Living in a Material World: A Critical Notice of *Suppose and Tell: The Semantics and Heuristics of Conditionals* by Timothy Williamson

DANIEL ROTHSCHILD
University College London, United Kingdom
d.rothschild@ucl.ac.uk

1 THE MATERIAL CONDITIONAL

Barristers in England are obliged to follow the 'cab rank rule', according to which they must take any case offered to them, as long as they have time in their schedule and can agree on fees. This rule is designed to ensure that unpopular people and causes can get legal representation. In philosophy, by contrast, we know we need no structural rule to ensure that seemingly hopeless causes receive representation from the brightest minds. And so we find Timothy Williamson, who famously defended epistemicism about vagueness over a quarter century ago, turning in his new book to another unpopular cause, the view that the meaning of 'if' is given by the material conditional.

The material conditional view of 'if' is the idea that linguistic expressions of the form 'if A then C' (which I'll write '$A \rightarrow C$') are false if A is true and C is false, and true otherwise. This makes $A \rightarrow C$ equivalent to 'not A or C' ('$\neg A \vee C$'). To take an example, 'If Lena has wings, she is an angel' is false, on the material conditional view, only if Lena has wings and is not an angel; the sentence also means the same as the sentence, 'Either Lena has no wings or she's an angel'. The material conditional is, of course, the standard conditional connective in propositional and first-order logic. It is reasonable to hope that the material conditional might correspond to the natural language conditional, as this connection would strengthen the ties between natural language and logic.[1]

The proponents of the idea that the material conditional is adequate to describe conditional expressions in natural language have always been an illustrious lot, including such luminaries as Paul Grice, David Lewis, and Frank Jackson. However, even these seasoned advocates could never do the view much service. Grice's seminal account of conversational implicature in his 1967 William James lectures gave such a robust defence of truth-functional (i.e. material) views of 'and' and 'or' that these have dominated ever since. But the less-read lecture in which Grice argues for the material conditional view of 'if' is convoluted and unconvincing.[2] Jackson's [1979] views were at least clear: the natural language conditional, $A \rightarrow C$, is the material conditional, but its assertibility goes with

[1] At the least, this is the hope for those uses of conditionals called 'indicative', which include all examples I discuss until §6.

[2] His treatment of conditionals, unlike many other sections of the lectures, was unpublished until *Studies in the Ways of Words* appeared in [1989]. Even Grice himself describes his arguments there as following a 'tortuous course' (p. 83).

the conditional probability of C given A (i.e. $A \rightarrow C$ is only assertible if you have a high conditional confidence in C on the supposition that A). Nonetheless Jackson's view, too, has been roundly rejected in the literature.[3]

In the preface to this latest attempt to shoehorn the material conditional into natural language, Williamson describes the material conditional view as 'unfashionable', suggesting an element of whim in its demise. However, its unpopularity is not due to shifting trends, but rather to the simple fact that language users do not use the conditional in the way we should expect were the material conditional view correct. The convolutions required to get the semantics (i.e. the meaning) of 'if' to depart so far from ordinary usage can make even the cleverest and most distinguished advocates of the material conditional view look ridiculous. Here are some standard examples of ways in which the natural language conditional does not behave like the material conditional:

i. While it's usually fine to infer $A \rightarrow C$ from C, it's unacceptable to infer $A \rightarrow C$ from $\neg A$. For the first inference: if you know that the music box is broken, then you can infer that if the music box looks fine it's (still) broken. For the second inference, if the music box doesn't look fine, you can't thereby infer that if the music box looks fine, it's still broken. This is puzzling if $A \rightarrow C$ has the same semantics as $\neg A \vee C$, since, given the symmetry of disjunction, one should expect the two inferences to be equally good.

ii. Our confidence in a conditional of the form $A \rightarrow C$ seems to rise or fall with our conditional confidence in C given A. By contrast, our confidence in the material conditional ought to be at least as high as our confidence in $\neg A$. As an example, my confidence in the conditional 'if the car is broken, the mechanical fault light will be on', depends just on how confident I am that the mechanical fault light will be on given that the car is broken. I do not become confident in the conditional just in virtue of being confident that the car is not broken. This is not what you would expect if $A \rightarrow C$ has the same truth-conditions as $\neg A \vee C$, since being confident that the car is not broken makes you confident that either the car is not broken or the mechanical fault light is on.

iii. If we believe a material conditional is false, we ought to believe its antecedent is true and its consequent is false. But this does not accord with our practice with natural language conditionals. To think that the statement 'either the Senate passed the bill or the House did' is false, I have to think neither of them passed the bill. But thinking neither of them passed the bill is not necessary for thinking the statement 'if the Senate didn't passed the bill, the House nontheless did' is false.

[3] For decisive problems with both views see Dorothy Edgington's magisterial 'On Conditionals', which, though the size of a short book, appeared in this journal in 1995 and gave prominence to the no-truth-value (NTV) view of conditionals.

The gradual decline of the material conditional view has coincided with a rich literature in semantics exploring how natural language conditionals actually work. Important views arising from the ashes of the material conditional are Robert Stalnaker's [1968, 1975] variably strict conditional, Angelika Kratzer's [1981, 1986, 2012] hypothesis that 'if' is a device for restricting modal operators, Dorothy Edgington's No-Truth-Value [1995] approach, and recent dynamic and expressivist views such as those of Antony Gillies [2004] and Seth Yalcin [2007]. There is currently no canonical view of conditionals, but there is, on all sides, at least agreement that the material conditional view is dead. In *Suppose and Tell*, however, Williamson tries to revive the view.

2 THE MATERIAL CONDITIONAL IS DEAD, LONG LIVE THE MATERIAL CONDITIONAL!

Williamson does not give the material conditional a significant role in his discussion of conditionals until over a third of the way into the book. He starts, instead, by focusing on what he calls the 'heuristics' that guide the use of conditionals. It is only after the heuristics are presented that Williamson introduces the material conditional semantics to complete his account.

The main heuristic introduced is a procedure governed by this rule:

> **Suppositional Rule (SR)** Take an attitude unconditionally to 'If *A* then *C*' just in case you take it conditionally to *C* on the supposition that *A*. (p. 19)

Chapters 3 and 4 expound with ingenuity the implications of systematically following the rule, providing insightful explanations of various features of our use of conditionals.[4] Cases i-iii, above, for example, can easily be explained by our use of suppositional procedures in evaluating conditionals.

In putting the suppositional procedure at the heart of his account of conditionals, Williamson follows a long tradition whose origin is this observation of Frank Ramsey's [1929]:

> If two people are arguing 'If *p*, will *q*?' and are both in doubt about *p*, they are adding *p* hypothetically to their stock of knowledge and arguing on that basis about *q*'...

This suppositional procedure, often called the Ramsey test, plays a crucial role in our process of *evaluating* conditionals, but the big question is how we build that into our semantic account of conditionals. The natural answer is to try and see what kind of semantic account of conditionals would explain the use of the suppositional procedure for evaluating conditionals.[5]

[4] Some of his arguments follow in the footsteps of Adams [1965] and Edgington [1995], but he focuses less exclusively on probabilities.

[5] This is the approach taken by Stalnaker [1968] in his classic paper 'A Theory of Conditionals'.

Williamson, though, does not use the Suppositional Rule as a guide to the semantics in the usual sense. Rather, he puts the Suppositional Rule in the driving seat of his treatment of conditionals, explaining our use of conditionals directly in terms of it without any recourse to, or even agreement with, the actual semantics. In this respect, Williamson knowingly rejects standard methodology in semantics.

Williamson's position is closest in spirit to that of Edgington, who puts the probabilistic version of the Ramsey test at the heart of her account of conditionals.[6] For both Edgington and Williamson, the semantics of the conditional plays no role, direct or indirect, in how we normally assess the truth or probability of a conditional. For both of them, the reason for this reduced role for semantics is the same: suppositional procedures lead to contradictions.[7] There is a crucial difference, however: for Edgington, following Lewis [1976, 1986], the contradictions only arise when you assume conditionals express propositions whose probabilities we report using the suppositional procedure. For this reason, she denies that conditionals have truth values or express propositions, making herself the most prominent advocate of the no-truth-value (NTV) account of conditionals. Williamson's suppositional heuristic, by contrast, licenses enough bad inferences to generate contradictions without even assuming that conditionals express propositions.[8] The fact that the suppositional heuristic is such a ready source of contradiction is why Williamson calls it a *heuristic*, referencing the useful but fallible psychological mechanisms posited by psychologists Daniel Kahnemann and Amos Tversky to explain systematic failures in human reasoning.

Williamson keeps company with NTV theorists like Edgington to the extent that he thinks that there is no semantics that captures the suppositional procedures. Those familiar with Williamson's earlier work will not be surprised to learn that he thinks conditionals nonetheless have truth-conditions.[9] They just do not have the kind of semantics that explains why we use the suppositional rule. In this respect Williamson's position is closest to that of Jackson, who treats the semantics of the conditionals as incapable of explaining much of their use.[10]

The main reasons Williamson gives for endorsing the view that conditionals have truth values stems from his view that there is another, *secondary* heuristic

[6] She follows and expands upon earlier work by Ernest Adams [1965, 1975].

[7] Judging the probability of a conditional to be its conditional probability is an instance of using Williamson's suppositional procedure. There are other uses of it, however, not all of which Edgington signs up for.

[8] Chapter 3. Williamson argues (p. 42–45) that no-truth-value theorists should accept this as well. The crucial assumption he makes here is that we define conditional probabilities for conditionals (p. 42), something NTV theorists, such as Edgington, can costlessly deny, since they do not think conditionals express propositions.

[9] Williamson [1994, §7.2] gives an argument that all assertions are either true or false. If the NTV position is correct, this argument fails.

[10] His view is unlike Jackson's, however, in that it does not limit the application of the Ramsey test to assertion and it gives a different role to the semantics.

associated with conditionals. This is the testimonial heuristic, according to which users of the conditional can come to believe 'If *A* then *C*' on the basis of testimony or memory. So, for Williamson, to believe 'If *A* then *C*' one can either suppose *A* and conclude *C* or pick up the sentence via some epistemically sound process general enough to apply to any sentence, conditional or not (e.g. by testimony or memory).

Here is where things get really weird. It is not that Williamson argues that the semantics is what actually does the work of communication. For in many cases what we communicate with conditionals is closer to what the supposition heuristics would predict; for example, when I say 'It's likely that if Alma is a reservist, she's been called up', I don't communicate that the material conditional is likely (i.e. that it's likely that either she's not a reservist, or she's been called up). Rather, what the semantics does, according to Williamson, is to somehow mediate between the two heuristics. Such mediation is needed because, Williamson argues, the two heuristics are in tension. For instance, in certain cases we can take conditionals on board using the testimonial heuristics even if they are not acceptable on the suppositional heuristic.[11] The semantics, however, somehow keeps the house divided standing:

> Thus a natural question arises: what holds the mixed bag of heuristics together? We cannot expect a mere random assortment to have much cognitive value. This is the natural point for the truth-conditional semantics to enter, and explain how the mixed bag is unified as a practice of using given kinds of sentence to express, retain, and communicate given kinds of information. Without such a truth-condition, the practice falls apart. (p. 116)

> As theorists, we arguably have to ascribe the standard truth-table to 'if' in order to best understand how the mix of heuristics used for 'if' hang together as a coherent and pointful cognitive practice. (p. 121)

This reduced, and oblique, role for semantics allows Williamson to explain why our intuitions about what conditionals mean depart so wildly from the material conditional. Our evaluation of sentences with conditionals generally follows the suppositional rule rather than the semantics. This fact accounts for why we have what Williamson calls 'illusions' of erroneous truth values for conditionals.[12]

These illusions, moreover, are very stable since we lack even reflective access to the material conditional semantics—as the stability in our judgments in cases like i.–iii above indicates. This makes Williamson's borrowing of the expression

[11] This is argued for in Chapter 5, e.g. pp. 102–103. I will challenge some aspects of these arguments below in my discussion of communication.

[12] See Chapter 6.

'fast and frugal heuristics' (pp. 4, 23, 250, 265) from the psychology literature potentially misleading. For heuristics in the psychology literature (such as those posited by Kahnemann and Tversky, including the famous conjunction fallacy) can be overruled by slower, more careful reasoning. Our alleged illusions about the probability of conditionals cannot.

One might say that for Williamson conditionals are material in the sense that the British government is a monarchy. Yes, there's the Queen who meets other heads of states and signs bills and is structurally critical to certain aspects of government, but ultimately the monarch does not govern. It is only by having a comparably reduced view of the role of semantics that Williamson is able to escape what seems to be damning evidence against the material conditional view of 'if'.

3 PREDICTIVENESS AND UNDERDETERMINATION

One of the appeals of semantics in the broad tradition engaged in by philosophers and linguists since the days of Lewis, Montague, and Kaplan is that because semantic accounts are systematic (usually in virtue of their compositionality) once you have the semantics of a particular expression you can, given the semantics of other bits of the language, determine the meaning of complex sentences that include the particular expression.

The suppositional heuristic, however, like the Ramsey test that inspired it, covers only the attitudes one should hold to sentences of the simple form 'If *A*, *C*'. Conditionals, though, appear throughout our language, often embedded in complex sentences such as this one:

(1) A dairy farmer usually sells a cow if she thinks it's not producing enough milk.

The suppositional heuristic, even combined with the semantics of the embedding expressions such as 'usually', is silent on what to do with such sentences. Williamson is aware of this concern and addresses it directly:

> The Suppositional Procedure and Rule explicitly mention only *unembedded* occurrences of conditionals. What about *embedded* occurrences, where conditional sentences figure as constituents of more complex sentences? Surely a story needs to be told about cognition of them too. But consider the crucial role of the usual logical connectives, including a conditional one, in mathematical reasoning. At least to a first approximation, it is adequately codified by the introduction and elimination rules for each connective in a standard system of natural deduction. The introduction and elimination rules for a given connective explicitly mention only its unembedded occurrences, where it occurs as the main connective of the given

> sentence; that applies in particular to the natural deduction rules for the conditional. Nevertheless, when the rules for all the connectives are combined, the result is an adequate background logic for mathematical reasoning, which handles embedded occurrences of conditionals just as well as embedded ones. The Suppositional Rule serves the analogous purpose for 'if' in the wider setting of general language use, where non-logical considerations usually dominate. There is no special lacuna for embedded conditionals. Of course, they may occur under other connectives whose associated epistemology remains to be understood, but that problem is not for an account of conditionals to resolve. (p. 26)

In other words, parallel to the compositional semantics of natural language there is another, related but distinct system of the 'epistemology' of expressions, and in this system the rules and heuristics for evaluating simple sentences combine to produce rules for evaluating complex sentences.

Williamson is right to point to natural deduction as a proof of concept for such an enterprise. But note that the embedding expressions in English that felicitously accept conditionals go well beyond 'connectives' in the usual sense to include noun phrases such as 'the fact that' and 'the rumour that', attitude verbs such as 'hopes that' and 'fears that', epistemic modals such as 'might' and 'must', adverbs of quantification such as 'usually' and 'sometimes', quantifiers such as 'none', and 'almost every', adverbs like 'only', and connectives such as 'because' and 'except'. While the truth-conditional semantics of all these expressions are debated, there are plausible accounts available in the vibrant natural language semantics literature. There is comparatively little known about the procedures and heuristics we use to deal with all these constructions, let alone how to combine such procedures together in complex sentences.[13]

To recap: What Williamson is proposing is that alongside the semantics of natural language we build up a related but distinct account of the procedures we use to evaluate expressions. These procedures also need to be described in such a way that they combine together to yield procedures for evaluating complex sentences. Moreover, until we have built up this edifice, the suppositional account of 'if' that Williamson supports will not be able to tell us how 'if' behaves when embedded.

This is not a good state of affairs. Compare the case of the connective 'or'. It is reasonable to think that 'or' just expresses standard logical disjunction. Of course, the facts about our use of 'or' are not captured by the truth-table definition alone. For example, as with conditionals we can invoke suppositions to determine the truth of disjunctions: one can come to know that 'A or B' by supposing $\neg A$ and

[13] Relevant here are Fodor and Lepore's [1991, 2002] arguments about the difficulties of building up a compositional theory of meaning out of inferential roles for a fragment of language that goes beyond standard logical systems.

then concluding that B. In addition, there are Gricean implicatures associated with the use of disjunctions.[14] Surely as well, cognitive heuristics abound in our use of disjunctions. Nonetheless, at a first approximation, if you want to know the meaning of a complex sentence including a disjunction you look to the semantics to give you an answer.[15] You can supplement your answer with some pragmatics and psychology, but these will only play their role after the semantics has been taken into account.

Parallel to the worry about the unpredictiveness of the account of conditionals, there is a worry that, given the mainly ceremonial role of the semantics in governing our use of conditionals, the semantics is both explanatorily idle and, therefore, underdetermined. Williamson directly addresses this worry:

> A more general methodological concern with the postulation of basic fallible heuristics for assessing sentences, in particular the Suppositional Rule, is what they do to the empirical status of truth-conditional semantic theories for natural languages. Such theories are intended to explain the data of language use, and thereby to be confirmed or disconfirmed. The fear is that the postulated fallible heuristics will take over the role of explaining the data, leaving the truth-conditional semantics with no explanatory task to perform. That would leave the truth-conditional semantics not just empirically unconstrained but redundant, an idle wheel. (p. 115)

Williamson goes on to dismiss this worry as depending on a 'simplistic understanding of the relationship between theory and evidence in science'. Simplistic though it may be, standard methodology in semantics is to use semantic theories to generate predictions about our use of sentences, in terms of both truth-conditional judgments and inferences. Such predictions are, of course, standardly mediated by pragmatics and cognitive limitations.

Nonetheless the role Williamson introduces of 'hold[ing] the mixed bag of heuristics together' (p. 116) is both obscure and at a far remove from the actual facts about how we use conditionals. While Williamson provides a kind of mathematical argument that the material conditional is the best candidate for playing this role (pp. 103-109), it is hard to see what is at stake if we were to reject some of the assumptions of his argument and argue for a different proposition instead as the semantics for conditional. Perhaps the role of the

[14] In particular, our sense that disjunctions are exclusive is probably the result of what are called 'scalar' implicatures [Horn, 1972].There is a sophisticated literature explaining the relationship between scalar implicatures in simple sentences and those in complex sentences [see, e.g. Chierchia, 2004, Spector, 2006]. In other words, there are credible and predictive theories scaling up the story of scalar implicatures to disjunctions in embedded environments.

[15] The most serious challenge to the story actually comes from the interaction of disjunction and possibility modals—the so-called free-choice reading—but there is a rough consensus that it's not the meaning of 'or' that needs to be altered to model this phenomenon [Zimmerman, 2000, Klinedinst, 2007].

semantics could be better defined to make it more clear what is at stake, but as things stand it looks like dead weight.

Williamson defends this remove from data as a form of sophisticated scientific methodology:

> . . . the semantics is by no means immune to testing by the data of natural language use. Its support comes from those data, such as the evidence that estimates of the probability of indicative conditionals tend to correlate with estimates of the conditional probability of the consequent on the antecedent. It is just that the testing does not take a naïve falsificationist form, because native speakers are not assumed infallible. It takes a more sophisticated abductive form typical of developed sciences with a large gap between theory and observation. (p. 116)

I disagree here: the distance between usage and semantics in Williamson's account is a sign not of developed science but rather of an underdeveloped theory.

4 A PREDICTION ANALYZED

Admirably, Williamson dives deep into the messy facts about how we use conditionals. He makes a serious effort to show that his account does not suffer from the same problems that bedeviled previous defenses of the material conditional. Consider this extract from his discussion of quantified conditionals:

> The quantifier 'no' provides more drastic challenges to the material interpretation. Consider (2):
>
> (2) No student failed if he worked hard.
>
> When the conditional is read materially, the truth of (2) requires the falsity of (3) for each student as a value of the variable pronoun 'he' in both occurrences:
>
> (3) He failed if he worked hard.
>
> But (3) is false on the material reading just in case the student assigned to 'he' worked hard and did not fail, in other words, he succeeded (for simplicity, we treat 'succeed' as contradictory to 'fail'; readers who wish to do so can instead substitute 'did not fail' for 'succeeded'). Thus (2) is true on the material reading just in case (4) is true:
>
> (4) Every student worked hard and succeeded.

> Needless to say, we normally do not assess sentences like (2) as even roughly equivalent to sentences like (4). Instead, we may use the Suppositional Procedure. On the supposition that an unspecified student worked hard, we ask whether he failed; in effect, we assess the open sentence '*x* failed' on the foreground supposition '*x* worked hard' and the background supposition '*x* was a student'. If the verdict is negative, our attitude to (2) is correspondingly positive. Consequently, if one seeks a conditional-free sentence similar in truth-conditions to (2), the obvious candidate is not (4) but (5):
>
> (5) No student who worked hard failed.
>
> (p. 148–149, my example numbering)

The problem with this style of reasoning is that there are too many *different* ways of using the suppositional procedure to derive truth-conditions for (2). For instance, we *could* reason as follows: *For (2) to be true it has to be the case that the suppositional procedure when applied to every student lets us arrive at the conclusion that they did not fail on the assumption they worked hard. The simplest way to guarantee the success of this procedure for an arbitrary student is to ensure that he worked hard and did not fail. Thus, (4) is the most natural conditional-free sentence similar in truth-conditions to (2).* Reasoning this way gives a completely different meaning for (2) than the one Williamson argues for.[16]

The details of the procedure we use to connect the meaning of 'no' with the suppositional heuristic for the conditional determine the reading of (2). What Williamson presents as a good prediction of his theory, as opposed to a bad prediction of the material conditional view, is mostly a demonstration of the wide range of data that can be accommodated by exploiting the underspecification of heuristics. This problem is compounded by the fact that Williamson also takes instances in which conditionals seem to act like material conditionals to vindicate his overall account.[17] Given the freedom Williamson's account allows in explaining linguistic judgments, it is no surprise that the survey of empirical problems for the material conditional view in Chapters 7 and 8 is triumphalist.[18]

[16] Interestingly, just after the passage I quoted, Williamson (p. 150) considers something like the procedure I suggest above to argue that (2) might, in fact, be equivalent to (4). The difference in the procedures depends on how you treat cases in which a student did not take the test: if you see the case as potential failure of the supposition procedure you get the (4) reading, if you see it as irrelevant you get the (5).

[17] For example he writes on page p. 129, 'The material interpretation is well able to explain the phenomenona of conditional imperatives'. See also p. 152.

[18] This is not to say that rival theories of conditionals all make sharp predictions about what conditionals mean in embedded contexts. For example, Kratzer's restrictor view [1981, 1986, 2012] allows for a wide variety of readings depending on what modal operator the 'if'-clause is restricting. But at least here (and in many other cases) the parameters that can vary are clear and well-understood.

5 CONDITIONALS AND COMMUNICATION

The considerations above suggest that accepting Williamson's account would come at the considerable cost of sacrificing much of the predictive power of formal semantics and requiring a radical rethink of the purpose of semantics itself. Perhaps, though, this is just what is needed.

In order to assess whether we should adopt Williamson's account, we need to understand what he thinks is wrong with other accounts. Dotted all over the book are considerations and direct arguments against rival views of the conditional, such as Kratzer's restrictor account, Edgington's NTV account, and dynamic or expressivist accounts such as Antony Gillies's and Seth Yalcin's. There is no room in this review to respond to even a small proportion of these remarks and full-blown arguments, almost of all of which raise interesting points and deserve significant discussion.

I will concentrate on what I take be one of the most dialectically critical arguments in the book. This is the series of arguments about communication with conditionals, concentrated in Chapter 5. The main target of this argument is what Williamson calls 'contextualist' accounts of conditionals, which are accounts that take indicative conditionals to express context-dependent propositions that depend on the epistemic situation of the context of use. Stalnaker [1975] and Kratzer [1986, 2012] are the main proponents of this view. In its crudest version—an untenable one—the contextualist takes the conditional $A \rightarrow C$ to mean that the speaker has ruled out all epistemic possibilities in which A is true and C is false.[19] At various points Williamson indicates that his discussions of communication apply also to expressivist and other non-standard semantic views (pp. 94, 101).[20]

The goal of Chapter 5 is to show that contextualist semantics for conditionals cannot capture the role of conditionals in communication. The rejection of these 'contextualist' views plays a critical role in the book, since Williamson's argument for the material conditional view depends on conditionals being context insensitive.[21]

Before Chapter 5, some passages suggest a direct if non-decisive argument against the context-sensitivity of conditionals.[22] For example, on p. 82: '. . . contextualism about bare indicative conditionals has a significant cost, because it threatens to undermine our standard practice of freely passing them about in the same verbal form between radically different epistemic contexts, by memory and testimony'.

[19] Kratzer [1986] suggests in her discussion of Gibbard's Sly Pete case this is *one* reading of the indicative conditional.

[20] This is not surprising since most proponents of non-standard semantic views such as Yalcin and Gillies assume a close relationship between the epistemic modals 'might' and 'must' and conditionals. In general, Williamson's arguments are directed against epistemic readings of conditionals, whether they are contextualist or expressivist (e.g., p. 101).

[21] This assumption comes in on page 103, where he writes, 'To do justice to the secondary heuristics, we assume that "if" introduces no context-sensitivity of its own'. This assumption then goes on to play a key role in his derivation of the material conditional view of 'if' .

[22] Williamson [2005] suggests a similar point about contextualist theories of knowledge.

By 'freely passing', Williamson means roughly that if someone one trusts asserts a sentence of the form X, one can warrantedly assert the sentence oneself without changing the form X. In this sense, indexicals do not allow free passing, since if I assert 'I am tired' you do not learn that the sentence 'I am tired', is true of you. But it is hard to believe that the lack of free passing could be taken as a *cost* of context sensitivity. For example, on Williamson's own account, subjunctive conditionals are context-sensitive (as we will see), but they are clearly able to pass freely between speakers.[23]

In Chapter 5, Williamson gives direct arguments against the contextualist account of conditionals. We'll start with two of his central cases (pp. 93–98, cases III and IV), which are variants of Gibbard's [1981] famous Sly Pete case. Two reliable engineers are checking a nuclear reactor. Each checks a different detector light. According to the plan of the reactor, the detector lights are red if the core is overheating, green if it is not; alternatively, if the system is not working, each can be either color. The first engineer sees a green light and reports to the controller:

(6) If the detector is working the core is not overheating or the reactor was not built to plan.

The second engineer sees a red light and reports to the controller:

(7) If the detector is working the core is overheating or the reactor was not built to plan.

Williamson argues that the controller, who trusts the engineers though he does not know the underlying setup described above, should be able to conclude the following conditional:[24]

(8) If the detector is working, the reactor was not built to plan.

On a simplistic contextualist account, we can understand (6) and (7) to be expressing that the person asserting them finds the material conditional epistemically necessary. Williamson points out (on behalf of the contextualist semantics) that if the controller goes on to take (6) and (7) as epistemically necessary for herself, then she should go on to take (8) as being epistemically necessary for herself as well. This is reasonable as if you trust someone, you might reasonably take what they find epistemically necessary for themselves to be epistemically necessary for yourself.[25]

[23] Indeed the lack of serious discussion of how communication with subjunctive conditionals works when communication with indicatives is supposed to be so problematic strikes me as a serious lacuna in the book.

[24] It is not necessary here that the controller knows the setup of the lights.

[25] This view of communication with conditionals is also in the spirit of expressivist accounts of epistemic modals and conditionals like Yalcin's [2007, §5]. For on this style of account, assertions of indicative conditionals are used to urge the audience to accept a kind of epistemic necessity without actually putting forward any proposition.

Williamson, however, goes on to consider a variation of the case in which the controller has high but imperfect confidence in the engineers who tell her (6) and (7). Plausibly, Williamson suggests, we should still expect the controller to infer (8) with only slightly weakened confidence. It would seem that a probabilistic account should work here—one in which, say, 99% credence in each statement leads to 98% credence in the conclusion. However, Williamson argues that trying to get a probabilistic account of this inference is not possible on the contextualist view.

> At first sight, such a probabilistic story is easy to tell. The controller has very high credence just short of 1 in each of (6) and (7), almost equally high credence in their conjunction, and at least as high credence as that in (8). But assume that her credence in each conditional goes with her conditional credence in the consequent on the antecedent, in accordance with the Suppositional Rule. Assume also that she puts very little credence in the antecedent. The trouble is this. Her credence in the conditional all depends on the tiny bit of her probability space where the antecedent holds, which may overlap with the tiny bit she assigns to the possibility that the engineers are not trustworthy. Thus she may end up with very little credence in the conditionals (6), (7), and (8), even though she is almost certain that the engineers spoke truly. (p. 96, my example numbering)

A full discussion of this complicated case is not possible, but let me suggest one line of response. Trust is not a blanket attitude we have towards individuals: it is dependent on subject-matter. I can trust someone on the topic of supervaluationism without trusting them on the topic of conditionals. In the passage above, Williamson supposes the supervisor's trust is degraded in exactly those worlds in which the antecedent ('the detector is working') holds. If the controller doesn't trust the engineers in the circumstances in which the detectors are working, then, in some obvious sense, she doesn't trust the engineers on the subject-matter of the conditionals of (6) and (7), which are about what happens if the detector is working.[26] Put it another way: I wouldn't say I trust a mechanic who I think is generally right except in the event that the mechanical fault indicator on my car works. My point is that it is one thing to have a blanket 99% confidence in what the engineer says independently of whether the detectors are working: that's what we would normally mean by trusting the engineers. It's quite another to trust the engineers except in the event that the detectors are working. In the latter case, even when it's very unlikely the detectors are working, there is, in fact, lack of trust on the subject of what follows if the detectors are working. It is not even clear in such cases that the engineers' conditionals ought to be successful

[26] I am relying the idea that conditionals are *about* what happens when the antecedent holds. See Yablo [2014] for some ideas on how to model subject-matter and aboutness.

at communicating content to the controller given the structure of her trust in the engineers. So these sorts of cases cannot support Williamson's sweeping claim that "[Contextualist accounts] have systematic difficulty in explaining the probabilistic acceptance of conditional testimony across epistemically different contexts under conditions of imperfect trust" (p. 98).

Williamson would not find the response above adequate, as he thinks accounts of indicative conditionals need to explain how we can communicate conditionals in the peculiar cases in which accepting $A \rightarrow C$ goes along with assigning C a low probability conditional given A. These cases come in many varieties and have a long history of discussion in the conditionals literature since at least John L. Pollock [1981], though the variant Williamson focuses on originates with Vann McGee [2000]. Williamson gives the most compact version of McGee's type of case that I know:

> An expert psychiatrist is giving a lecture. Pointing somewhere on a slide of the brain, she says: 'This area of the brain is enlarged in all congenital liars. For example, if I am a congenital liar, this area of my brain is enlarged.' We are very confident, though not perfectly confident, of the generalization and so of its instances, including (9), on her testimony:
>
> (9) If she is a congenital liar, that area of her brain is enlarged.
>
> But our credence for that area of her brain being enlarged, conditional on her being a congenital liar, is not high, for we do not find the conjunction 'She is a congenital liar and that area of her brain is enlarged' more probable than 'She is a congenital liar and that area of her brain is not enlarged.' We thereby violate the Suppositional Rule for credences, overriding it again with the secondary testimonial heuristic. (p. 101, my numbering)

This is a difficult genre of case: we seem to be able to assert a conditional of the form $A \rightarrow C$ but we also seem to assign a low probability to C conditional on A. This should not be possible on the probabilistic version of the Ramsey test. McGee found this kind of case so puzzling that he was unable to account for it.[27]

The case plays two dialectically critical roles for Williamson: i) like the previous reactor case, it is meant to show that contextualist accounts (as well as a broad range of other accounts) do not give an adequate account of communication with conditionals, and ii) it is meant to show a tension between the

[27] His short paper memorably ends as follows:

> What I would like to do now is to provide an adapted version of the standard theory [i.e. Stalnaker's account of conditionals] that corrects the deficiency the example illustrates. Unfortunately, I don't have one to offer.

suppositional and testimonial heuristics, since (9) is unacceptable on the former, but still accepted by the latter.

If the critical problem about communication and conditionals centers on cases as obscure as this one, contextualists need not lose sleep.[28] It is hard to imagine telling a group of linguists that they need to fundamentally rethink the project of compositional semantics for conditionals because they can't account for communication in bizarre circumstances where communication might just fail.[29]

Taking inspiration from Williamson's own willingness to throw out many of our judgments about conditionals, the contextualist can simply claim we are mistaken in thinking that we are entitled to confidence in (9). To offer a parallel argument to that made above, the expert psychiatrist, while an expert on many things, is not to be trusted on the topic of what follows from her being a congenital liar, and so we ought to not to accept the conditional (9) coming from her mouth. (We can of course accept the conditionals of the same form about everyone except her.) While this may seem ad hoc, the contextualist is actually well-motivated in this response as I will argue.

To get a better view of the overall situation, it is useful to step back from conditionals to communication more generally. Regardless of what we say about conditionals, we *independently* need an account of how we communicate worldly information by way of reporting conditional probabilities.[30]

Consider this version of the reactor case. The first engineer reports:

(10) The conditional probability that the core is overheating given that the detector is working is one.

The second engineer reports:

(11) The conditional probability that the core is not overheating given that the detector is working is one.

Any account of communication should explain how if you almost completely trust both engineers in these statements, then it is reasonable (and perhaps necessary) for you to come to have a high probability that the detector is not working. Actually spelling out how this kind of communication works turns out to be a thorny problem in formal epistemology: there is no consensus on how one updates one's probabilities on the basis of reliable information from experts about about their conditional probabilities.[31] However, just because

[28] This is not to say the *general* phenomenon of conditional probabilities coming apart from conditional assertability is not of interest, see Kaufmann [2004] for discussion and citations.

[29] It's also worth noting the closeness in form to the liar paradox. The conditional is similar to 'If I'm a liar, then what I say is still true'.

[30] While most people do not communicate by relaying conditional probabilities, it is certainly feasible in the Bayesian era we live in.

[31] One version of this problem is Bas van Fraassen's [1981] Judy Benjamin problem. But the problem

explaining how such communication works is difficult doesn't mean there isn't such communication. And if there is, then we can communicate with conditionals even if their semantics is contextualist or expressivist.

Looking at communication with direct expressions of conditional probabilities also sheds light on what is going on in Williamson's psychiatrist case. Suppose the psychiatrist says, 'This area of the brain is very likely to be enlarged in all congenital liars. So, for example, the conditional probability that this area of my brain is enlarged, given that I am a congenital liar, is very high'. Hearing this little speech we will be tempted to accept (12).

(12) The conditional probability that this area of her brain is enlarged on given that she is a congenital liar is very high.

We will be tempted to accept (12) because we accept a theory that, if true, has (12) as a consequent. But of course, we cannot ultimately make the probabilistic judgment in (12) (since if we learned that she is a congenital liar, we would reject the whole theory). Whatever is going on here might explain Williamson's pyschiatrist case. For direct expressions of conditional probabilities might behave as strangely as assertions of conditionals. Thus, these cases do not provide a good reason to pull apart conditionals from expressions of conditional probabilities.

If I have spent a long time discussing what might seem like rather obscure cases, it is because Williamson tasks them to do heavy lifting in support of the material conditional reading of 'if'. They cannot bear the weight.

6 COUNTERFACTUALS

So far I have discussed only Part I of *Suppose and Tell* and thus focused only on *indicative* conditionals. The second and shorter part of Williamson's book is devoted to so-called counterfactual (or subjunctive) conditionals.[32] It is interesting to see that an extremely unusual theory of 'if' can form part of a fairly straightforward semantic architecture for the counterfactual. Williamson takes the currently dominant line that 'if' has a common meaning in indicative and counterfactual conditionals. So, sentence (13) must include the material conditional.

(13) If Aisha had bought bitcoins, she would be rich now.

Williamson gives a plausible analysis of English 'would' as a contextually-restricted necessity modal operator.[33] He then supposes that the modal 'would' gets scope over the conditional, so that (13) is really of this form:

is more general and relates to issues about peer-disagreement and Bayesianism.

[32]Both 'counterfactual' and 'subjunctive' are agreed to be grossly inaccurate labels, but long usage means we are stuck with them.

[33]He is not the first to base a semantics of counterfactuals on a semantics of 'would': Asher and McCready [2007] have a related account.

(14) Would (Aisha bought bitcoins → she is rich)

On the semantic level, this proposal makes counterfactuals context-dependent strict conditionals, a recognizably orthodox semantics.[34] Here too heuristics also play a key role; a new suppositional heuristic accompanies 'would' and combines with the heuristic for 'if' to give a suppositional heuristic for counterfactual conditionals. But this role for supposition seems to integrate well with the strict conditional semantics. Williamson makes some interesting suggestions about how the heuristics might explain various familiar peculiarities of counterfactuals such as the apparent validity of the law of conditional excluded middle and Sobel sequences. As readers familiar with Williamson's recent work on subjunctive conditionals will expect, there is also an extensive discussion of conditionals with impossible antecedents or consequents and a firm rejection of the use of impossible worlds in their treatment. In this review I have focused on Part I of *Suppose and Tell* because of its radical departure from semantic orthodoxy, but there is a wealth of strong arguments and interesting observations equally worthy of discussion in Part II.

7 SEMANTIC METHODOLOGY

In his methodological manifesto 'Must Do Better' [2006] Williamson is fulsome in his praise for natural language semantics: '...[O]ver recent decades truth-conditional semantics for natural languages has developed out of philosophical logic and the philosophy of language into a flourishing branch of empirical linguistics'. He goes on to chide the participants in the Dummettian debates over realism for paying scant attention to the developments in natural language semantics and failing to ensure that their accounts 'meet the actual needs of empirical linguistics'.[35]

Suppose and Tell makes clear that Williamson is nonetheless far from happy about the current state of natural language semantics. In his conclusion Williamson charges semanticists with being too fond of complex theories that hew too closely to surface phenomena:

> In the methodology of contemporary semantics, one may observe a widespread though far from universal tendency to overcomplicate. In well-meaning attempts to fit all the data, elaborate, rickety semantic structures are attributed to the ordinary nuts and bolts of natural language. For cognitive purposes, the whole edifice looks remarkably inefficient. This postulated complexity is often treated as an indicator of theoretical prowess, rather than a warning sign of a degenerating research programme. Of course, there is plenty of evidence from

[34] See Starr [2021, §2.2] for a good discussion of the strict conditional view of counterfactuals.
[35] See also Williamson [2020, §5], where he criticizes Kit Fine on similar grounds.

> other branches of linguistics that natural languages are much more complex than they first appear. Nevertheless, semantics could benefit from greater awareness of the danger of what natural scientists call overfitting: the willingness to add extra parameters to an equation until its curve goes almost exactly through all the data points. That may sound like a laudably empirical attitude, but hard experience shows its typical result to be theoretical instability. As new data points come in, still more parameters have to be added to the equation, and their coefficients oscillate wildly. (p. 264)

Earlier in the book Williamson humorously speculates that the reasons semanticists develop complex theories might sometimes be less innocent than a desire to accommodate all the existing data:

> There sometimes seems to be an element of macho competition over who can argue for the most complicated semantics for a simple expression. Of course, it would be quite unfair to implicate most semanticists in that tendency. (p. 6, fn. 4)

If the metric of simplicity is simply the semantic entry for 'if' itself, then Williamson's view of the conditional is certainly simple. However, that simplicity comes at the cost of failing to integrate conditionals with the compositional semantics of natural language. It is only by specifying the compositional heuristic procedures for evaluating the wide range of embedding expressions that we can recover predictions out of his theory. Once these procedures are developed—if they can be—it is not clear that Williamson's overall system will still look simple. Lean semantics are well and good, but they have to work.

Another aspect of contemporary semantics Williamson criticizes is its disconnection from psychology and cognitive science:

> Many semanticists (though not all) prefer semantic explanations. That is not surprising. Every discipline tends to favour its proprietary ways of thinking. That can include neglect of obvious alternatives distinctive of other disciplines. Although semanticists are usually alert to possible rival explanations from pragmatics, the danger here is neglect or dislike of potential explanations from less linguistic disciplines, such as psychology and cognitive science. (p. 76)

Williamson own account of conditionals is a proudly cognitive one: 'The Suppositional Conjecture is a psychological hypothesis, which in the end must live or die by psychological evidence' (p. 22). The passage goes on to clarify, however, that the brand of psychology to be found here is purely speculative:

> However, this book does not contain much discussion of experimental data. Before we can sensibly test a model against such data, we

> must properly understand the model itself, and what it does or does not imply. Science needs well-developed theories to make sense of its almost intractably messy data, which often result from dozens of interacting variables. Without such theories to guide it, data-gathering risks becoming a parody of directionless Baconian inductive inquiry. Properly developing a theory is no easy task. It is enough for one book. (ibid.)

Bringing cognitive science into the theory of the conditional is a laudable goal. A good theory of how we think using conditionals might be preferable to a complex semantic account attempting to encode these cognitive tendencies as they appear in natural language. There is, moreover, an important history of philosophers developing models to be taken up in cognitive science (e.g. Fodor's [1983] modularity hypothesis). I am skeptical, however, of Williamson's attempt to give a speculative account of our reasoning with conditionals with almost no reference to the cognitive science literature, either experimental or theoretical. Besides the borrowing of the term 'heuristics' and the mention of data supporting the relationship between conditionals and conditional probability, the cognitive science of reasoning plays almost no role in the book. For example, there is no discussion of the Wason selection task, a famous psychological paradigm about reasoning with conditionals.[36] Given the avowed cognitive orientation of the book, I would have hoped some of this literature could have been discussed despite the risk of getting lost in the 'almost intractably messy data' of science.

Suppose and Tell tells a story in which good scientific methodology and attention to the cognitive aspect of language join forces to rescue the material conditional view of 'if'. But this story is a fantasy: in the end, invocations of science are no substitute for the hard work of developing a workable theory of a natural language construction. The cost of believing this fantasy is to dismiss, without good reason, much of the progress that has been made in the study of conditionals over the last half-century.

Despite its flaws, however, the book is full of original ideas, powerful arguments, and trenchant observations. It will serve as a useful resource for those looking for unorthodox ideas about conditionals and the cognitive role of semantics. It also should give solace to the lonely souls who still favor the material conditional—if they do not look too closely.[37]

REFERENCES

Ernest W. Adams. A logic of conditionals. *Inquiry*, 8:166–97, 1965.

[36] `https://en.wikipedia.org/wiki/Wason_selection_task`.

[37] I would like to thank Susan Chambers, Simon Goldstein, Lucy O'Brien and Amia Srivivasan for comments on previous drafts. I am especially grateful to Matthew Mandelkern for discussion and comments on previous drafts and to Dorothy Edgington for detailed written comments.

Ernest W. Adams. *The Logic of Conditionals: An application of probability to deductive logic*. Dordrecht, 1975.

Nicholas Asher and Eric McCready. Were, would, might and a compositional account of counterfactuals. *Journal of Semantics*, 24(2):93–129, 2007.

Gennaro Chierchia. Scalar implicatures, polarity phenomenon, and the syntax/pragmatic interface. In Andrea Belleti, editor, *Structures and Beyond*. Oxford University Press, 2004.

Dorothy Edgington. On conditionals. *Mind*, 104(414):235–329, 1995.

Jerry Fodor. *Modularity of Mind*. MIT, 1983.

Jerry Fodor and Ernest Lepore. Why meaning (probably) isn't conceptual role. *Mind and Language*, 6(4):328–343, 1991.

Jerry Fodor and Ernie Lepore. *The Compositionality Papers*. OUP, 2002.

Allan Gibbard. Two recent theories of conditionals. In W. L. Harper, R Stalnaker, and G Pearce, editors, *Ifs: Conditionals, Belief, Decision, Chance, and Time*. Reidel, 1981.

Anthony Gillies. Epistemic conditionals and conditional epstemics. *Nous*, 38(4): 585–616, 2004.

Paul Grice. *Studies in the Ways of Words*. Harvard, 1989.

Laurence Horn. *The Semantics of the Logical Operators in English*. PhD thesis, UCLA, 1972.

Frank Jackson. On assertion and indicative conditionals. *Philosophical Review*, 88(4):565–589, 1979.

Stefan Kaufmann. Conditioning against the grain. *Journal of Philosophical Logic*, 33:583–64, 2004.

Nathan Klinedinst. *Plurality and Possibility*. PhD thesis, UCLA, 2007.

Angelika Kratzer. The notional category of modality. In H.-J Eikmeyer and H. Reiser, editors, *Words, Worlds, and Contexts*, pages 38–74. Walter de Gruyter, 1981.

Angelika Kratzer. Conditionals. *Chicago Linguistics Society*, 22(2):1–15, 1986.

Angelika Kratzer. *Modals and Conditionals*. Oxford University Press, 2012.

David Lewis. Probabilities of conditional and conditional probabilities. *Philosophical Review*, 8:297–315, 1976.

David Lewis. Probabilities of conditional and conditional probabilities II. *The Philosophical Review*, 95:581–589, 1986.

Vann McGee. To tell the truth about conditionals. *Analysis*, 60:107–111, 2000.

John L. Pollock. *Indicative Conditionals and Conditional Probability*, pages 249–252. Springer, 1981.

Frank Plumpton Ramsey. General propositions and causality. In Frank Plumpton Ramsey, editor, *The Foundations of Mathematics and other Logical Essays*, pages 237–255. Kegan Paul, Trench, Trübner, 1929.

Benjamin Spector. *Aspects de la pragmatique des opérateurs logiques*. PhD thesis, Universite Paris 7, Denis Diderot, 2006.

Robert Stalnaker. A theory of conditionals. In N. Rescher, editor, *Studies in Logical Theory*, pages 98–112. Oxford University Press, 1968.
Robert Stalnaker. Indicative conditionals. *Philosophia*, 5:269–286, 1975.
William Starr. Counterfactuals. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2021 edition, 2021.
Bas van Fraassen. A problem for relative information minimizers in probability kinematics. *The British Journal for the Philosophy of Science*, 32(4):375–379, 1981.
Timothy Williamson. *Vagueness*. Routledge, 1994.
Timothy Williamson. Knowledge, context and the agent's point of view. In G. Preyer and G. Peter, editors, *Contextualism in Philosophy*, pages 91–114. Clarendon Press, 2005.
Timothy Williamson. Must do better. In P. Greenough and M. Lynch, editors, *Truth and Realism*. Oxford University Press, 2006.
Timothy Williamson. Vagueness: A Global Approach, by Kit Fine. *Mind*, 2020.
Stephen Yablo. *Aboutness*. Princeton University Press, 2014.
Seth Yalcin. Epistemic modals. *Mind*, 116:983–1026, 2007.
Thomas Ede Zimmerman. Free choice disjunction and epistemic possibility. *Natural Language Semantics*, pages 255–290, 2000.