METAPHYSICS

Comments on Timothy Williamson's 'Overfitting and Heuristics in Philosophy'

RESEARCH

KIT FINE

]u[ubiquity press

ABSTRACT

I respond to some objections that Timothy Williamson makes to truthmaker semantics and the structuralist conception of propositional identity.

CORRESPONDING AUTHOR:
Kit Fine
NYU, US
kf14@nyu.edu

KEYWORDS:
intensionalism; hyperintensionalism; proposition; truthmaker; semantics; structuralism; identity; Abstraction; fundamentality; definition; williamson

TO CITE THIS ARTICLE:
Fine, K. 2026. Comments on Timothy Williamson's 'Overfitting and Heuristics in Philosophy'. *Metaphysics*, 9(2), pp. 33–44. DOI: https://doi.org/10.5334/met.305

DOI: 10.5334/met.305

Timothy Williamson's book is, as usual, a panoply of riches; and he presents those riches with his characteristic brilliance and verve. His general line is that heuristics that we adopt can sometimes lead to faulty data and that philosophers, in a misguided attempt to accommodate some such data, have been prone to develop theories that contain too many degrees of freedom to be of genuine explanatory value. One of his central illustrations of this tendency comes from the so-called 'hyperintensional revolution'. He believes that recent hyperintensional theories have been largely motivated by faulty data of this sort and that, once we recognize that this is so, we will no longer have good reason to accept such theories or to abandon the intensional framework with which they compete.[1]

Since my name has sometimes been associated with the hyperintensional revolution, I thought it might be helpful for me to say something by way of its defense. I shall focus in particular on sections 3.4 and 3.5 of Williamson's book in which he discusses truthmaker semantics and the earlier, structural, approach to propositional content of Bertrand Russell. I believe that Williamson raises very real issues in each of these cases but ones that are not as devastating to the proposed semantics as he seems to think. They call, not for an abandonment of the semantics, but greater care in how they are formulated.

Let me begin with truthmaker semantics. This semantics, at least in one of its incarnations, is bilateral. It states separate clauses for when a sentence is made true and for when it is made false; and it is not in general possible to understand what the truthmakers of a sentence are in terms of its falsitymakers or its falsitymakers in terms of its truthmakers. This leads Williamson (2024: 132) to make the following objection:

> 'The quasi-independence of verifiers and falsifiers clearly adds a degree of freedom to the models, by comparison with classical semantics, which in effect equates falsity with non-truth for sentences. This aspect of the semantics is most striking for atomic sentences. Imagine that you are being taught a foreign language. Your teacher explains to you exactly what would make a given atomic sentence true. Could you then complain to her: "You've only done half your job! You've told me what this sentence's verifiers are, but you haven't told me anything about its falsifiers"? That sounds quite unreasonable. Your teacher has already done enough to enable you to understand the sentence, by normal linguistic standards. You are not missing half its meaning. Normally, what would make a sentence true does determine what would make it false.'

The point is that in understanding what makes a sentence true I should normally be in a position to understand what makes it false. But if, as in the bilateral version of truthmaker semantics, falsifiers can vary independently of verifiers then it is hard to see how this could be so.

I myself raised a related problem in connection with what I called 'new linguistic constructions' in the paper (Fine 2017b: 635) that Williamson cites:

> 'But it looks as if, within the present [bilateral] framework, we should give an account, not only of when [the counterfactual] A > C [an example of a new construction] is verified by a given state, but also of when it is falsified by a given state, since the negative content of a bilateral proposition is not in general a function of its positive content. This is a source of considerable embarrassment, since it would appear to allow a kind of ambiguity in the proposition expressed by linguistic constructions which does not genuinely exist.'

And Williamson is quite correct in pointing out that this is also a problem for the treatment of atomic sentences.

I have said that this is a problem for the truthmaker semanticist, but there is a way in which it may be a problem for everyone. For let us grant that there is a reasonably intelligible notion of *exact* truthmaking or falsitymaking—of a verifier or falsifier that is relevant as a whole to the sentence it

1 I should like to thank Timothy Williamson, Jin Zeng and two referees for the journal for some very helpful comments.

DOI: 10.5334/met.305

makes true (or false). Then it looks as if we should be able to determine the falsifiers of an atomic sentence, i.e. the verifiers of its negation, on the basis of its truthmakers. Consider the sentence 'this patch is green' and let us suppose that the candidate truthmakers (in the 'state space' under consideration) concern the patch's shape and color. Then we have no hesitation in saying that the patch's being red would be an exact falsifier for the sentence whilst the patch's being red and round would not be. Whence this ability (which no doubt relates to our remarkable ability to discern the opposite of words even though the opposite is rarely a straight negation)? Thus, this is a problem for anyone who recognizes the intelligibility of exact truthmaking; and the worst that might then be said against the truthmaker semanticist, as such, is that his semantics *obliges* him to deal with a notion whose basis for application is somewhat problematic.

How might this problem be solved? In the previously cited paper, I suggested that we might make use of a primitive relation of exact exclusion between states. Thus, the state of the patch's being red would exactly exclude its being green whilst the state of its being red and round would not exactly exclude its being green. Armed with the notion of exact exclusion, we may then provide a general explanation of how the falsifiers of an atomic sentence (or of new linguistic constructions) might be determined on the basis of its verifiers. For the falsifiers will be the states that exclude its verifiers (or, to be more precise, if s_1, s_2, are the sentence's distinct verifiers, then its falsifiers will be the states of the form $t_1 \sqcup t_2 \sqcup \ldots$, where each t_i is a state that excludes s_i).

This solution is to be clearly distinguished from the unilateral version of truthmaker semantics, which also makes use of the notion of exclusion. Under the unilateral semantics, the notion of exclusion provides a *general* explanation of the falsifiers of a sentence in terms of its verifiers; the falsifiers of any sentence whatever will be those states that exclude its verifiers. However, the sentences α and $\neg\neg\alpha$ will not in general express the same proposition under the unilateral semantics (Fine 2017b: 634–5; Williamson 2024: fn. 8 §3.4), and so it will not be well suited to someone who wishes to adopt a classical perspective on negation and yet does not wish to buy into such a high degree of hyperintensional difference.[2]

But the present suggestion is that we only appeal to exclusion in accounting for the falsifiers of atomic sentences (or of new linguistic constructions). We still apply the standard bilateral clauses to logically complex sentences—conjunctions, disjunctions, negations and the like; and so, given the clause for negation, we will still be able to maintain that the sentences α and $\neg\neg\alpha$ will express the very same proposition.

Williamson calls the appeal to exclusion a 'complication' (2024: §3.4 fn. 8); and I suspect that he might think of it as little better than the original bilateral semantics in that it also introduces a new and unnecessary degree of freedom. For even though the falsifiers of an atomic sentence can, with its help, be determined on the basis of its verifiers, we have some choice as to what we take the relation to be. Thus, the semantics introduces a new "parameter", the exclusion relation, *E*, which comes along with the ever attendant danger of over-fitting.

However, I think that Williamson somewhat overstates the contrast between truthmaker semantics and the standard intensional semantics in this regard. He correctly points out that 'a classical semantics … in effect equates falsity with non-truth for sentences' (2024: 132). Thus, there is in this case no need for an additional parameter. But this is the classical *extensional* semantics. Williamson is an advocate of the classical *intensional* semantics. Within this semantics, the semantic value of a sentence—the proposition it expresses—can be identified with a class of possible worlds (those, intuitively, in which the sentence is true) and the proposition expressed by a negative sentence $\neg\alpha$ will then be the complement, within the domain *W* of all possible worlds, of the proposition expressed by the sentence α that is negated. But then, in contrast to the extensional version of the classical semantics, we will still have need of an additional parameter, viz., the domain *W* of possible worlds. The choice of *W*, moreover, can be extremely impactful. For example, depending upon what we take it to be, a given proposition might or might not be a necessary truth or one proposition might or might not be the negation of another. It looks as if our

2 Although under certain special circumstances, one can provide a unilateral semantics that uses an appropriate notion of exclusion and also conforms to double negation (Randriamahazaka 2025: 237–9).

DOI: 10.5334/met.305

poor student might even, by the intensionalist's own lights, have some difficulty in understanding what it would take for the given sentence to be false!

Indeed, the parameter associated with the choice of the domain *W* can naturally be regarded as a special case of the parameter associated with the choice of an exclusion relation *E*. For we might take the states to be worlds and, given a particular choice of *W*, we might take *E* to be the relation which relates two members of *W* when they are distinct from one another.[3] It is thereby built into the identity of the exclusion relation that it only relates possible worlds; and, under this understanding of exclusion, the negation of a proposition will be true in a world just in case the proposition is not true in that world—just as the intensional semantics requires.

There is a way, though, in which thinking of the domain *W* as a parameter is highly misleading. For it is not as if, in applying intensional semantics, we are free to give it this or that value. Rather, it is the nature of a constant (like the speed of light). For the purposes of applying the theory, we take it to be fixed (standardly, by the class of metaphysically possible worlds); and so there is no actual danger of over-fitting, of manipulating its value to fit the data.

But by the same token, it is also highly misleading to think of the exclusion relation *E* as a parameter. For let us grant that there is a reasonably intelligible notion of exact falsification (just as the intensionalist should grant that there is a reasonably intelligible notion of a metaphysically possible world). Then we might define a state *s* to *exclude* a state *t* if it is an exact falsifier for the proposition that *t* obtains—or for the sentence 't obtains', where t is a name for *t* (Fine 2017b: 664–5). This then serves to fix the notion of exclusion; and so we are within our rights to treat it as a constant rather than as a parameter of our theory. Of course, this still leaves us with the question of determining what the actual value of *E* is—of determining when one state does indeed exclude another. But likewise, in the case of the possible worlds semantics, we are left with the question of determining the actual value of *W*—of determining when a given world is indeed metaphysically possible.

There remains a worry, however. For given what makes a sentence true, whether the sentence be atomic or not, then should I not normally be able to determine what makes it false? Williamson takes the degree of freedom arising from 'the quasi-independence of verifiers and falsifiers' to be 'most striking for atomic sentences' (2024: 132); and this suggests that he would also find it troublesome when logically complex sentences are in play and, indeed, what he goes on to say about 'the slack between α and $\neg\alpha$' suggests as much. But within a version of the bilateral semantics under which it is only the falsifiers of atomic sentences that are determined through exclusion, it will still not be possible to determine the falsifiers of a logically complex sentence on the basis of its verifiers; two sentences, one of which is logically complex, may have the same verifiers and yet differ in their falsifiers. So do we not face the very same problem as before, though at the level of logically complex sentences rather than simply at the level of atomic sentences?

I am not so sure. Go back to the extensional semantics. Suppose we allow truth-value gaps, as has been argued to arise in the case of presupposition. Then we do not expect to be able to determine whether a sentence is false on the basis of whether or not it is true; for, given that it is not true, it could either be false or be neither true nor false. Similarly in regard to truth-value gluts, should they exist; for, given that the sentence is true, it could either fail to be false or be both true and false. In these cases, we need to state how the connectives (and other logical operations) behave with respect to both truth and falsity—the semantics is, in this respect, bilateral; and there is, moreover, some choice as to what we might take this behavior to be. In the case of gaps, for example, we could go with "weak Kleene" and maximize the assignment of gaps, or go with "strong Kleene" and minimize the assignment of gaps.

A similar point holds in regard to the intensional semantics. But in this case it is not merely the actual existence of truth-value gaps (or gluts) that would force the semantics to go bilateral, but also the possible existence of truth-value gaps (or gluts) —as might be the case, for example, with

3 Strictly speaking, we should also include the null state and the full state and allow the full state to exclude any state.

DOI: 10.5334/met.305

possible failures of presupposition or singular reference; and so the need to go bilateral would arise for a much broader class of sentences.

The general point is this: given whether or when a sentence is true, there may be a separate question as to whether or when it is false; and in such cases, we would naturally expect the semantics for the logical connectives (and the like) to be bilateral, taking account both of whether or when the component sentences are true and also of whether or when they are false.

Now, in the case of truthmaker semantics, we are not simply interested in *whether* or *when* a sentence is true or false but also in *how* it is true or false. But given how a sentence is true, there will in general be a separate question of how it is false; and so, again, we should naturally expect the semantics for the connectives to be bilateral, taking account both of how the components sentences are true and of how they are false.

Of course, the need to adopt a bilateral approach will arise for a much broader class of sentences under the truthmaker semantics than under the intensional semantics. For even when a sentence is necessarily bivalent, we may still need to distinguish between how it is made true and how it is made false. But similarly, the need for a bilateral approach will arise for a much broader class of sentences under the intensional semantics than under the extensional semantics; and it is hard to see how the legitimacy of adopting a bilateral approach should somehow depend upon the extent to which its application is required. One would naturally expect the required application to be broader under a more refined semantics than under a less refined semantics.

The apparatus of exclusion is relevant to another problem that Williamson raises. He (2024: 133) writes:

> 'From the truthmaker semantics, Fine extracts a general account of a proposition, a candidate for being expressed by a declarative sentence in a context, as an ordered pair of sets of states the first conceived as the set of verifiers, the second as the set of falsifiers'

It is true that I sometimes talk of a proposition as if it could be identified with an ordered pair of sets of states (with each component set itself perhaps being subject to certain conditions). But that is only because for the purposes at hand (such as the definition of analytic implication) a more refined conception of what a proposition might be is not required. I certainly do not think in this way of a proposition *in propria persona*, one that is 'a candidate for being expressed by a declarative sentence in a context'. Indeed, I say as much (Fine 2017b: 629):

> 'Once we allow a proposition to have both positive and negative content, there should be some constraints on how the two contents are related; one should not allow an arbitrary set of verifiers to be paired with an arbitrary set of falsifiers.'

I add:

> 'One source of constraint comes from the modal connections between the verifiers and the falsifiers (we shall later deal with constraints that arise from considerations of relevance).'

Williamson considers how one might make 'use of a modal constraint to demarcate propositions' (2024: 134). He focuses on Exclusivity, that no verifier is compatible with a falsifier, although he might also have considered Exhaustivity, that any possible state is compatible with a verifier or with a falsifier. He finds such a solution lacking on two counts. The first is that it violates what I have taken to be the spirit of my approach, under which 'there is nothing in the general notion of content or meaning or in the most general logical devices that requires us to draw a distinction between possible and impossible states' (Fine 2017a: 566). The second is that, through being 'ad hoc', it introduces 'yet more complexity, and more degrees of freedom to his [Fine's] semantic framework, thereby intensifying the danger of over-fitting' (2024: 134).

DOI: 10.5334/met.305

But if our interest is in the proper demarcation of the class of propositions, then there is perhaps a more serious difficulty with the proposal, since it appears to let in far too many ordered pairs of sets of states as propositions. Suppose, for example, that one pairs the singleton of the state of the patch being green with the set of all possible worlds in which the patch is not green. This pairing then satisfies Exclusivity and Exhaustivity; and yet it is hard to believe that there is any such proposition. And what goes for this pairing goes, of course, for many other pairings in which the falsifiers of a putative proposition do not appear to be relevantly matched to its verifiers.

In fact, I never intended the modal constraints to be the whole solution (or even part of the solution) to the demarcation problem, though looking back at what I wrote I see that the passage cited above is not altogether clear in this regard. The solution to this problem was meant to come from the second source ('relevance'), which is discussed in the technical appendix to Fine (2017b: 664–5). I there attempt to characterize the domain of propositions under a bilateral conception of what the propositions are. Simplifying somewhat, the idea is to start with atomic propositions. These are propositions of the form $<\{s\}, \{t_1, t_2, \ldots\}>$, where s is a state and t_1, t_2, ... are all the states that exclude s. The propositional domain will then consist of all the propositions that can be obtained from the atomic propositions by forming negations and arbitrary conjunctions—although, of course, it is always on the cards that we might for theoretical reasons wish to consider a broader domain of proposition-like entities.

As far as I can see, this characterization does not introduce any further degrees of freedom; for once given exclusion, the propositional domain is fully determined.

The characterization of the domain of propositions *is* more complicated than in the intensional case. But again, it is easy to exaggerate how more complicated or different it is. Suppose I ask: what are the propositions under the intensional approach? Then a liberal answer, corresponding to the liberal answer I gave in the truthmaker case, is that a proposition is given by an arbitrary set of possible worlds. But suppose I attempt to justify this answer, to explain why each set of possible worlds should correspond to a proposition. Then an especially good answer, it seems to me, can be seen to rest upon the following two assumptions. The first is that to each possible world there should be a proposition true in exactly that world. The second is that the domain of propositions should be closed under negation and arbitrary conjunctions.

But this is exactly the same kind of answer that we gave in the truthmaker case. The only difference lies in the starting point. Instead of taking each possible world to correspond to a proposition we take each state to correspond to a proposition (under a suitable pairing with falsifiers). Indeed, we might even see the answer given under the intensional approach as a special case of the answer we gave under the truthmaker approach, with the states now being worlds and exclusion being distinctness.

The question of whether the characterization is implicitly modal is trickier. It is implicitly modal if exclusion is. But is exclusion implicitly modal? A reason for thinking it is not is that truthmaking is not. For given any states s and t, s will exclude t just in case s is a truth-maker for the sentence 't does not obtain' (for t a name of t).[4] Not that it would much bother me if the characterization were implicitly modal. For the point would still remain that the truthmaker framework enables one to distinguish, in an especially illuminating way, between those elements of logic and semantics that are essentially modal in character and those that are not.

It is therefore not at all clear that the charges Williamson lobbies against the modal characterization of the propositional domain—of its being ad hoc, needlessly complex, and subject to undue degrees of freedom—will hold much sway against the alternative, relevance-based, account.

I turn to Russell (as discussed in §3.5 of Williamson 2024). 'The picture', Williamson tells us, behind Russell's view 'is that a declarative sentence expresses a proposition with an overall constituent structure ... like that of the syntactic constituent structure of the sentence' (2024: 135). But Williamson adds an important qualification; 'in some cases, the syntactic form of the sentence

4 But see Champollion and Bernard (2024) for a contrary view.

misrepresents its underlying *logical form*'—a point familiar, of course, from Russell's theory of definite descriptions.

Russell's structural conception of propositions can be seen to be embodied in the principle that two propositions are the same only if they are constructed in the same way from the same constituents. This leads Williamson to lay down the following principle on behalf of the structuralist (2024: 135):

> FINEGRAIN If $O(C_1, ...,C_m) = O^*(C^*_1, ...,C^*_n)$ then $O = O^*$, $m = n$, and $C_i = C^*_i$ for $1 \leq i \leq n$', where O and O* are m- and n-place operators that 'figure in the build-up of propositions' and where $C_1, ...,C_m$ and $C^*_1, ...,C^*_n$ are, respectively, suitable inputs for these operators.[5]

Williamson argues that this principle, along with some natural background assumptions, gives rise to an inconsistency. Before discussing how this inconsistency is thought to arise, let me make two general comments about the nature of the difficulty. First, the connection with the overall theme of overfitting in his book is somewhat tenuous. It is merely that the ways he himself considers of 'patching it [the theory] up in a semi-fine-grained way promises much greater complexity for no commensurate gain—some kind of half-hearted over-fitting.' Second, he (2024: 137) remarks, 'the underlying problem [as given by the difficulty he is about to raise] is both simpler and deeper than that [the difficulty raised by the Russell-Myhill paradox]'. But I believe that these two issues are largely, though not completely, orthogonal to one another. We can perhaps all agree that there are natural solutions to the Russell-Myhill paradox (by means, for example, of potentialism or ramification) which provide no solution to the difficulty Williamson raises. But, as I hope will become clear, there is a natural solution to his difficulty which, under the assumption that there is a basic operation for forming arbitrary conjunctions, still provides no solution to the paradox.

I turn now to the difficulty (p. 138). Williamson considers 'the common definition of the necessity operator □ as the dual of the possibility operator ◊'. We may write this as:

$$\text{DF}\Box \quad \Box p =_{df} \neg\Diamond\neg p$$

Given DF□, it would appear to follow that:

$$\text{DEF}\Box \quad \Box p = \neg\Diamond\neg p$$ [6]

But applying FINEGRAIN to DEF□ gives the absurd result:

$$\text{DEF}\Box! \quad \Box = \neg \text{ and } p = \Diamond\neg p$$

To avoid this difficulty, Williamson considers the possibility that the definition, DF□, is not intended to yield the identities of the propositions, as in DEF□, but the identity of the operator, as with:

$$\text{DEF}\Box\lambda \quad \Box = \lambda p\,(\neg\Diamond\neg p)$$

But, whether the definition DF□ can legitimately be understood along the lines of DEF□λ, surely it can be understood in such a way that DEF□ is a consequence of the definition. Indeed, take the view (championed by Russell) that the left-hand side of DF□ can simply be regarded as an abbreviation for the right-hand side.[7] Then surely DEF□ will hold, since then $\Box p = \neg\Diamond\neg p$ will simply be an abbreviation for the self-identity, $\neg\Diamond\neg p = \neg\Diamond\neg p$.

---

5 Williamson expresses himself somewhat loosely. It is not, of course, the symbolic operators O and O* that will 'figure in the build-up of the propositions' but the operations signified by those operators. I will follow Williamson in being somewhat loose over how I express myself, especially in regard to the use-mention distinction and the distinction between operators and operations.

6 Williamson does not appear to distinguish between the definition itself, DF□, and the truth, DEF□, to which it gives rise. But I assume that he would have no difficulty in allowing such a distinction.

7 Strictly speaking, we have an abbreviation *scheme*, one in which 'p' could be any sentential expression. But we may put such subtleties to one side.

So how, in this case, can the disastrous application of FINEGRAIN be blocked? The proper response, it seems to me, is that DF□ is, as in Russell's own theory of descriptions, a definition in use; it shows how □ is to be understood in certain contexts, viz. those of the form □..., where '...' is a sentence. But this then means that the required application of FINEGRAIN:

$$\text{If } \Box p = \neg\Diamond\neg p \text{ then } \Box = \neg \text{ and } p = \Diamond\neg p$$

is not even meaningful, since the definition provides no way of understanding □ in the context of the identity-statement □ = ¬. Of course, one might be tempted to think that □ is a propositional operator that signifies a propositional operation but, in taking 'the syntactic form of the sentence' to represent in this way its 'underlying logical form', we are being misled by the syntactic form in the manner that Williamson has already warned us against.

So the most natural way of understanding the definition—at least for Russell—is one in which the identity of the propositions can be secured with zero danger from the application of FINEGRAIN.

I said before that we are not *required* to take DF□ to be a definition of the operator □. But surely we can provide a definition of the operator:

$$\text{DF}\Box\lambda \quad \Box =_{df} \lambda p.\,(\neg\Diamond\neg p)$$

And we might, as before, even regard the left-hand side of the definition as a mere abbreviation for its right-hand side. And from this definition, we can derive:

$$\text{DEF}\Box\lambda \quad \Box = \lambda p.\,(\neg\Diamond\neg p)$$

What then comes of Williamson's argument? Williamson regarded this alternative understanding of the definition as a way out of the previous difficulty, one which need not commit us to the identity of the propositions in DEF□. But I am inclined to regard it, not as a way *out* of the previous difficulty, for which it is not required, but as a way *into* a new difficulty. For FINEGRAIN now has meaningful application. Suppose, in addition, that we accept the following form of Abstraction:

$$\lambda p.\,(\neg\Diamond\neg p)p = \neg\Diamond\neg p$$

From DEF□λ, we obtain:

$$\Box p = \lambda p.\,(\neg\Diamond\neg p)p$$

By the logic of identity:

$$\Box p = \neg\Diamond\neg p$$

And so we are all set to obtain the untoward instance of FINEGRAIN and the unwelcome conclusion that □ = ¬.

Since we may regard □ simply as an abbreviation of λp.(¬◊¬p), we see that the present difficulty has nothing to do with definition as such.[8] The difficulty simply derives from the following instance of Abstraction:

$$\lambda p.\,(\neg\Diamond\neg p)p = \neg\Diamond\neg p$$

Since then, from FINEGRAIN, it would appear to follow that λp.(¬◊¬p) = ¬ and p = ¬◊¬p.

What are we to say? I agree with Williamson (2024: 139) that we should not jettison Abstraction. Even if we can understand the application of a λ-term in such a way that the principle fails,[9] surely we can also understand it in such a way that it holds. The problem must therefore lie with FINEGRAIN; and the issue, it seems to me, turns on what, in the application of FINEGRAIN, we should take the proposition-building operators to be. It looks as if we must deny that λp.(¬◊¬p) is

8 Williamson (2024) writes 'the problem is not confined to definitions' (139). But one might wonder whether the problem really has anything to do with definitions, except in so far as their implementation is already tied to the acceptance of Abstraction.

9 Fine (2012: 70–71) discusses some of the ways in which the principle may fail.

a proposition-building operator. But how can that be? If we can build a proposition by successively applying ¬, ◊ and ¬ in that order to p, then why can we not build the proposition in a single step by applying λp.(¬◊¬p) to p?

In discussing this issue, it will be helpful to get a little clearer on what one might mean by a 'proposition-building operator'. This term might be understood in a broad sense, under which it is taken for granted that proposition-building operators are closed under composition. Thus, if ◊ and ¬ are proposition-building operators, then so is ¬◊¬. This seems to be Williamson's intended use. He writes 'given proposition-building operators O_1 and O_2, there is a proposition-building operator $O_{1,2}$ whose result when applied to a proposition p is the result of applying O_2 to the result of applying O_1 to p' (2024: 139). And for 'proposition-building operator' so understood, Williamson is quite right in thinking that FINEGRAIN cannot be maintained.

But there is another, narrow, understanding of 'proposition-building operator' under which the class of proposition-building operators will *not* be taken to be closed under composition. In this narrower sense, the proposition-building operators—or what I shall now call the *basic* proposition-building operators—should actually be *used* in the construction of the propositions. Thus even though ◊ and ¬ may be basic proposition-building operators, ¬◊¬ will not be, since it is ◊ and ¬, not ¬◊¬, that are actually used in constructing the proposition ¬◊¬p from p.

Indeed, proposition-building operators in this narrow sense would appear to be to *anti*-closed under composition, if I may put it that way. If O_1, O_2, ..., O_n are basic proposition-building operators then their composition $O_{1,2,...,n}$ will not be. Indeed, one might think, more generally, that if $\lambda x_1 x_2 ... x_n.(\Phi(x_1, x_2, ..., x_n))$ is any suitably complex lambda term, then there is no basic proposition-building operator O for which $O = \lambda x_1 x_2 ... x_n.(\Phi(x_1, x_2, ..., x_n))$, since any proposition built through the application of $\lambda x_1 x_2 ... x_n.(\Phi(x_1, x_2, ..., x_n))$ will, in contrast to O, involve the application of more than one basic operator.[10]

What I therefore propose is that FINEGRAIN be restricted to basic proposition-building operators.[11] FINEGRAIN will then, indeed, help provide a structural criterion of identity for all propositions. Moreover, this formulation avoids all the difficulties that Williamson raises against his own more liberal formulation of FINEGRAIN, since all of those difficulties depend upon appealing to propositional operators which are evidently not basic.

An analogy with the formation of singleton sets may help to drive the point home. One might adopt a structuralist conception of singleton sets according to which singletons are the same only if they are constructed in the same way from the same constituents; and such a conception might then be taken to be embodied in the following principle:

> FINEGRAIN-FOR-SINGLETONS If O(C) = O*(C*) then O = O* and C = C*, where O and O* are operators that figure in the build-up of singletons and where C and C* are respective inputs for these operators.

One might then object to this structuralist concept of singletons on the grounds that both O = {-} and O* = {{-}} are singleton building operators, yet O(O(Socrates)) = O*(Socrates) in spite of the fact that O ≠ O* and C ≠ C*. Clearly, this is not a good argument against the structuralist conception of singletons; it only goes to show that the proposed version of FINEGRAIN does not give proper expression to the conception. And it seems to me, that Williamson's argument against the structuralist conception of propositions is no better than this argument against the structuralist conception of singletons.

In the case of propositions, we might think of ourselves as recognizing a new way in which a logical symbolism may be misleading as to logical form. Recall the case in which □p is an *abbreviation*

10 I say *suitably* complex, since we may wish λp.(¬p), for example, to be a basic proposition-building operator, given that λp.(¬p)p = ¬p.

11 A similar critique might be made of the formulations of structuralism to be found, for example, in Dorr (2016) and Goodman (2017). I certainly would not wish to claim that my proposal is the only way to respond to the difficulties raised by FINEGRAIN. See Bacon (2023) for a rather different response.

DOI: 10.5334/met.305

for ¬◊¬p. Then □p is misleading as to its logical form, since □ is not being used to represent a propositional operator. But consider now the case in which □ is an abbreviation for λp.(¬◊¬p). Then we no longer have the same reason for thinking that □p is misleading as to its logical form, since □ will indeed represent a propositional operator. However, it will still be misleading as to its logical form, though in a different way. For □p, in contrast to ¬◊¬p, will not represent the "canonical" manner in which the resulting proposition q = □p is built from p. And this point will still stand if we replace □ by a variable O. For Op, under the assignment of □ to O, will likewise not represent the canonical manner in which the proposition, under that assignment, has been built.

Williamson himself seems to recognize the possibility of restricting the class of relevant proposition-building operators. He writes that, for the structuralist, 'FINEGRAIN holds only within some unspecified, privileged subclass of proposition-building operators; it is not the advertised general principle for individuating propositions' (2024: 143). One might wonder who is guilty of false advertising here, though I must confess to often finding it difficult to ascertain exactly what form of structuralism structuralists wish to advocate. But, be that as it may, it seems to me that Williamson's discussion makes the task of formulating a more satisfactory principle seem much more daunting and difficult than it actually is. For one thing, he seems to think the default is one under which every suitable λ-term will represent a proposition-building operator. As he puts it, 'The λ operator seemed to provide for just such a general account of proposition-building operators' (2024: 142). But the exact opposite is the case; the use of λ-terms is a barrier rather than a portal to the proposition-building operators in the sense that is relevant to the formulation of FINEGRAIN.

In the second place, his attempt to save the structuralist from his objections to FINEGRAIN puts them in a needlessly unflattering light. In response to those objections, he initially has the structuralist proposing that Abstraction be taken to be an equivalence rather than an identity. He then shows that even when Abstraction is taken to be a *material* equivalence, the existence of a "Russellian" proposition-building operator R for which:

$$ROp = \neg OOp$$

will give rise to a contradiction (upon setting O = R). The 'friends of a fine-grained view of propositions', he then suggests 'may reply that ... there is no such operator as R'. 'But that response', he continues, 'is inadequate, since it was the fine-grained view itself that made it seem natural for there to be such an operator as R in the first place. ... We need a more positive, general account ...; *post hoc* denials in particular cases are not enough.' (2024: 142).

This all makes the structuralist look pretty bad. They must deny an extremely plausible form of Abstraction; they must deny the existence of operators that clearly seem to exist; and, to top it all, they are forced to do so on ad hoc grounds.

But none of these moves made on behalf of the structuralist are necessary and none of the objections to the structuralist have any force against the present view. Abstraction in the form of an identity can be retained; and the existence of R and of other operators that appear to have perfectly good specifications need not be denied. Nor, under the present structuralist view, should R be taken to be a basic proposition-building operator, since its application to the proposition Op clearly involves the successive application of O and ¬ to Op. There is therefore nothing on the present view that would make it seem natural for R to be such an operator, nor anything ad hoc about denying its status as such an operator.

I would not wish to maintain that there are no difficulties in developing a reasonable form of structuralism under the present approach (even putting the Russell-Myhill paradox to one side). We can say in a general way what a basic proposition-building operator is: it will not be the composition of other proposition-building operators; and it will not correspond to a complex λ-term. We might even suppose that proposition-building operators in the broad sense are built up from other proposition-building operators by means of suitable *operator*-building operators (again in a broad sense). The basic proposition-building operators will then be those that are not themselves built up from operator-building operators or the like. This provides a very natural understanding of the basic proposition-building operators within the structuralist framework; they, like basic constituents in general, are not themselves built up.

DOI: 10.5334/met.305

But one would like to be able to say and do more by way of developing the structuralist position. One task is to provide a rigorous formulation. But have we not done this already through the restricted version of FINEGRAIN? No, for FINEGRAIN only tells us when two propositions of the form $O(C_1, ...,C_m)$ are the same. It tells us nothing about propositions of the form $O(C_1, C_2, ...)$ where O is an infinitary operator or about propositions not of either form, if such there be; and, of course, for the account to be complete, it should also provide a structural account of the identity of the non-propositional constituents. This is by no means trivial, since it is far from clear what kind of account one should give of how complex properties, such as the property of being red and round, are built up.[12]

Additional problems arise from attempting to achieve the desired degree of generality of formulation within the confines of a higher-order language. This appears to be Williamson's preferred means of expression (2024: fn. 10), though it is by no means integral to the structuralist project. But staying within higher-order logic, FINEGRAIN would have to be stated by means of infinitely many schemes rather than a single compendious proposition. For each specific pair of distinct numbers m and n, we should have:

$$O(C_1, \dots, C_m) \neq O*(C*_1, \dots, C*_n)$$

and, for each specific number m, we should have:

$$\text{If } O(C_1, \dots, C_m) = O*(C*_1, \dots, C*_m) \text{ then } O = O*, C_1 = C*_1, C_2 = C*_2, \dots, C_m = C*_m$$

If we wished to round out the account by saying that each proposition was of the form $O(C_1,C_2, ...)$, then this could not be done (even ignoring the case in which O was infinitary), both because the constituents $C_1,C_2, ...$ might be of arbitrary finite length and because they might be of arbitrarily high type.

All of this is in marked contrast to the formulation of an identity criterion for propositions under either an extensional or intensional conception of their identity. In the extensional case, we can simply say:

$$\text{if } p \equiv q \text{ then } p = q$$

and in the intensional case, we can say:

$$\text{if } \Box(p \equiv q) \text{ then } p = q,$$

perhaps defining $\Box A$ by $A = T$. I do not know if Williamson's unhappy formulation of FINEGRAIN somehow arose from his desire to stay within the usual confines of a higher-order language, as with the extensional and intensional criteria above.

A further problem arises from the need to use the notion of being basic in the formulation of FINEGRAIN. To this end, one could either take it to be a primitive (at each type) or to be defined. In the latter case, it would be natural to take some notion of constituency as primitive. A basic propositional operator could then be defined as one without any constituents (though how this is to be formulated within higher-order logic, given that the constituents can in principle be of arbitrary type, is another problem). But in either case, one should round out the theory by laying down appropriate axioms for the additional primitives. If, for example, the notion of being basic is taken as a primitive within the theory then we would need to say whether or not it itself is basic.

A final issue is to *specify* the basic proposition-building operators and the other building operators. It is one thing to say what it is to *be* a basic operator or *how* they are, but another to say *what* they are. Again, no easy task. Consider the operators $\Box$ and $\Diamond$ and the operator ∇ for being contingent. Are there three basic operators here but also a number of corresponding derived operators—such as $\neg\Diamond\neg$, $\neg\Box\neg$ and $\Diamond \wedge \Diamond\neg$? Or is there just one basic operator, $\Diamond$ say, and two derived operators, $\neg\Diamond\neg$, $\neg\Box\neg$ and $\Diamond \wedge \Diamond\neg$? Or are $\Box$ and $\Diamond$ both basic operators but ∇ a derived operator? Each position has

12 See Fine (2010: §10) for one suggestion as to how this might be done. The incorporation of neutral relations in the sense of Fine (2000) adds further complications.

its difficulties. My own inclination, for what it is worth, is to adopt the third, intermediate position under which $\Box$ and $\Diamond$ are both basic (or have basic forms) whilst ∇ is not. This might appear to be invidious since, from a logical point of view, any two of these operators can be defined in terms of the other. However, metaphysics is not logic; and it might be argued that, from a more metaphysical perspective, $\Box$ and $\Diamond$ have a kind of intrinsic simplicity which is lacking in ∇; and similarly for other cases of this sort. Thus, $\wedge$, $\vee$ and $\neg$ might well be regarded as basic whilst the Sheffer operators (not-both, neither-nor) are not.

What I believe the line of argument pursued by Williamson shows is that it is far harder to provide a reasonably explicit and rigorous formulation of the structuralist position than one might have thought. A simple formulation within the confines of higher-order logic is not possible; further resources and further questions arising from the presence of those resources must also be taken into account. Williamson writes, 'At first sight, the picture of structured propositions seems to present a clear, systematic vision of the structure of propositions. On closer inspection, it turns out to be all, or almost all, smoke and mirrors' (2024: 142). But I do not think such a harsh verdict is called for on the basis of the objections he raises. We do not have 'smoke and mirrors' but, at worst, a promissory note that has every reasonable chance of being redeemed.

COMPETING INTERESTS

The author has no competing interests to declare.

AUTHOR AFFILIATIONS

Kit Fine
NYU, US

REFERENCES

Bacon, A. 2023. 'A Theory of Structured Propositions'. *Philosophical Review*, 132(2): 173–238. DOI: https://doi.org/10.1215/00318108-10294409

Champollion, L and **Bernard, T.** 2024. 'Negation and Modality in Unilateral Truthmaker Semantics'. *Linguistics and Philosophy,* 47: 517–555. DOI: https://doi.org/10.1007/s10988-023-09407-z

Dorr, C. 2016. 'To Be F Is To Be'. *Philosophical Perspectives,* 30(1): 39–134. DOI: https://doi.org/10.1111/phpe.12079

Fine, K. 2000. 'Neutral Relations'. *Philosophical Review,* 109(1): 1–33. DOI: https://doi.org/10.1215/00318108-109-1-1

Fine, K. 2010. 'Towards a Theory of Part'. *Journal of Philosophy,* 107(11): 559–589. DOI: https://doi.org/10.5840/jphil20101071139

Fine, K. 2012. 'Guide to Ground'. In: Correia, F and Schneider, B (eds.) *Metaphysical Grounding.* Cambridge: Cambridge University Press. pp. 37–80. DOI: https://doi.org/10.1017/CBO9781139149136

Fine, K. 2017a. 'Truthmaker Semantics'. In: Hale, B, Wright, C, and Miller, A (eds.) *A Companion to the Philosophy of Language*. Wiley-Blackwell. pp. 556–577. DOI: https://doi.org/10.1002/9781118972090

Fine, K. 2017b. 'A Theory of Truthmaker Content I: Conjunction, Disjunction and Negation'. *Journal of Philosophical Logic,* 46(6): 625–674. DOI: https://doi.org/10.1007/s10992-016-9413-y

Goodman, J. 2017. 'Reality is Not Structured'. *Analysis*, 77(1): 43–53. DOI: https://doi.org/10.1093/analys/anw002

Randriamahazaka, T. 2025. *Modelling Meaning: A Logico-Algebraic Investigation of Hyperintensionality*. Unpublished thesis (PhD), University of St Andrews.

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

TO CITE THIS ARTICLE:
Fine, K. 2026. Comments on Timothy Williamson's 'Overfitting and Heuristics in Philosophy'. *Metaphysics*, 9(2), pp. 33–44. DOI: https://doi.org/10.5334/met.305

Submitted: 05 March 2026
Accepted: 29 September 2026
Published: 04 November 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.