proposition in a new way. Williamson can be thought of as applying this idea to intensionalism's problem of metaphysical equivalence.[3]

There is a lot more detail in these rich chapters, some of which I'll touch on below. But this summary is enough to get us to my main questions.

3. THREE QUESTIONS ABOUT WILLIAMSON'S VIEW

3.1. QUESTION 1: WHAT ARE GUISES?

Guises play a central role at two points in Williamson's positive story. First, they are a key component of his description of the heuristics which explain our temptation to make mistaken ascriptions, as in the case of Pierre. Second, they are central to his account of cognitive achievement in mathematics and metaphysics. Surprisingly, however, Williamson does not say anything systematic about the notion of a guise, or about the related notion having-an-attitude-under-a-guise (e.g. believing that Superman can fly under the guise of the sentence 'Superman can fly'). Indeed, as we saw above, he employs the unfamiliar (and equally untheorized) notion of 'as-if-saying', and seems to require a notion of 'as-if-answering', which must apply to non-linguistic animals, among others.

---

2 Williamson rightly rejects the idea that the simple semantics he endorses for names should be defended on pragmatic grounds, by invoking some kind of Gricean mechanism (pp. 169–171). But he does not consider other prominent accounts of errors we make in assessing attitude reports (and the attitudes themselves). I would have been interested to hear about the comparison to Braun (1998). Braun (also an anti-Fregean) suggests an analogy between our assessment of attitude reports (and presumably also the corresponding attitudes) and simple sentences. Even people who know that Superman is Clark tend to (in Braun's view, mistakenly) think that 'Superman flies' and 'Clark Kent flies' express different propositions. For Braun, this is because, in assessing these sentences, people deploy guises closely associated to the sentences, and the relevant guises differ in much the way that the sentences do. Similarly, even people who know that Superman is Clark tend to (in Braun's view, mistakenly) think that 'Lois Lane believes that Superman flies' and 'Lois Lane believes that Clark Kent flies' differ in meaning. For Braun, this is because, in assessing these sentences, people deploy guises closely associated to the sentences, and these guises differ in much the way that the sentences do. At first blush, this view seems quite different from Williamson's in Chapter 4, since it does not involve the 'Just Ask' heuristic and is in the first instance presented as an account of how we assess natural language reports, not the corresponding facts. But Williamson's later view, in Chapter 5, is very close to this traditional component of 'neo-Russellianism', and looking back from Chapter 5, one wonders how much his view ultimately does differ from Braun's.

3 The old-school defenders of this view typically endorsed a structured view of propositions, which Williamson rejects. I agree with Williamson (p. 250, see also autobiographical remarks on pp. xxi–xxii) that this position is not as well-motivated by semantic data related to attitude ascriptions as it might have seemed, once we are employing guises anyway. Guises allow us to draw cognitive distinctions without drawing them at the level of content, making it unclear what role is left for structured propositions in these contexts (see also Goodman & Lederman 2021, n. 46). (I don't take a stand on the separate question of whether structured propositions could be independently motivated on metaphysical grounds.).

DOI: 10.5334/met.302

At some moments, I wondered if Williamson's reticence about guises was due to his seeing guises as only a fiction of our imagination, a product of erroneous heuristics. But Williamson clearly says that we often give *true and informative* explanations of behavior in terms of attitudes-under-a-guise (Ch. 4.12, esp. p. 227). Even more strikingly, Williamson elsewhere appeals to guises *in propria persona* to provide an account of the cognitive achievement of learning an old proposition in a new way (Ch. 5.5). He writes:

> linguistic guises are not *what* we think, and not normally what we think *of*; they are what we think *with*, when we think in words. Similarly, in speech, when you make an assertion, a guise is not *what* you assert, and not normally what you assert it *of*; it is more like what you assert it *with* (though the hearer may receive it under a different guise). We must keep track of linguistic or, more generally, representational differences, without confusing them with differences at the level of reference or content. (pp. 250–251)

Neither of these aspects of the book would make sense on the proposed idea that guises are just a product of erroneous heuristics.

An alternative reason for the reticence about guises, which Williamson suggested to me in conversation, might be that it's just not possible to give a systematic story about guises or attitudes-under-guises. We talk and think about attitudes-under-guises (though not always under that guise), and our sentences and thoughts about these facts can be true, even if there may not be anything interesting for a systematic theory to say about the nature of guises, or structural relationships among them.

If true, this would be a pretty disappointing result for the account of cognitive achievement given in the book. It's tempting to think that there's rational pressure for a typical person who learns that George Eliot is Mary Evans (having known all along that George Eliot was a great writer), to conclude that Mary Evans was a great writer. Or again, it's tempting to think that there's rational pressure for someone who learns that the axioms of second order ZFC imply Fermat's last theorem (having known all along that the axioms are true), to come to believe Fermat's last theorem. This rational pressure seems to be of a piece with other forms of rational pressure in epistemology that involve one proposition providing evidence for another. One might think that it is part of the aim of epistemology to explain what is going on here.

For Williamson, of course, the intuitive descriptions of the cases that I just gave are incorrect. The first person would not really *learn* that George Eliot is Mary Evans; they knew that all along (under a different guise). The interesting change is that they come to know this claim under a new guise. Similarly, they would not 'conclude' (for the first time) that Mary Evans was a great writer, since they knew that all along. The interesting change would involve their coming to know this claim under a new guise. (Similar points hold for the mathematical example.) These facts highlight, however, that Williamson cannot provide a content-based explanation of the felt rational pressure described above. Instead, he would naturally appeal to features of guises and believing-under-certain-guises to explain what is going on.

One natural approach, given other things Williamson says, would be to explain this by focusing on what he calls 'linguistic guises' (mental correlates of sentences) and use syntactic relationships among the sentences corresponding to the linguistic guises to explain the felt rational pressure. But, first, it's not clear that this will work, since Williamson is open about the fact that there can be multiple guises, even for a single thinker, corresponding to a single sentence. (In another famous example of Kripke's, that of Paderewski, we judge that Peter knows that Paderewski is a pianist, and also that he does not know that Paderewski is a pianist, apparently associating different guises with the very same sentence in the complement clause.) So there may not be a single 'linguistic guise' (even for a single thinker) that we can appeal to, in making sense of the relevant relationships. And second, even if this strategy did work for the linguistic case, it seems that we should hope for relevant connections that extend to non-linguistic guises, since we can concoct similar examples using spatial reasoning puzzles or various non-linguistic sounds.

I've focused here on 'rational pressure', but we can raise the challenge for many other topics in epistemology: for instance, memory, testimony, and learning. To see this, consider a damning

DOI: 10.5334/met.302

challenge that Williamson presents for an individualistic Frege-inspired view of senses (Chapter 4.5, pp. 174–177). On the view Williamson considers there, different individuals associate different senses with different sentences of natural language. Williamson observes that this view does not yield a systematic story about what happens if a person 'merges' her senses of a given individual or property. If someone who (intuitively) does not know that furze is gorse is considering whether to come to believe that furze is gorse, she is considering, from the inside, dropping the two senses she associates with these different words, and instead using a single new sense which would correspond to both. But the relevant sort of individualistic Fregeanism makes it unclear whether this thinker can even entertain the thought that she would thus come to stand in the Fregean belief-relation to.

I was nodding along at the concluding passage of this discussion:

> The required theoretical apparatus of individualistic Fregean senses is elaborate, unclear, and ill-developed. It is not useful for…epistemology, given the problem of incommensurability of senses between different cognitive states. After all, a significant part of epistemology concerns the communication of knowledge across such differences—by memory, from earlier to later states of the same individual, and by testimony, from states of one individual to states of another. Learning itself—the acquisition of knowledge—constitutes a change in cognitive state. What gets preserved through such transactions is likely to be at least somewhat coarse-grained. (p. 177)

These are important points. But they seem to apply just as much to Williamson's own story about cognitive achievement with respect to necessarily equivalent propositions as they do to the Fregean. On Williamson's view, in such cases, the relevant forms of 'memory' and 'learning' can't be understood only in terms of preservation of the same coarse-grained content (which one may know under other guises). The interesting questions concern instead whether the person continues to know a proposition under the same guise, or comes to know it under a new one. But, to ask Williamson's question: what is a person contemplating exactly, when they are contemplating (as in the example above) changing the set of guises that they use to 'think *with*'?

Williamson's quietism about guises and attitudes-under-guises leaves us with no answer to these questions. We can't make sense of the kind of rational pressure described earlier, or of how memory, testimony and learning work in an important class of cases. Williamson does not tell us almost anything about the conditions under which one person is thinking with the same guise as another (or if that is even possible). He does not tell us under what conditions the same person is thinking with the same guise at different times. He does not tell us under what conditions guises are related in such a way as to generate felt rational pressure. He does not tell us whether there are experiments we could do to clarify structural facts about guises, or whether they are the type of thing best studied from the armchair.

These questions are not primarily about the *nature* of guises; I am not asking for an *analysis* of them. They are instead about structural constraints on the notion of guises, which might help to show how this notion operates within the theory, and thus give us a better purchase on it (compare the discussion of safety in Williamson 2000). It would be answered, for instance, by having a model which illuminates their application to cases of interest.

3.2. QUESTION 2: WHAT IS THE CONNECTION BETWEEN HAVING AN ATTITUDE AND HAVING THAT ATTITUDE-UNDER-A-GUISE?

A natural way of understanding Williamson's locution 'believing-under-a-guise' would be that if a person believes a proposition under a guise *g*, then they believe that proposition simpliciter. After all, if a person has a heart attack under a full moon, they have a heart attack *simpliciter*. If a person screams profanities under a bridge, they scream profanities *simpliciter*. The same pattern holds if the actions occur under false pretenses, under an awning, or for that matter under the great blue sky.

DOI: 10.5334/met.302

But, though he does not say so in so many words, Williamson does not understand the relationship between believing a proposition-under-a-guise and believing a proposition in this way. In the book, he is explicit about the lack of such a relationship in the case of credences (p. 213), but not explicit about his views in the case of belief and knowledge. In fact, in the book, he sometimes sounds as though he believes there is at least a somewhat reliable connection here for belief and knowledge. For instance, in discussing the case of Penny, who learned the word 'furze' from samples of the shrub in flower and, separately, the word 'gorse' from samples not in flower, he writes:

> One can know something by knowing it under one linguistic guise without knowing it under all its linguistic guises. Penny has the knowledge attributed in (17) ['Penny knows that her evidence includes that furze has yellow flowers'] under the guise of sentence (19) ['My evidence includes that furze has yellow flowers'], even though she does not have it under the guise of sentence (20) ['My evidence includes that gorse has yellow flowers'], and that seems enough for the truth of (17). (p. 204; bracketed insertions give the sentences to which Williamson's numbers refer)[4]

And later on, he writes:

> In the case of belief, we may consistently suppose that one believes a proposition p if and only if one believes p under at least one guise; such a guideline was tacitly followed in the discussion of Frege cases above. (p. 213)

But through our correspondence, I've come to understand that, while Williamson thinks this is a *consistent* supposition, he does not endorse it. In fact, Williamson is open to the idea (though not committed to it) that there is *no systematic relationship* between Φ-ing that *p* under a guise and Φ-ing that *p simpliciter*. In his view, the former is needed to characterize our *heuristics* for mind-reading (and also, as we saw, some true and informative action explanations, as well as the cognitive achievement of coming to know old propositions in new ways), but these facts may not be connected to what attitudes we have directly. (It is his quietism about the relationship between these two notions that drives the striking lack of a semantic clause for attitude verbs in Chapter 4. Williamson's view is that attitude verbs can be given a simple possible-worlds semantics in terms of what attitudes the person has *simpliciter*, and without any reference to what attitudes they have under various guises.)

A consequence of this picture, as I've come to understand it, is that a variety of standard cases will be undecided by Williamson's theory. For instance, his theory does not predict either that Pierre believes that London is pretty or that Pierre does not believe that London is pretty. Similarly, to take another of Kripke's examples, it does not predict that Peter believes that Paderewski is a pianist or that Peter does not believe that Paderewski is a pianist. These results are perhaps not so surprising. The cases are hard cases, and one might even think it a virtue of a view not to make predictions about them.

But there are more surprises to come. Williamson's theory in the book also does not make a prediction as to whether the Greeks believed that Hesperus rose in the evening. They did believe this proposition under one guise. But they also failed to believe it under another. Williamson does not give us guidance on how to move from these facts to what they believed *simpliciter*.

My first reaction to realizing all this was to suspect that the same issue would arise for necessary truths. After all, these also are claims that people tend to know or believe under one guise and fail to know or believe under another (for instance, I know that it is raining or it is not raining under the guise of the sentence '2+2=4', but I do not know that fact under the guise of the sentence 'the Riemann Hypothesis is true'). So it might seem that we would then be in a position of not

4 In other passages Williamson is similarly cautious: 'In the scenario, ["Penny doubts that her evidence includes that gorse has yellow flowers"] **seems** true, because she doubts the proposition under the guise of sentence (20), even though she does not doubt it under the guise of sentence ["Penny doubts that her evidence includes that furze has yellow flowers"].' (p. 205, emphasis mine); 'In this scenario, ["Penny believes that her evidence does not include that gorse has yellow flowers"] **seems** true because Penny believes the proposition under the guise of the negation of sentence (20), even though she does not believe it under the guise of the negation of sentence (19).' (p. 206, emphasis mine). In both passages, bracketed insertions replace Williamson's sentence numbers, except for (19) and (20), which appear in the main text above.

DOI: 10.5334/met.302

knowing whether people know/believe these claims. But in discussion, Williamson emphasized to me that he does view these cases as settled. To settle them, he appeals to the predictive power of normal modal logics for belief and knowledge. These theories include the claim (often as an axiom) that agents *do* know/believe the necessary truths. Endorsing this powerful theory for people would therefore settle that we too believe the necessary truth. (Presumably Williamson's commitment to normal modal logics would also commit him to the claim that we know all consequences of what we know, as this is also included in normal modal logics.)

This move eliminates one sort of gap in Williamson's theory. But it doesn't remove others. It seems plausible that I know that the sky is blue. But suppose that the Riemann hypothesis is true, but (as seems natural) that I do not know this under the guise of 'The Riemann Hypothesis is true'. Then plausibly I do *not* know that [the Riemann hypothesis is true and the sky is blue] under the guise of the 'The Riemann Hypothesis is true and the sky is blue'. The problem is that, given intensionalism, the proposition I fail to know under that guise *is* the proposition that the sky is blue (since, given intensionalism, the necessary truth conjoined with any proposition is just that proposition). So I know that very proposition under one guise and fail to know it under another, and Williamson's theory does not tell us whether I know it *simpliciter*. In the book, Williamson provides a nice explanation of why we feel the temptation to judge related contradictory reports ('I know that the sky is blue', 'I do not know that the Riemann Hypothesis is true and the sky is blue') to be true: the generalization of 'Just Ask' to knowledge explains just that. But again he does not tell us which of the facts obtains: do I or do I not know that the sky is blue? The point is perfectly general; the book seems to make no predictions about the truth or falsity of just about any report of an attitude towards a contingent claim.

There are several flat-footed ways of developing a more predictive theory. One option, inspired by Salmón (1986), would follow the standard logic of 'under' in English, and hold that a person *X*s that *p* if and only if they *X* that *p* under some guise. This theory would make strong predictions in the cases above. It predicts that the *positive* reports are true and the negative ones are false: Pierre believes London is pretty, Peter believes that Paderewski is a pianist, the Greeks believed that Hesperus rose in the evening, and I believe that the sky is blue. These are fine predictions, as such things go. But (as has been much discussed in the last forty years) the theory also makes bad, or anyway pretty surprising predictions. For instance, it predicts that Lois Lane knew that Clark Kent could fly all along. In conversation Williamson has rejected it.[5]

A different theory, still more predictive than full-blown quietism, would hold that for any contingent *p*, whenever a person *X*s that *p* under some guise, and fails to *X* that *p* under another guise, it is borderline or indeterminate whether they *X* that *p*. This view would still make predictions in cases where Williamson does not. In particular, it would *predict* that it is borderline or indeterminate whether Pierre believes that London is pretty, indeterminate whether Peter believes that Paderewski is a pianist, indeterminate whether the Greeks believed that Hesperus rose in the evening, and indeterminate whether I believe that the sky is blue. I do think that making some prediction is, ultimately, important, but it is not clear to me that this view would be better or for that matter very different, since it implies that, while there are facts of the matter, these facts are unknown or even unknowable to us.

My main goal in this subsection is just to observe that the book is silent on such questions, which is surprising given Williamson's emphasis on strength as a virtue of theories. It would be helpful to see his picture developed in more detail, so that we could assess its predictions in such cases.

I suspect that Williamson might feel that the demand for more detailed development is unfair. Lexical entries for propositional attitudes are often given disquotationally, without engaging with

5 The Salmón-inspired view has some interesting consequences if we extend it to include Williamson's background commitments. For instance, in Williamson's classic discussion of evidential probability, he defines it as the probability of the unique prior, conditioned on the conjunction of everything you know (Williamson 2000: 221–222). (In the present work, he calls this 'epistemic probability'.) But we can generate a problem for this view if we allow certain 'cheap' guises. Consider for instance 'Every proposition is true if and only if actually true'. By the reasoning above, this sentence expresses the conjunction of all truths. Moreover, it is guaranteed to be true, so, if there is an associated guise, we can presumably know the conjunction of all these actual truths under this guise. On the Salmón-inspired view, we would then know that conjunction *simpliciter*, and every truth would have evidential probability 1.

DOI: 10.5334/met.302

the details of what will be predicted in particular examples. So, Williamson might say that *no* existing theory *predicts* that ordinary people believe that the sky is blue. In a sense this is true. But it is based on an uncharitable reading of the usual practice, which makes the usual practice quite different from what Williamson is doing here. Usually such lexical entries are given in a context where Frege puzzles and the problem of metaphysical equivalence are deliberately set aside. There is an implicit, background goal of respecting intuitive judgments about natural language sentences as much as possible. When semanticists and philosophers *do* turn their attention to these more difficult cases—the problems that Williamson considers here—much, much more detail is given about how such cases will be handled (as in the case of the Fregeans that Williamson criticizes, or in the case of Salmón or Braun). Williamson's view by contrast presents serious challenges for the truth of 'normal' attitude reports in even typical cases; he is clear that we cannot simply follow intuitive judgments where they lead. But this is precisely why we should hope for a more predictive story from him about when relevant reports are true. Presumably such a theory would take the form of a structural characterization similar to (say) the notion of similarity in the theory of counterfactuals, that might help to connect Φ-ing that *p* to Φ-ing that *p* under a guise. Even better would be (again) if this theory were connected with a model of guises that allows us to draw out (simplistic) predictions in particular cases.[6]

Williamson's proposals in these chapters seem *much* simpler than the overwrought proposals he criticizes. But the book does not provide a systematic theory about one of its foundational notions (guises), nor does it make predictions about what attitudes people have in even basic cases, in spite of itself presenting serious reasons for doubt about our ability to know what attitudes they have. It may be that Williamson's theory will be simpler and more predictive than the alternatives. But the book does not contain a sufficiently detailed proposal for us to know whether this is true. And there is a lurking fear that Williamson's account only seems simpler because it does not offer answers to hard questions that other approaches attempt to resolve.

3.3. QUESTION 3: HOW DOES SEMANTICS RELATE TO PSYCHOLOGICAL PROCESSING?

Williamson is at pains to keep semantics simple, even when the resulting story about psychological processing, in terms of guises, may be quite complicated. This often means that Williamson's semantics looks pretty different from what he suggests people do psychologically in processing utterances. But Williamson is clear that he thinks semantics has an important connection to psychological processing. Indeed, part of his argument for his theory, and against alternatives, is that the alternative views will impose overly large computational costs on speakers. For instance, he writes:

> In semantics, apart from the usual need to avoid overfitting, there is the additional concern that an over-complicated semantic framework may impose infeasible computational burdens on ordinary speakers. In contemporary semantics, one often sees very complex semantic accounts presented with no apparent sense that their complexity might be a theoretical cost. (p. 83)

Here, the complexity of semantic entries is claimed to directly correspond to the 'computational burdens' imposed on speakers and hearers, suggesting an important relationship between semantics and linguistic processing. Williamson elsewhere argues against contextualism on related grounds:

6 Williamson himself seems to recognize the need for such a model. In an important passage (p. 210), he says that it is often valuable to model those subject to a Frege puzzle by including worlds where there are distinct individuals for each of the guises the subject possesses (even if they do not corefer). But he does not go on to accept a version of impossible-worlds semantics that allows this variation. I would have been interested to hear more about why we should reject this limited form of hyperintensionalism, which does not fall to Williamson's objections in Chapter 3 (it preserves compositionality along with a simple semantics for logical connectives; see e.g. Lederman 2021). A theory which uses a single class of models for all cases has an explanatory advantage over one which has an official library of models and then appeals to unofficial models in an unsystematic way to give explanations in cases where the official ones give out. This form of limited hyperintensionalism can give unified, content-based explanations of more behavior than Williamson's preferred option.

DOI: 10.5334/met.302

> contextualist strategies open up myriads of new error-possibilities for speakers who do not carefully store lots of information about the contexts in which they originally acquired linguistically encoded information. Speakers unaware of such contextualist features of the semantics of their language will be especially liable not to do the hard work of storing all that information. (p. 48)

I found these comments puzzling, in light of Williamson's own views about linguistic processing as they emerge later in the book. In Chapter 4, the BDP heuristic and the related 'Just Ask' heuristic are apparently claimed to explain the psychological processing of relevant expressions. Semantics appears completely psychologically inert in these cases. I wondered if the above remarks were intended merely as an *ad hominem* against those who hold views about semantics on which it has a more intimate connection to cognitive processing. But if so, it would have been helpful to hear why Williamson himself (not taking on board these premises adopted *ad hominem)* opposes contextualism or a more complex semantics. Perhaps the idea is that, since semantics is effectively idle, we have a relatively free hand to adopt a simple powerful theory here.

But, setting these points aside, I also found these comments unsuccessful as an *ad hominem*. I find contextualism quite attractive. But Williamson's comments were a *non sequitur* with respect to my favored view of semantics. In my view, formal semantics does not directly seek to provide characterizations of how, cognitively, information is stored or processed. For me, it aims to describe truth-conditions and entailment relations, without taking a stand on the question of how information about these facts is stored or processed. I would say (and I don't think the view I'm describing is particularly idiosyncratic) that speakers tend to (implicitly) know the facts captured by a good semantic theory. But in my view questions about the ease or difficulty of processing concern implementation at a much lower level of description (e.g. neural, circuit-level) than the one described by the propositional attitude 'know'. So even granting that we know these facts about entailment and truth, I think we can (and should) mostly leave questions about processing to one side while doing semantics, and focus on what makes for good descriptions at the level of propositional attitudes.

My main question here, then, has two parts. First, how does Williamson himself view the relationship between semantic values and psychology given that semantics seems psychologically inert for him? Are his comments here just intended *ad hominem* or do they capture his own views? Second, can these comments be patched up to still have force against my sort of view about the function of semantics? Or do they leave such views untouched, as I have suggested they do?

This question is related to a broader one about Williamson's overall approach. At times it seems that he sees complexity in a theory as primarily a problem when it arises in a semantic theory, and less of a problem for a psychological theory. He criticizes semantic theories for lacking strength, but seems less concerned if a psychological theory is complicated or lacks predictive power—a charge that could be easily raised against the theory of guises, the use of separate heuristics for belief as opposed to other attitudes, and his use of a further heuristic for factive attitudes. This raises the question: why should complexity be a special problem in semantics, especially once the psychological story is complicated enough to easily make sense of more complex semantic values?

4. A WILLIAMSONIAN CASE FOR PERSPECTIVISM

In this last section, I turn to an 'in-house' debate. As I've said, Williamson's main goal in these chapters is to defend intensionalism (together with anti-Fregeanism about proper names). At times, he is *officially* open to a contextualist view about attitude reports. But it is clear that, in his heart, he is not open to any such thing. In fact, he argues against contextualism in several places. Here, however, I want to articulate what is, in my view, a strong Williamsonian case for contextualism.

In 'Perspectivism', Jeremy Goodman and I sought to develop and defend a form of contextualist anti-Fregeanism (or, as we put it there, 'Millianism'). *Perspectivism*, as we call it, endorses anti-Fregeanism about names, together with the following lexical entry for 'believe' (at an appropriate

DOI: 10.5334/met.302

level of abstraction):

$$[\![\text{believes}]\!]^c = \lambda p.\lambda x.\lambda w.\text{ there is a mentalese sentence } s \text{ that expresses } p \text{ which is an element of } \pi(c) \text{ and which is in } x\text{'s 'believe'-box in } w.$$

Here c is meant to be a parameter that gathers relevant features of context, and $\pi(c)$ is a contextually supplied set of mentalese sentences. For present purposes, when we say 'mentalese sentence' you can read 'guise' in Williamson's sense.

More generally for other attitude verbs 'Φ', perspectivists hold:

$$[\![\Phi]\!]^c = \lambda p.\lambda x.\lambda w.\text{ there is a mentalese sentence } s \text{ that expresses } p \text{ which is an element of } \pi(c) \text{ and which is in } x\text{'s '}\Phi\text{'-box in } w.$$

Both of these entries appeal to the notion of a 'box' for mentalese-sentences/guises. Instead of 'boxes', Williamson speaks of 'Φ-ing p under a guise g'. This language *looks* less technical, but it's not really: the 'folk' don't explicitly use the notion of Φ-ing under a guise anymore than they use the notion of a Φ-box. And like Goodman & Lederman (2021), on a first pass, Williamson is (tacitly) assuming that each attitude verb has its own box. Given all this, we could equivalently put our entry in Williamsonian language as:

$$[\![\Phi]\!]^c = \lambda p.\lambda x.\lambda w.\text{ there is an element of } \pi(c), g, \text{ and } x\ \Phi\text{'s that } p \text{ under } g \text{ in } w.$$

Which is all to say that really the only difference here between us and Williamson is that we have a context-sensitive, but systematic relationship between Φ-ing that p under g and just Φ-ing that p, where Williamson does not specify what the relationship is.

Williamson's primary objection to contextualist theories of attitude verbs like ours is familiar from his earlier work (cf. Williamson 2005; Williamson 2020, 82). There he argues that, given contextualism, disquotation will no longer be a good way of remembering what was said (since in the new context, repeating 'John knows it's raining', say, might mean something different). Like many others, I'm not moved by this argument. First, it is unclear to me whether Williamson rejects the context-sensitivity of words like 'tall' or 'rich' on similar grounds. If Williamson accepts *these* instances of context-sensitivity, he must think that we have another way of remembering what was said in these cases, and it is not clear to me why the relevant method would be unavailable in the case of attitude reports. (In *O&H*, on p. 49, he suggests that these cases could become more palatable if there's a 'default' resolution of the context-sensitivity, but he doesn't consider the possibility of such a default for attitude verbs.)

Second, Williamson's view of the psychology of communication *already* relies on a flexible relationship between the words used and the guises that are relevant in a given circumstance. To see this, note that Williamson's story about cognitive achievement in metaphysics and mathematics naturally goes along with a view of how communication works when we use sentences that express propositions necessarily equivalent (and hence identical) to propositions the hearer already knows under other guises. In typical cases, such communication will be successful if and only if the hearer comes to know the proposition under the guise of the sentence used. But Williamson explicitly acknowledges that this default will fail in cases which feature context-sensitive expressions (p. 253). For instance, in an example from Kaplan (1989) where a person asks themselves (gesturing twice) whether 'that ship is the same as that ship', the guises operative for the demonstratives must be different in each occurrence.

Williamson does not acknowledge this further point, but the flexibility in which guise is associated with which sentence cannot be restricted to obviously context-sensitive sentences of English. Kripke's Peter can come to know that Paderewski (the pianist) is a politician by hearing an appropriate utterance of 'Paderewski is a politician', a sentence he himself was already inclined to assent to in other circumstances, prior to the revelation. This same sentence must thus be able to correspond to a guise different from the one he deployed via the sentence before, *even though the sentence is not obviously context-sensitive*.

DOI: 10.5334/met.302

The point seems to apply to other cases as well. Williamson motivates the introduction of guises on the basis of the example of the cock-robin, a case which is broadly parallel to that of Puzzling Pierre. In the latter case, we can convey something interesting by saying both 'Pierre thinks that London is pretty', and 'Pierre does not think that London is pretty' (as well as by 'Pierre thinks London is not pretty'). By Williamson's own lights, then, it must be that our heuristics for assessing these sentences reference distinct guises related to the complement clause and in particular to the name 'London'. So again we have a flexible association of guises with sentences. Goodman and I developed another such case, sharper in certain ways, in Goodman & Lederman (2021, 626), building on the 'Thelma' case of Schiffer (1977). But there are many relevant cases once one starts to look. Consider Bruno the dog and his owner Dave:

> Bruno knows that Dave is friendly. But, one day, hearing Dave make noise on his way in, Bruno begins barking, because he doesn't know that Dave is friendly.

A natural extension of Williamson's discussion to this case would hold that, in reading the vignette, we distinguish different guises of the proposition that Dave is friendly, and make implicit notional use of BDP, imagining off-line that Bruno as-if says sentences associated with these guises. But this procedure once again requires that we associate *different* guises with the sentence 'Dave is friendly'.

These examples will be important again below. But the main point to draw here is that in all of them, the natural Williamsonian story about successful communication will come apart from the disquotational procedure. What we take up from these utterances will in part be the associated guise (which won't be fixed by the mere words used). Thus Williamson already seems committed to something that goes beyond the disquotational procedure.

My third and most important reason for being unmoved by Williamson's emphasis on disquotational memory is that it sits oddly with somewhat obvious facts about our psychology (an oddity that is all the more striking given Williamson's welcome emphasis on the psychology of utterance-processing in *O&H*). It is just not a plausible theory in cognitive science that we store the information obtained from utterances by storing the words said. We don't need fancy studies to show this: in many cases it's possible to remember what someone said without remembering the exact words they used (I bet you can remember the idea of the Bruno case from a moment ago, but I bet you can't remember the form of words used there).

So this first line of attack against contextualism doesn't move me, really at all. A second issue with contextualism, from Williamson's perspective, is that it is too complicated. But I'm unclear on why this complexity in the semantics is so bad. First, as I said before, I'm not convinced that contextualism does really impose unreasonable or even particularly significant computational demands on speakers and hearers. There may be computationally low cost ways of storing information, which are hard to characterize in a perspicuous format in formal semantics. Second, and more importantly, Williamson's view about the psychology of communication already has basically enough structure to make contextualism go. As we have seen, in a wide array of not obviously context-sensitive expressions, Williamson should hold that sentences are flexibly associated by speakers and hearers with different guises in different settings. In other words, he already thinks that there is a (partial) mapping between sentences and (pluralities of) guises in a context. This is basically what we need for contextualism.

By comparison to Salmón's invariantist theory discussed under question 2 in the last section, the only extra ingredient required for Goodman and my theory, is a restriction on which guises are salient in which context. This is a version of contextualist quantifier domain restriction, which Williamson seems open to in other work (Williamson 2003). This form of contextualism does not seem to me to complicate the theory significantly, since it involves a mechanism we must already postulate.[7]

7 In fact, as I mentioned, Williamson says that contextualism is less costly if it comes with a 'default' interpretation of the relevant context-sensitive vocabulary (p. 49). And the kind of contextualism that we develop could be paired with such a 'default'. One might take the default to be the reading on which every *natural* guise is salient, setting aside

DOI: 10.5334/met.302

Going beyond these objections, a third objection Williamson might make to contextualism for attitude reports is that it does not make predictions about when utterances are true. Williamson might correctly observe that since contextualism allows both true and false readings of certain sentences, it is only when paired with a further theory about which readings will be operative for a given use, that it makes predictions about particular utterances. I agree that this is an important challenge for contextualism in general and I'll say more about it in a moment. (Indeed, this objection can be seen as a 'formal' version of a charge I levied against Williamson in the material mode in my 'second question' above, where I faulted his theory for lack of predictiveness.) But I think even without giving that further theory, contextualism of the kind just described makes many more predictions than Williamson's theory as it stands. For contextualists are happy to accept ordinary speaker judgments about sentences where possible and to use them as evidence about which readings must be operative for which utterances. So, for instance, there is no problem in allowing for 'I believe that the sky is blue' to express a truth. In other words, the challenge for the contextualist is not as powerful in my view as the challenge I raised for Williamson earlier. Since Williamson casts doubt on speaker judgments and then provides a theory which does not settle in a principled manner which of those judgments can be accepted, he is left without a default prediction in even 'normal' cases. The contextualist does not have this problem.

At the very high level of description I've given so far, Williamson could accept perspectivism and preserve a closer connection between the judgments of speakers and what sentences communicate and express, than he is otherwise able to. Indeed, for all I have said so far, he could do so while respecting *all* of the usual Fregean judgments.

But Goodman and I went further than this abstract specification of the view, adding more constraints on our form of contextualism. In the first instance, we were concerned that, given a completely flexible version of contextualism like the one above, natural principles about propositional attitudes would not come out as true even when uttered in a single context (Goodman & Lederman 2021: Section 4). We sought to develop a form of the theory which would vindicate the claim that instances of such principles are true when assessed holding context fixed. It is not clear to me to what extent Williamson would be moved by these considerations, since in the present work he seems happy to rely ultimately on guise-theoretic explanations when natural-language reports give out (see especially Chapter 4.12). If it is not a concern to yield a theory with a strong 'logic' of natural language attitude reports, Williamson could stop with the relatively unconstrained version of contextualism above (or with the slightly more constrained theory we present in Section 9.1 of our paper). But if it is a concern, perhaps he would be drawn to the idea that we should go further.

Second, we were concerned to say something about which perspective is operative for particular utterances (addressing the challenge, mentioned above, for contextualists to say something predictive about which context-sensitive utterances are true or false in which contexts). Using a toy model of guises as sentences in a mental language, we proposed that the set of guises which are salient in a given context is generated by a set of elementary mentalese expressions including the mentalese correlates of expressions in the complement clause of the ascription.

The theory that we produced as a result of these demands, against the background of a toy (but highly constrained) model of guises, has several virtues. First, it preserves the structural laws about belief and knowledge that would be preserved by the Salmónian semantic clause described in section 3.2, which are, in my view, pretty significant. And, second, it allows us to typically give content-based explanations of communication and action without explicitly mentioning guises. It's true that, like many strong theories, ours delivers surprising results. Like Williamson, we must attribute error to speakers in some cases. In particular we deny that there are relevantly informative true readings of negated attitude reports whose complement clauses contain coreferential expressions associated with distinct guises, such as 'Lois doesn't know that Superman is Clark Kent' (Section 11 of our paper). This clearly delimited class of sentences is far

unnatural guises like the one suggested by 'every actual truth is true', and taking as given an intuitive distinction between natural and unnatural guises.

smaller than the enormous class of sentences in which Williamson (or neo-Russellians) attribute error to normal speakers, but it is still some error.

Williamson dismisses our paper in the book on the grounds that a form of contextualism which allows for attribution of error smells like overfitting (p. 171). He observes that, if the mechanism which explains the error is the same in relevant cases, then attributing 'more' error is not costly. In the abstract, I agree. But since, as I have argued, contextualism itself should not be thought of as especially costly, I see matters quite differently. Overfitting is certainly a problem. In fact, our concern to avoid overfitting—our goal of preserving a simple theory of the parameters of context sensitivity, while also preserving a strong 'logic' of propositional attitude reports—is precisely what led us to accept that speakers commit errors in some cases. But *under*fitting is also a risk. Williamson's theory of the psychology of linguistic processing, together with a commitment to domain restriction, already provides the materials to make sense of contextualism. Why attribute *extra* error when one doesn't have to? Williamson can't respond that his theory is more predictive about which attitudes a person has simpliciter, since it isn't.

Indeed, we arrived at the verdict that speakers commit error precisely through the hard work of building a model and deriving its predictions systematically in a wide array of cases, something that Williamson spends much of his book praising in the abstract. We strove to build a theory that balanced strength, simplicity, and fit to the data. We derived the consequences of that model in a simple, systematic way. If the result seems complicated, it is because we did not adopt a quietist stance, and strove to answer questions that Williamson leaves unaddressed in the book.

In sum, 'error plus contextualism' arises from a reasonable way of taking into account the different desiderata of avoiding both overfitting and underfitting. And, even for those who disagree, one of the other contextualist proposals that we discuss in the paper might still be better than what I believe is the *over*simple proposal Williamson leaves us with in the book.

ACKNOWLEDGEMENTS

Thanks to two anonymous referees for this journal, Kyle Blumberg, Cian Dorr, Dan Greco, Ben Holguín, Daniel Kodsi, Matt Mandelkern, Daniel Rothschild, and especially Jeremy Goodman and Timothy Williamson for comments and conversations.

COMPETING INTERESTS

The author has no competing interests to declare.

AUTHOR AFFILIATIONS

Harvey Lederman orcid.org/0000-0001-9438-097X
University of Texas at Austin, US

REFERENCES

Braun, D. 1998. Understanding Belief Reports. *The Philosophical Review*, 107(4): 555–595. DOI: https://doi.org/10.2307/2998375

Crimmins, M. and **Perry, J.** 1989. The Prince and the Phone Booth: Reporting Puzzling Beliefs. *The Journal of Philosophy*, 86(12): 685–711. DOI: https://doi.org/10.2307/2027014

Dorr, C. 2014. Transparency and the Context-Sensitivity of Attitude Reports. In: García-Carpintero, M & Martí, G (eds.) *Empty Representations: Reference and Non-Existence*. Oxford: Oxford University Press. pp. 25–66. DOI: https://doi.org/10.1093/acprof:oso/9780199647057.003.0002

Goodman, J. and **Lederman, H.** 2021. Perspectivism. *Noûs*, 55(3): 623–648. DOI: https://doi.org/10.1111/nous.12322

Kaplan, D. 1989. Demonstratives. In: Almog, J, Perry, J & Wettstein, H (eds.) *Themes from Kaplan*. New York: Oxford University Press. pp. 481–563.

Kripke, S. 1979. A Puzzle about Belief. In: Margalit, A (ed.) *Meaning and Use*. Dordrecht: Reidel. pp. 239–283. DOI: https://doi.org/10.1007/978-1-4020-4104-4_20

Lederman, H. 2021. Fine-Grained Semantics for Attitude Reports. *Semantics and Pragmatics*, 14(1): 1–45. DOI: https://doi.org/10.3765/sp.14.1

Salmón, N. 1986. *Frege's Puzzle*. Cambridge, MA: MIT Press.

Schiffer, S. 1977. Naming and Knowing. *Midwest Studies in Philosophy*, 2: 28–41. DOI: https://doi.org/10.1111/j.1475-4975.1977.tb00026.x

Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press. DOI: https://doi.org/10.1093/019925656X.001.0001

Williamson, T. 2003. Everything. *Philosophical Perspectives*, 17(1): 415–465. DOI: https://doi.org/10.1111/j.1520-8583.2003.00017.x

Williamson, T. 2005. Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge. *The Philosophical Quarterly*, 55(219): 213–235. DOI: https://doi.org/10.1111/j.0031-8094.2005.00396.x

Williamson, T. 2020. *Suppose and Tell: The Semantics and Heuristics of Conditionals*. Oxford: Oxford University Press. DOI: https://doi.org/10.1093/oso/9780198860662.001.0001

Williamson, T. 2024. *Overfitting and Heuristics in Philosophy*. New York: Oxford University Press. DOI: https://doi.org/10.1093/oso/9780197779217.001.0001

TO CITE THIS ARTICLE:
Lederman, H. 2026. Comments on *Overfitting and Heuristics*. *Metaphysics* 9(2), pp. 1–15. DOI: https://doi.org/10.5334/met.302

Submitted: 01 March 2026
Accepted: 05 October 2026
Published: 26 October 2026

COPYRIGHT:
© 2026 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See https://creativecommons.org/licenses/by/4.0/.

Metaphysics is a peer-reviewed open access journal published by Ubiquity Press.