'Might' Counterfactuals, Thickened Semantics, and Agglomeration

August 24, 2010

Matthew A. Benton
mbenton@philosophy.rutgers.edu

1 Introduction

Much attention has been given to counterfactuals in the recent philosophical literature. Most of this attention focuses on "would"-counterfactuals: subjunctive conditionals of the form

(1) If he had tagged up, he would have scored the winning run.

But comparatively little work has been devoted to the would-counterfactual's unfortunate cousin, the "might"-counterfactual, which looks like the following:

(2) If he had tagged up, he might not have scored the winning run.

When philosophers do discuss might-counterfactuals, it is usually to provide an account of the relation between would- and might-counterfactuals. This essay considers three prominent accounts of might-counterfactuals, and then confronts a puzzle involving quantum mechanics and the semantics for counterfactuals. I propose a revision to the standard semantics which eludes the problem posed by the puzzle, and finally, defend that revision's denial of the agglomeration inference and show how it reestablishes that account to a better standing relative to the other two.

2 Three Rival Versions of Might-Counterfactuals

2.1 Duality

David Lewis initially proposed that the would- and might-counterfactuals are duals. That is, where the symbols "$\Box\!\!\rightarrow$" and "$\Diamond\!\!\rightarrow$" denote the would- and might-counterfactual, respectively, the Duality Thesis (DT) states that

$$p \Diamond\!\!\rightarrow q \ =_{df} \ \sim (p \Box\!\!\rightarrow \sim q)^1$$

[1]David Lewis (1973a): 21. Thus they are sometimes called "Lewisian equivalencies."

DT has the advantage of capturing the logic of our usual conversational practice. For example, an easy way to deny a would-counterfactual is to assert its might-not counterpart, as (1) and (2) above demonstrate. One virtue of DT is that it predicts this, for the negation of a would-counterfactual looks equivalent to its might-not counterfactual. That is, given (i),

i. $p \mathrel{\Box\!\!\rightarrow} q$,

one may deny it by asserting (ii) below; and the following steps show it to be equivalent, on DT, to (i) negated.

ii. $p \mathrel{\Diamond\!\!\rightarrow} \sim q$ (denial of i)

iii. $\sim (p \mathrel{\Box\!\!\rightarrow} \sim\sim q)$ (DT)

iv. $\sim (p \mathrel{\Box\!\!\rightarrow} q)$ (double negation)

Asserting (iv) is just the other way to deny a would-counterfactual: negate it outright. And of course, the reverse can be done: to deny a might-counterfactual $p \mathrel{\Diamond\!\!\rightarrow} q$, assert its would-counterfactual cousin with the same antecedent and negated consequence: $p \mathrel{\Box\!\!\rightarrow} \sim q$.[2] So the might- and would-counterfactuals intuitively operate as duals, interdefinable as specified by DT. But DT encounters some problems, as will be seen below.

2.2 Two Epistemic Accounts

Robert Stalnaker was the first to suspect that there may be something epistemic about might-counterfactuals. He points out that DT treats "the apparently complex construction *if . . . might*, as an idiom which cannot be explained in terms of the meanings of if and might." Since "might" can also be assessed in nonconditional constructions, Stalnaker's approach begins with them, noting that while "might" on its own obviously expresses some kind of possibility, it most commonly expresses epistemic possibility.[3] He considers

(3) John might come to the party

and

(4) John might have come to the party,

in order to register both the epistemic and nonepistemic readings "might" can generate. (3) seems to be a clear case of epistemic possibility, which is why (3),

[2]See Lewis (1973b), repr. in (1986): 8.
[3]Stalnaker (1984): 143.

conjoined with its will-not denial, resulting in

(5) John might come to the party, although he won't

sounds contradictory. In (3), the epistemic possibility of John coming to the party "is compatible with the speaker's knowledge," which is why (5) sounds so bad: the first conjunct asserts a possibility which the second conjunct withdraws.[4]

On Stalnaker's view, "might" is treated as a possibility operator on the would-conditional such that $p \Diamond\!\!\rightarrow q$ expresses some kind of possibility that $p \Box\!\!\rightarrow q$, with the context determining which kind of possibility is dominant. And Stalnaker's very reasons for adopting an epistemic reading of many might-counterfactuals are related to his rejection of DT:

> The main evidence that *might* conditionals are epistemic is that it is unacceptable to conjoin a *might* conditional with the denial of the corresponding *would* conditional. This fact is also strong evidence against Lewis's account, according to which such conjunctions should be perfectly normal.[5]

Stalnaker cites the following dialogue to make this point.

> X: Would President Carter have appointed a woman to the Supreme Court last year if a vacancy had occurred? [i.e., $\mathit{vacancy} \Box\!\!\rightarrow \mathit{woman}$?]
>
> Y: No, certainly not, although he might have appointed a woman. [i.e., $\sim (\mathit{vacancy} \Box\!\!\rightarrow \mathit{woman})$ & $(\mathit{vacancy} \Diamond\!\!\rightarrow \mathit{woman})$]

Y's response sounds odd in the same way that (5) does; but DT sanctions Y's reply, since $\sim (\mathit{vacancy} \Box\!\!\rightarrow \mathit{woman})$ is equivalent to $(\mathit{vacancy} \Diamond\!\!\rightarrow \sim \mathit{woman})$; and $(\mathit{vacancy} \Diamond\!\!\rightarrow \mathit{woman})$ looks compatible with this, for Y's second conjunct appears to assert that Carter might have, but also implies that he might *not* have, appointed a woman. The lesson is supposed to be that DT fails to account for the epistemic character of some might-counterfactuals.[6,7]

---

[4] Stalnaker notes that (4) can express a non-epistemic possibility, namely that "it was within John's power to come, or that it was not inevitable that he not come." But (4) can also express the epistemic possibility that John came to the party, if uttered by someone who can't remember whether or not John was there. So (4), and "might have" constructions generally, seem genuinely ambiguous. See Keith DeRose (1998): 67–9.

[5] (1984): 144.

[6] Ibid. Stalnaker also discusses what he calls a "quasi-epistemic" possibility of certain might-counterfactuals, which he thinks can capture the non-epistemic sense Lewis adverts to in his penny example of (1973a): 80. A quasi-epistemic possibility is compatible not just with "my knowledge, but what would be compatible with it if I knew all the relevant facts," that is, "possibility relative to an idealized state of knowledge. If there is some indeterminacy in the language, there will still remain some different possibilities, even after all the facts are in, and so this kind of possibility will not collapse into truth" (1984): 145.

[7] However, it isn't at all clear that these examples undermine DT. The DT advocate has two responses available, which reinforce one another. The first response is to disambiguate Y's

Keith DeRose follows Stalnaker but goes even further, defending a view on which might-counterfactuals always express an epistemic possibility. His epistemic thesis (ET) suggests that

$$p \diamondto q \ =_{df} \langle e \rangle \ p \boxright q$$

where "$\langle e \rangle \phi$" represents an epistemic possibility that ϕ, and where epistemic possibilities are understood as possibilities of the kind that sentences of the form "It is possible that P_{ind}" typically express (the subscript "$_{ind}$" indicates that the P expressed is in the indicative mood).[8] DeRose's definition of what this possibility amounts to says that any assertion of "It is possible that P_{ind}" is governed by the "flexible hypothesis,"

> (FH) S's assertion, "It is possible that P_{ind}" is true if and only if (i) No member of the relevant community knows that P is false, and (ii) There is no relevant way by which members of the relevant community can come to know that P is false,

where both who is and who is not a member of the relevant community and what is and what is not a relevant way of coming to know are very flexible matters, variable according to the context of utterance.[9]

So ET posits a close connection between might-counterfactuals and sentences of the form "It is possible that P_{ind}" and FH explains the epistemic components of such possibility assertions. DeRose argues for the invariably epistemic reading of might-counterfactuals by stressing that when an assertion appears to have a

reply in the Court vacancy discourse. Arguably Y's reply sounds bad because its opening "No, certainly not . . . " has a scope ambiguity, and on its most natural reading those words function not to negate X's entire conditional, but to assert the stronger claim according to which the negation takes narrow scope by attaching to the consequent: (*vacancy* $\boxright \sim$ *woman*) (see Lewis (1973b), repr. (1986): 7). On this reading, Y's full reply considers the antecedent proposed by X, that a vacancy on the Supreme Court had occurred, and then claims that Carter would not have appointed a woman, although he might have appointed a woman. But, importantly, on this disambiguation DT has a direct explanation for why it's odd: it's flatly inconsistent, because (*vacancy* $\boxright \sim$ *woman*) contradicts (*vacancy* $\diamondto$ *woman*).

Second, given the (intended) wide scope reading of Y's response, what Y affirms has the form $\sim(p \boxright q)$ & $(p \diamondto q)$. Now on DT, that is equivalent to $\sim(p \boxright q)$ & $\sim(p \boxright\sim q)$, which is itself equivalent to $((p \diamondto q)$ & $(p \diamondto\sim q)$. But this just amounts to the claim that Carter might and might not have appointed a woman in the event of a vacancy, that it isn't clearly settled either way. Thus Y's response as found in Stalnaker's example is problematic on pragmatic grounds: it sounds odd because it is a convoluted (and ambiguous) way of saying something for which more succinct and smooth responses will do, such as "He might have, but he might not have." (Or if Y wants to downplay the inevitability suggested by the "would" of the question, Y can simply reply "He might not have," which serves to deny the would-counterfactual.) Bringing these two lines of response together: if the narrow scope reading is intended by Y, then DT rightly explains why it sounds bad; and if the wide scope reading is intended, it is pragmatically infelicitous given the easier and clearer ways for Y to respond to X's question (e.g. "He might have, but he might not have") which dont generate the more natural narrow scope interpretation, and which sound fine on DT. Either way, DT remains untouched.

[8]DeRose (1991): 581ff. and (1998): 67–8.

[9](1991): 593–4, and (1999): 398.

non-epistemic reading, it is due to clause (ii) of FH. Speakers who utter non-conditional possibility statements like

> (6) It's possible that the ball will veer left,

or a conditional possibility sentence like

> (7) If we had rolled the ball into the (indeterministic) zone, it might have veered left,

mean to convey the content of clause (ii); this is because contexts in which (6) or (7) are typically uttered are ones where the speaker and listeners already acknowledge that FH's clause (i) is fulfilled, and so clause (ii), and particularly the modal force of its "can,"[10] generates a non-epistemic (perhaps metaphysical) sense of possibility. But due to the obviously epistemic clause (i), DeRose maintains that such sentences are invariably epistemic.

2.3 The Duplicity of Ambiguity

Lewis later retreated from DT in favor of the "ambiguity thesis" (AT), on which might-counterfactuals are ambiguous: they can take on DT's original "not-would-not," but can also express a "would-be-possible" reading. Compare:

> (nwn) It is not the case that: if A had been, then it would not have been that B.

> (wbp) If it were that A, it would be that: B is possible.

On AT, the difference is that (nwn) means that some of the most similar worlds where A obtains are worlds where B obtains, whereas (wbp) means that all of them are worlds where it is possible that B happen. If the sense of possibility in (wbp) is that of a non-zero chance at the time of the antecedent, then we can get compatible statements that sound contradictory.[11] What prompted Lewis to consider AT was the possibility of a world in which a "quasi-miracle" occurs, where a quasi-miracle is "a pattern of lawful outcomes of many different chance processes" which achieves a perfect convergence with the way the world would have been had some event not transpired: e.g., if some pivotal event took place, such as Nixon's pressing the nuclear launch button, a quasi-miracle could effect a perfect "cover-up" such that the world reconverges to how it would have been had he not pressed it at all.[12]

[10] DeRose clarifies the modal force of the "can" in clause (ii): it isn't just that there is no relevant way by which members of the relevant community can come to know that P is false, and that this fact is due to their present epistemic limitations; rather, it's that they can't come to know that P is false because it is not knowable or not deducible. See (1999): 399 and 401.

[11] Lewis, Postscripts to (1979) in (1986): 63–4.

[12] Ibid., 60.

3 Minute and Remote Possibilities

One concern that makes the possibility of quasi-miracles so salient involves indeterminism at the level of quantum mechanics. This led Lewis to recognize the threat to DT:

> But if there would have been some minute probability of a quasi-miracle, does it not follow that there might have been one? And if there *might* have been one, then is it not *false* to say that there *would not* have been one? True, it would have been overwhelmingly probable that there not be one. But may we flatly say that this improbable thing *would not* have happened?[13]

This illustrates a broader problem to be faced when we accept quantum mechanical physics, for it seems to force upon us the truth of any number of might-counterfactuals. John Hawthorne puts the problem succinctly:

> On those interpretations of quantum mechanics according to which the wave function for a system delivers probabilities of location, it seems that in any mundane situation, there is always a small chance of some extremely bizarre course of events unfolding. Suppose I drop a plate. The wave function that describes the plate will reckon there to be a tiny chance of the particles comprising the plate flying off sideways.[14]

Adherence to such a theory should render many, if not all, might-counterfactuals true. In particular, Hawthorne points out, well have to accept

(8) If I had dropped the plate, it might have flown off sideways.

But this will tempt us to think that (9) is false:

(9) If I had dropped the plate, it would have fallen to the floor.

This temptation is undoubtedly motivated by something like DT, and given the standard semantics offered by Lewis, the truth of (9) will be secured only if it ends up being impossible, contrary to quantum mechanics, that the plate flies off sideways upon being dropped. Lewis's semantics for the would-counterfactual stipulates that

> A $\Box\!\!\rightarrow$ C is true at a world w iff there is an A & C world closer to w than any A & $\sim$ C world,[15]

[13]Ibid., 61 (emphasis mine).

[14]John Hawthorne (2003), repr. (2006): 255.

[15]Though this formulation ignores the vacuously true case in which there are no A&C worlds; from here on I state the semantics continuing to ignore the vacuous condition. I've

where *closeness* is based on a primitive two-place relation of comparative overall similarity between worlds.[16] But because there may be ties in similarity, there can be ties for which worlds are closest; so Lewis's semantics is reformulated in terms of a system of spheres of worlds that are equally close to the world of evaluation. A counterfactual is then true "if there is some antecedent-permitting sphere in which the consequent holds at every antecedent-world, and false otherwise."[17] More precisely,

> A $\Box\!\!\rightarrow$ C is true iff C obtains at every member of some class W of A-worlds such that every member of W is *closer* to the actual world than is any A-world not in W,[18]

where W is the aforementioned sphere. The truth of the would-counterfactual eliminates the possibility of $\sim$ C from the closest worlds where A obtains.

This semantics requires then that any world in which the plate *does* fly off sideways not be on the sphere of closest worlds in which the plate is dropped. This just means that such a world is *less* similar, making it not as close. But quantum physics urges that it *is* as close: the indeterminacy is such that a world where it flies off sideways is exactly like our own up to the moment the plate drops. So given the physics, and given DT and this semantics, we will have to conclude that virtually all ordinary would-counterfactuals are false;[19] indeed, Alan Hájek argues for this very conclusion in his paper "Most Counterfactuals are False."[20]

So we appear to be faced with three options. First, we can follow Hájek, keep DT and the semantics, and maintain that most would-counterfactuals are false. Second, we can reject DT and posit an epistemic thesis for might-counterfactuals. Or third, we can retain DT but modify the semantics. There are good reasons for exploring this third option rather than pursuing the first two, because each of the first two options brings serious drawbacks.

The second option of rejecting DT is unappealing insofar as the other accounts do not offer a simple explanation of why there is the intuitive logical interplay, of the kind exhibited by (i) through (iv) in section 2.1, between would- and might-counterfactuals. Going epistemic for might-counterfactuals may accommodate some of the conversational data, but much is lost when the precise

---

borrowed this wording from Alan Hájek (ms): 29. Lewis's own formulation in (1973b), repr. (1986): 10, is very close to this: A $\Box\!\!\rightarrow$ C is true at i iff some (accessible) world AC-world is closer to i than any (A&$\sim$C)-world, if there are any (accessible) A-worlds. For Lewis's formal analysis, see (1973a): 16ff.

[16](1973a): 91–95.

[17]Ibid., 16. Lewis's formal definition requires that A $\supset$ C holds at every world on the sphere.

[18]From Jonathan Bennett (2003): 165.

[19]"Ordinary" counterfactuals being those that don't include a hedge like "would *probably*" or "would *very likely*."

[20]Hájek actually accepts a weaker version of DT, according to which the would-counterfactual and might-counterfactual are not contradictories but rather contraries, i.e., A $\Box\!\!\rightarrow$ C and A $\Diamond\!\!\rightarrow\sim$ C cannot both be true, though they can both be false. See (ms): 3ff.

relation between the two counterfactuals remains unaccounted for. In particular, given our current puzzle with quantum mechanics, epistemic accounts can't easily explain the temptation to deny (9) once we've accepted (8), which DT predicts. And the first option, siding with Hájek, is even more difficult to stomach. Not only do we, *contra* Hájek, count many would-counterfactuals as true, his view suffers from the following odd result. Hájek thinks most would-counterfactuals are false, because he accepts both DT and that nearly all might-counterfactuals are true; but then given any true might-counterfactual, its contraposed[21] counterpart is true as well. That is, the truth of $p \Diamond\!\!\rightarrow q$ means that q is true at some p-world on the relevant sphere; and so *ipso facto* there is some q-world at which p is true. But this is implausible: even though it is true that

If I were 15 feet tall, I might be the tallest person on record

it may well be false that

If I were the tallest person on record, I might be 15 feet tall.

So on to option three.

4 "Thin" and "Thick" Counterfactuals

We seek a semantics which makes would-counterfactuals true even though there may be a few worlds on the antecedent-permitting sphere where the consequent does not obtain, and which likewise does not automatically make might-counterfactuals true due to the presence of one or a few consequent worlds on the sphere. Here is one route to such a revision, inspired by some remarks of DeRose. While discussing a debate in philosophy of religion, in which Robert M. Adams deployed DT against Molinism,[22] he writes:

> I think of would counterfactual conditionals as being "thinner" (easier to make true) than Adams does. What for Adams grounds the truth of a "would probably" counterfactual will often for me be enough to ground a "would" conditional.[23]

DeRose clarifies this as follows:

> Suppose that before the time when A would occur, it's not certain that if A occurs then C will occur—there's some objective chance

[21] I follow Lewis (1973a): 35, who calls this "contraposition" for might-counterfactuals, though this label might seem odd because the negations are left off. Lewis says this inference "has no plausibility at all."

[22] See Robert Merrihew Adams (1977), esp. section I.

[23] (1999): 409, note 4.

> that C won't occur if A does. That seems to be enough, for Adams, to show that later, after the time when A would have occurred, "If A had occurred, C would have occurred" is not true. I'm easier. I think A $\Box\!\!\rightarrow$ C can be true even though there was some objective chance that A might occur without C occurring.[24]

The outlined strategy permits a small objective chance of the consequent not occurring at some of the relevant antecedent worlds. A natural way of understanding this "thinned" semantics for the would-counterfactual is that

> $p \Box\!\!\rightarrow q$ is true iff q obtains at *most* members of some class W of p-worlds such that every member of W is closer to the actual world than is any p-world not in W,

or, more informally,

> $p \Box\!\!\rightarrow q$ is true iff *most* of the closest p worlds are q worlds.[25]

This will also allow for modifying the might-counterfactual semantics so as to "thicken" its truth-conditions, which would go correspondingly:

> $p \Diamond\!\!\rightarrow q$ is true iff q obtains at *enough* members of some class W of p-worlds such that every member of W is closer to the actual world than is any p-world not in W

or, more informally,

> $p \Diamond\!\!\rightarrow q$ is true iff *enough* of the closest p worlds are q worlds.[26]

On this modified semantics for might-counterfactuals, it takes more than a minute possibility, like a quantum mechanical one, to make ordinary might-counterfactuals true. That is, it takes more to make $p \Diamond\!\!\rightarrow q$ true than simply that there be some slight objective chance that q occurs when p occurs. Making might-counterfactuals thicker in this way avoids the falsifying threat from minute, quantum mechanic possibilities, for such possibilities don't occur in the vast majority of worlds in which the antecedent obtains.

This modified semantics retains DT when the threshold for 'most' and the threshold for 'enough' mirror one another. When an utterance of "$p \Box\!\!\rightarrow q$"

[24] Personal correspondence. DeRose thinks of the ordinary would-counterfactual more like a simple arrow "$\rightarrow$", which to his ears doesnt carry the "suggestion of a corresponding necessity." The would-counterfactual of DT, "$\Box\!\!\rightarrow$", which does carry this certainty, he thinks of as better communicating "would necessarily" or "would certainly". See (1999): 390–1.

[25] These informal analogues are, of course, not purely Lewisian since they make the Limit Assumption. Yet they are less cumbersome than the Lewisian framework, so I will use these versions from here on. Proper definition of these thinned and thickened clauses would be done in terms of measures over an infinite set of worlds; I ignore such technicalities here, but see Daniel Berntson (2010 ms) for such an account.

[26] Bennett (2003): 249ff. endorses something like this account, but does not address agglomeration.

meets the contextually determined truth-conditional threshold for 'most,' its might-counterfactual denial "$p \Diamond\!\!\rightarrow \sim q$" will *ipso facto* fail to meet the threshold for 'enough';[27] so $\Diamond\!\!\rightarrow$ and $\Box\!\!\rightarrow$ remain interdefinable as duals. Here's an example using percentages: if the threshold at a context required for 'most' to make a particular assertion of a counterfactual "$p \Box\!\!\rightarrow q$" true is that at least 95% of the worlds on the p-permitting sphere are q-worlds, then a true assertion of "$p \Diamond\!\!\rightarrow \sim q$" in that context requires that 5% or more of the worlds on the p-permitting sphere are q-worlds. Now suppose that in this context "$p \Box\!\!\rightarrow q$" meets this threshold, since, in this case, 98% of the worlds on the p-permitting sphere are q-worlds; then an assertion of "$p \Box\!\!\rightarrow q$" is true and any assertion of "$p \Diamond\!\!\rightarrow \sim q$" is false, since only 2% of the worlds are $\sim q$-worlds, and this falls short of the 5% needed to make the might-counterfactual true. Since it is false, its negation "$\sim (p \Diamond\!\!\rightarrow \sim q)$" is true, which is just the dual of "$p \Box\!\!\rightarrow q$".[28]

The thickened and thinned semantics proposed here is incomplete, for no specification of the thresholds for 'most' and 'enough' has been offered: how much is enough, in each case? Presumably, the thresholds will be context-dependent, varying according to the content of the antecedent and consequent, and perhaps also according to the salience of certain possibilities.[29] I think we can postulate that we often begin a conversational context (most ordinary conversations, anyway) with certain "default" thresholds, based on our experience with and usage of such counterfactuals—though I shall take no stand here on what that default threshold typically is: I merely use the defaults 95/5% for the sake of example. Whether such defaults can actually be ascertained is presumably a question of empirical psychology; and because my purpose here is only to offer a preliminary sketch of the view, I will proceed on the assumption that language users operate with such defaults, and leave it at that for now.

There are two main objections to this revised DT semantics. The first is that it, like its original DT predecessor, denies conditional excluded middle. I shall briefly argue in the next section that this is not a great cost. The second and more significant problem is a failure of agglomeration; this will be addressed in the subsequent section.

[27] This will be true in all cases, save one special exception: when the percentage of worlds for each of "$\Box\!\!\rightarrow$" and "$\Diamond\!\!\rightarrow$" are exactly at their thresholds. I discuss this special case below in section 6.

[28] I find it plausible to think of these sphere-world percentages as correlating with the probabilities of those consequent worlds occurring (given the physical laws and the past). For those who prefer a probabilistic semantics, the revised proposal can be adapted to fit this framework: the "similarity" of the spheres is just exchanged for (or perhaps determined by) the probability of the consequent on the antecedent and the laws and the past, such that $p \Box\!\!\rightarrow q$ is true iff $\mathrm{P}(p \mid q \mathbin{\&} \mathrm{L} \mathbin{\&} \mathrm{Past})$ is sufficiently high, meeting the threshold.

[29] For example, it seems plausible that when physicists are discussing quantum mechanics as physicists, these possibilities are much more salient (even if they aren't thereby more objectively likely to occur), thus moving the threshold for "$\Diamond\!\!\rightarrow$" down to something much smaller, perhaps 0.005 or so.

5 Against Conditional Excluded Middle

The thickened and thinned semantics for Duality denies the principle of conditional excluded middle (CEM), which has it that $[(p \;\Box\!\!\rightarrow q) \vee (p \Box\!\!\rightarrow \sim q)]$ always expresses a truth. Lewis's semantics denies CEM because it does not assume there is a unique closest antecedent world, and so it cannot countenance the idea put forth by CEM: for if there is a tie for closest between two p-worlds, and q holds at one of them but $\sim q$ holds at the other, then CEM fails. Lewis went further, using might-counterfactuals to level an argument against CEM; but recent commentators have urged us to reconsider the merits of the principle.[30] Like Lewis, our revised Duality semantics denies CEM. Lewis's own objections to CEM are forceful, and I will not rehearse them here. I will merely offer two more considerations favoring a denial of CEM, in response to two recent attempts to revive its plausibility.

The first reason to deny CEM is that it is only plausible on the assumption that $p \;\Box\!\!\rightarrow \sim q$ is equivalent to $\sim (p \Box\!\!\rightarrow q)$, a principle which Robbie Williams calls

Equiv: (A $\Box\!\!\rightarrow C) \Leftrightarrow \;\; \sim (A \Box\!\!\rightarrow$ C).

Equiv principle has some appeal, as Lewis himself noted.[31] But accepting it requires serious revision to the semantics; and thus, if our choice is between two equivalency principles, Equiv and a Duality equivalence between might- and would-counterfactuals, we should choose the latter.

Williams considers three counterfactual statements:

(a) If I were to jump on the ice, I would fall in. [$p \;\Box\!\!\rightarrow q$]

(b) Its not the case that: if I were to jump on the ice, I would fall in. [$\sim (p \Box\!\!\rightarrow q)$]

(c) If I were to jump on the ice, I wouldnt fall in. [$p \;\Box\!\!\rightarrow \sim q$]

He notes that one can negate (a) by using the "clumsy" (b), but thinks it "far more natural" to express disagreement by way of (a)'s "opposite conditional" (c).[32] And Equiv sanctions the idea that one can use either to deny (a).

But it is a mistake to equate conditionals like (b) with those like (c), since the outright negated conditional (b) is different from the "opposite conditional" (c). This is because on the Lewisian semantics, (b) can be true while (c) is false: to get the falsity of a counterfactual A $\Box\!\!\rightarrow$ C, all that is required is for there to be some $\sim$ C world amongst the class W of A-worlds, where every member of W is closer to the actual world than any A-world not in W. So the truth of (b) merely requires there to be a few, or even *one*, world in that set W where I jump on the ice, to be one where I don't fall in. By contrast, the truth of (c)

[30] Lewis (1973a): 80–82. See J. Robert G. Williams (forthcoming) and Sarah Moss (ms) for recent demurrals.

[31] Lewis, (1973b), repr. in (1986): 7.

[32] Williams (forthcoming): 2.

requires that *every* member of W, where I jump on the ice, to be worlds where I don't fall in.

There is, of course, a sense in which (c) can be used to "express the negation" of (a). But we shouldn't assume that expressing its negation *requires* opposing (a)'s strong claim about *all* the closest A-worlds by affirming a different but equally strong claim about all those worlds; another way to express its negation is to falsify it by making a weaker claim about a few of those worlds. Of course, we can grant that (c) implies (b), because on the standard semantics, $p \mathrel{\Box\!\!\rightarrow} \sim q$ implies $\sim (p \mathrel{\Box\!\!\rightarrow} q)$[33]: if it's true that if I were to jump on the ice, I wouldn't fall in, then surely it's false that if I were to jump on the ice, I would fall in. But the converse implication from (b) to (c), from $\sim (p \mathrel{\Box\!\!\rightarrow} q)$ to $p \mathrel{\Box\!\!\rightarrow} \sim q$, doesn't hold, unless one is already assuming CEM or Equiv.

Why would one accept Equiv? The main support cited by Williams is the behavior of conditionals under extensional quantifiers. The following are thought to be equivalent:

A. No x: (x goofs off $\Box\!\!\rightarrow$ x passes)

B. Every x: (x goofs off $\Box\!\!\rightarrow$ $\sim$ x passes)

"No student would have passed if they had goofed off" seems equivalent to "Every student would have failed to pass if they had goofed off." But now consider C:

C. Every x: $\sim$ (x goofs off $\Box\!\!\rightarrow$ x passes)

The unregimented (C) is best rendered as: "For every student, it's false that: if they had goofed off, they would have passed." Does (C) imply (B)? It's tempting to say that it does; but consider what we'd say if brilliant Brianna is in the class. Brianna aces all her classes when she studies diligently; however, she hates history, so she's inclined to goof off in that class. Nevertheless, Brianna is so smart that there's a decent chance that she'd still pass history even with goofing off. So, restricting our quantifier to just the students in Brianna's history class, suppose that every other student in the class is such that if they engage in any goofing off they will surely fail; Brianna is the only one who can get away with goofing off yet still have a chance of passing. (C) is true: in that history class, for every student it's false that: if they had goofed off, they would have passed. The counterfactual is true for Brianna because there are some goofing-off worlds where she passes, and many where she doesn't, and thus it's false in her case that if she had goofed off, she would have passed. But (C) doesn't imply the stronger (B), according to which every student would have failed if they had goofed off: brilliant Brianna might have passed even if she had goofed off. This suggests that an argument for Equiv from conditionals under extensional quantifiers won't work because it's only by assuming that equivalency at the outset that any such argument goes through. And it won't help to appeal to Williams's initial thought that it sounds natural to deny a counterfactual $p \mathrel{\Box\!\!\rightarrow}$

[33]Except in the vacuous case where p is impossible. I ignore such exceptions here.

q by asserting its opposite conditional $p \;\Box\!\!\rightarrow\!\sim q$, for it is even more natural, and colloquial, to deny it by affirming its might-not counterpart $p \;\Diamond\!\!\rightarrow\!\sim q$; as we noted above in section 2, such facts led Lewis to endorse Duality in the first place.[34]

Another complaint made against deniers of CEM is that it appears to settle "from the armchair" the outcome of certain "chancy" propositions.[35] This complaint is made vivid when we consider what the Duality theorist must say when the antecedent involves flipping a fair coin: the uttered instance of CEM,

> (Chancy CEM) "If I had flipped the dime, it would have come up heads, or if I had flipped it, it would have come up tails"

sounds pretty plausible. Lewis, and our Lewisian semantics sketched above, must deny it, because both disjuncts come out false; and it sounds especially bad because the Duality theorist accepts "If I had flipped it, it would have either come up heads or come up tails," but cannot infer from this to (Chancy CEM).[36] More fundamentally, the complaint asks why we should tolerate a semantics which tells us, in effect, that a truly indeterministic outcome is false.

The obvious response here is that a Duality semantics *doesn't* tell us that an indeterministic outcome is false, nor does it settle anything about outcomes from the armchair. What it tells us is that we must consider false certain chancy counterfactuals whose antecedent event would precipitate an indeterministic process but whose consequent tells us which outcome would result: if the process leading from the antecedent event to the consequent is truly indeterministic, a counterfactual which purports to tell us how things would have gone is false. If it's indeterminate whether C or ∼C given an antecedent A, then it's not indeterminate but *false* to say that if it had been that A, then C would have been, and likewise false to say that ∼C would have been. For if it's indeterminate, that is because it could have gone either way, and if it could have gone either way, then it's obviously false to maintain that it *would* have gone a given way even though it could have gone either way. In short, from the fact that a process is indeterministic it does not follow that a counterfactual describing that process is indeterminate in truth-value.

6 The Problem of Agglomeration

The proposed thickened semantics results in a failure of the Agglomeration inference for would-counterfactuals,

[34] See especially Lewis (1973b) in (1986): 7–8.

[35] See Moss (ms).

[36] It isn't immediately clear why the plausibility of $p \;\Box\!\!\rightarrow (q \vee \sim q)$ should make us accept $[(p \;\Box\!\!\rightarrow q) \vee (p \Box\!\!\rightarrow \sim q)]$. Accepting the latter on the basis of the former relies on a principle which distributes $\Box\!\!\rightarrow$ over such disjunctions. But similar principles which distribute necessity operators like $\Box$, or tense operators like WILL, over disjunctions, have been contested: see Aristotle, *De Int.* IX, 19a28–34 for the former, and Prior (1957) for the latter.

$p \mathrel{\Box\!\!\rightarrow} q$

$p \mathrel{\Box\!\!\rightarrow} r$

Therefore: $p \mathrel{\Box\!\!\rightarrow} (q \,\&\, r)$.[37]

The plausibility of this inference suggests that one ought to able to infer from the would-counterfactual premises to a would-counterfactual, the consequent of which contains the conjunction, however long, of the consequents of the premises; call this the "consequent-conjunction." Consider the following example:

If I had dropped the plate, it would've hit the floor

If I had dropped the plate, it would've shattered on impact

If I had dropped the plate, the pieces would have to be cleaned up.

From these three conditionals, shouldn't it follow, by the most obvious inference, that

> If I had dropped the plate, it would've hit the floor, shattered on impact, and the pieces would have to be cleaned up ?

The standard semantics affirms agglomeration, because at all of the closest worlds, without exception, where the plate drops, the plate also hits the floor, shatters on impact, and the pieces have to be cleaned up. But on the thickened semantics proposed, there will be a small percentage of the closest plate-dropping worlds where the consequent-conjunction doesnt hold.

What must the defender of our revised semantics do? Our thickened theorist must, it seems, deny agglomeration. But there are two ways to soften the blow. On the first way, a limited agglomeration-like inference can be preserved, and the proponent of the thickened semantics can explain why we think agglomeration holds in most cases in which it is put to use. On the second way, agglomeration can be considered a reasonable, but not valid, inference.

First way: begin by assuming that our default threshold for 'most' in affirming a would-counterfactual is 95%: we countenance would-counterfactuals that are "thin" in virtue of their antecedent-permitting sphere being comprised of at least 95% worlds at which the consequent holds. Now suppose that these two would-counterfactuals are true:

$p \mathrel{\Box\!\!\rightarrow} q$ (0.98)

$p \mathrel{\Box\!\!\rightarrow} r$ (0.98)

Therefore: $p \mathrel{\Box\!\!\rightarrow} (q \,\&\, r)$.

[37] Hawthorne (2006): 256, makes this point against any such weakening of the standard semantics.

These two premises guarantee a true conclusion, since the percentage for $p \Box\!\!\rightarrow (q \,\&\, r)$ is 0.96 or better.[38] But given more premises with the same percentages, the conclusion goes false, for the conclusion falls below the threshold:

$p \Box\!\!\rightarrow q$ (0.98)

$p \Box\!\!\rightarrow r$ (0.98)

$p \Box\!\!\rightarrow s$ (0.98)

$p \Box\!\!\rightarrow t$ (0.98)

Therefore: $p \Box\!\!\rightarrow (q \,\&\, r \,\&\, s \,\&\, t)$ (0.92)

So we will only be comfortable affirming the conclusion in cases where the premises are *safely* true, viz., where the percentages are far enough above the threshold to yield a true conclusion. In such cases, the conclusion is true because all the premises are true *and* each is safely above the threshold.[39] Hence the thickened theorist can sanction certain agglomeration-like inferences even though she denies agglomeration as invalid.

How significant a drawback this is depends on the importance we assign to such an inference, especially when we load up on counterfactual premises. But plausibly, we hesitate to agglomerate even on the standard semantics once we take on too many premises. This is because the standard semantics dictates that the worlds on the p-permitting sphere must all be worlds in which the consequent-conjunction holds, and when such premises are true, with each additional premise we arrive at a more thoroughly specified world. Suppose each of these premises are true.

O $\Box\!\!\rightarrow$ A If Otto had gone to the party, then Anna would have gone.

O $\Box\!\!\rightarrow$ B If Otto had gone to the party, then Benny would have gone.

O $\Box\!\!\rightarrow$ C If Otto had gone to the party, then Carter would have passed out.

O $\Box\!\!\rightarrow$ D If Otto had gone to the party, then Dean would have danced.

O $\Box\!\!\rightarrow$ E If Otto had gone to the party, then Eve would have left early.

O $\Box\!\!\rightarrow$ F If Otto had gone to the party, then Frank would have stayed late.

[38] "Or better" because it could turn out that the $(p\& \; q)$-worlds and the $(p \,\&\, r)$-worlds either converge or overlap: if they converge, they will be the same set of worlds, and the percentage for $p \Box\!\!\rightarrow (q \,\&\, r)$ will end up being 0.98 as well; if they merely overlap, the percentage for $p \Box\!\!\rightarrow (q \,\&\, r)$ will be between 0.96 and 0.98.

[39] Hence we will affirm a conclusion with the consequent-conjunction only when we sense that the premises are safely true—and these will likely be simple cases with few premises, or possibly, cases with multiple premises so long as each has a percentage far over the threshold. Safety in this sense is not to be confused with safety constraints on knowledge, as in Ernest Sosa (1999).

Are we actually comfortable inferring to O $\Box\!\!\rightarrow$ (A & B & C & D & E & F), that is, to

> If Otto had gone to the party, then: Anna and Benny would have gone, Carter would have passed out, Dean would have danced, Eve would have left early, and Frank would have stayed late ?

When asked to agglomerate so many consequents, we might, even on the standard semantics, second-guess whether the premises really are all true. Would all of those consequents occur, without fail, had Otto gone to the party? The difficulty is that each additional premise specifies, with greater and greater precision, the possibilities that must obtain at the antecedent worlds: we are surprised to learn that this level of convergence obtains between the set of O & A worlds and the set of O & B worlds (& C worlds, & D worlds . . .), a result that may not have set in when evaluating each would-counterfactual on its own. And we ought to be thus surprised, because "as the set over which an existential modal quantifies gets *smaller*, it gets *harder* for such modals to be true."[40]

But if the foregoing is right, then isn't our real discomfort with the plausibility of the initial long list of would-counterfactuals? This is just another way of putting the concern: agglomerating on the standard semantics conveys the necessity of the consequent(s), which in turn forces more precise convergence in cases of multiple premises. The point is that *if* the standard semantics is correct, then that gives us reason to pause when working with several would-counterfactuals to agglomerate. Agglomerating just two consequents can be easier to stomach; but agglomerating more than this means weve countenanced three or more premises, which commits us to an increasingly specific world. So we may be disinclined to agglomerate in cases beyond the simple ones, regardless of which semantics is right; but if so, then the standard semantics and the thickened semantics dont yield very different inferential procedures, at least with respect to agglomeration.[41]

Second way to soften the blow: follow Stalnaker's (1975) lead in distinguishing between a "reasonable inference" and a valid inference, and claim that agglomeration is only the former. A reasonable inference is a pragmatic relation between speech acts rather than a semantic relation between the propositions that are their contents:

> an inference from a sequence of assertions or suppositions (the premises) to an assertion or hypothetical assertion (the conclusion) is *reason-*

[40]See Kai von Fintel and Anthony S. Gillies (2008): 85–6, though their point is made concerning epistemic modals and the "time lag" involved when points of assessment become increasingly later than the context of utterance.

[41]The standard semanticist could try to explain the disinclination to agglomerate when confronting a long list of would-counterfactual premises by claiming that when we look at them in succession, our credence in each one may go down. In so doing, the standard semanticist affirms that the truth of the premises results in a true conclusion, and thus affirms the validity of agglomeration. Our thickened theorist, on the other hand, argues for the legitimacy of this disinclination based on the conclusions being judged false; and this explains *why* our credences go down as we are led through multiple premises toward a conclusion containing the consequent-conjunction.

> *able* just in case, in every context in which the premises could appropriately be asserted or supposed, it is impossible for anyone to accept the premises without committing himself to the conclusion. (1999: 65)

This would explain why agglomeration looks so good whenever we're inclined to use it, and why we'd shy away from it at times: the inference doesn't depend solely on truth-preservation from premises to conclusion, but on how well positioned we are for judging their truth. Thus, on this line, sometimes agglomeration is reasonable *and* valid, because in those cases it preserves truth; but in other cases, even though it doesn't preserve truth, it's still reasonable. Yet we often can't tell the difference between these, because in the latter, merely reasonable cases, the conclusions will be very close to making the threshold.

Of course, agglomeration will have to be denied for *might*-counterfactuals, since it is obviously invalid. Even on the standard semantics, the consequent-conjunction won't follow. Consider an everyday case:

> If it had rained, the game might've been cancelled.
>
> If it had rained, you might've been miserable.
>
> Therefore: If it had rained, it might have been that: the game is cancelled and you're miserable.

This conclusion wouldn't follow if your biggest wish is for the game to be cancelled and its being cancelled is what pulls you out of your sodden misery. Even more decisive is the following, where the coin is fair:

> If the coin had been flipped, it might've landed heads.
>
> If the coin had been flipped, it might've landed tails.
>
> Therefore: If the coin had been flipped, it might've been that: it lands heads and lands tails.

So might-counterfactuals don't admit of agglomeration.

There is one lingering worry. Above we specified default thresholds of 95% for 'most' and 5% for 'enough'; on this assumption, the semantics we typically operate with look like:

> $p \;\Box\!\!\rightarrow q$ is true iff at least 95% of the closest p-worlds are q-worlds;
>
> $p \;\Diamond\!\!\rightarrow q$ is true iff at least 5% of the closest p-worlds are q-worlds.

But what happens when *exactly* 95% of the closest p-worlds are q-worlds, and the other 5% are $\sim q$-worlds? In such a case, both

(i) $p \;\Box\!\!\rightarrow q$ (0.95)

and

(ii) $p \mathbin{\Diamond\!\!\rightarrow} \sim q$ (0.05)

are true. But this presents a problem, since a contradiction arises: (ii) is equivalent, by DT, to $\sim(p \mathbin{\Box\!\!\rightarrow} \sim\sim q)$, which is just the denial of (i). Won't this involve us in asserting contradictions?

I don't think so. In rare cases where there is an exact fit between the antecedent-permitting sphere's world-percentages and our thresholds, we will probably withhold counterfactual assertions due to vagueness considerations. When an exact fit obtains, it's just too close to call. If the threshold is at 95%, a p-permitting sphere comprised of exactly 95% q-worlds is so close to having 94% q-worlds that speakers will likely be uncomfortable asserting the would-counterfactual (or its might-counterfactual counterpart), since truths that are very nearly false tend to make us quiet.[42]

7 Back in Contention with ET

Modifying the semantics in the way envisioned helps DT avoid a chief problem posed by its epistemic rival, ET. Recall that DT, as formulated initially by Lewis, maintains that utterances of the form $(p \mathbin{\Box\!\!\rightarrow} q)$ & $(p \mathbin{\Diamond\!\!\rightarrow} \sim q)$, such as (*) are semantically inconsistent, because the two conjuncts cannot be true together:

> (*) If p, then it would've been that q, and if p, then it might not have been that q.

DeRose argues that ET permits the consistency of such conjunctions while explaining why utterances of them sound so bad: ET explains why such conjunctions "seem inconsistent while they're in fact consistent."[43] Because it has not been shown how DT can handle this, ET is thought to be the superior theory.

But the thickened and thinned semantics proposed here plausibly saves DT from this difficulty. On our suggested theory, sentences of type (*) *can* be consistent, for there are some contexts where the conjuncts can be true together; indeed, we just considered one at the end of §6, the case of exact fit between the antecedent-permitting sphere's world-percentages and our thresholds. In such a case both conjuncts will be true together, though above I maintained that

[42] The same effect could be achieved by the following alternative semantics, one of which uses a "greater than" and the other a "greater than or equal to" clause. For example:

> $p \mathbin{\Box\!\!\rightarrow} q$ is true iff 95% or greater p-worlds are q-worlds;
> $p \mathbin{\Diamond\!\!\rightarrow} q$ is true iff greater than 5% of p-worlds are q-worlds.

Here, in the envisioned case of exact fit, exactly 95% of the p-worlds are q-worlds, $p \mathbin{\Box\!\!\rightarrow} q$ is true and $p \mathbin{\Diamond\!\!\rightarrow} \sim q$ is false (since *exactly* 5% of the p-worlds are $\sim q$-worlds). Of course, it will be somewhat arbitrary which counterfactual receives which clause. But we can let there be no fact of the matter about which one has the equal-to clause and which one doesn't; then the Lewisian equivalencies can be supertrue even though in a given borderline case theres no fact of the matter as to which of the pair is the true one and which is the false one.

[43] DeRose (1999): 391. I already diffused above, in fn. 7, the problem presented by Stalnaker concerning $\sim(p \mathbin{\Box\!\!\rightarrow} q)$ & $(p \mathbin{\Diamond\!\!\rightarrow} q)$ conjunctions.

such a context would render most speakers reluctant to assert the conjoined counterfactuals. Two other such contexts come to mind: one will be realized if there is a strange threshold shift mid-sentence.[44] Another such context is one in which the semantic thresholds for $\Box\!\!\rightarrow$ and $\Diamond\!\!\rightarrow$ don't mirror one another. Above we noted that duality is retained only when these thresholds are complementary, as when the threshold for $\Box\!\!\rightarrow$ is at least 95% and the threshold for $\Diamond\!\!\rightarrow$ is at least 5% of the worlds on the sphere. But if the threshold complementarity breaks down, both conjuncts could be true together.

Consider an example of the first of these three contexts, involving a mid-sentence threshold shift, using (1) and (2) with which we began.

> (1) If he had tagged up, he would have scored the winning run.
>
> (2) If he had tagged up, he might not have scored the winning run.

(1) and (2) seem inconsistent together, as predicted by DT; and typically they are inconsistent. But on our revised semantics, they can both be true together, and hence consistent: if conjoined into

> (1+2) If he had tagged up, he would have scored the winning run;
> but had he tagged up, he might not have scored the winning run,

they would be consistent in the case of a mid-sentence threshold shift, since plausibly, the sphere of nearby worlds changes across the conjunction in order to make both (1) and (2) true in (1+2). Such an utterance could be plausible when the order is reversed: (2+1) can sound okay to a listener but typically it will require certain conversational cues which indicate the mid-sentence shift: "If he had tagged up, he *might* have tripped, and so *might not* have scored the winning run; but c'mon!—if he had tagged up, he would have scored." Here the tonal emphasis on "might" and the "c'mon!" serve to signal the threshold shift it takes to conjoin (1) and (2) in an acceptable way: it shifts up (e.g. from 5% to 6%) for the "$\Diamond\!\!\rightarrow$" claim, to let in to the antecedent-permitting sphere enough worlds where our sure-footed baserunner trips, but then dips back down to the default threshold position after the "c'mon!" pulls the context back to normal considerations, bouncing such worlds off the sphere.

Notice, however, that these three envisioned contexts are extraordinary, and infrequent enough to account for their oddity: the vast majority of the times, our conversational contexts don't feature mid-sentence threshold shifts, non-mirroring semantic thresholds, or exact fits between world-percentages and semantic thresholds. Their uncommonness can create the sense that something is wrong with such utterances, and this accounts for their apparent inconsistency. Our thickened proposal counts most such utterances as false due to DT, but it also can explain why even the rare consistent utterances of them seem inconsistent to our ears. So DT, on the thickened and thinned semantics, seems to fare as well as ET with regard to conjunctions like (*), and proponents of ET

[44] On the felicity of some mid-sentence semantic shifts, see Jason Stanley (2005): chap. 3. But see DeRose's arguments against such mid-sentence shifts, via the methodology of "all-in-one-breath conjunctions," in (1998): 70–1.

cannot obviously claim the advantage over this version of DT.

References

Adams, Robert M. (1977) "Middle Knowledge and the Problem of Evil," *American Philosophical Quarterly* 14: 109–117.

Bennett, Jonathan. (2003) *A Philosophical Guide to Conditionals.* Oxford: Clarendon Press.

Berntson, Daniel. (2010) "Saving Ordinary Counterfactuals." ms.

DeRose, Keith. (1991) "Epistemic Possibilities," *Philosophical Review* 100: 581–605.

____________. (1994) "Lewis on 'Might' and 'Would' Counterfactual Conditionals," *Canadian Journal of Philosophy* 24: 413–418.

____________. (1998) "Simple 'Might's, Indicative Possibilities, and the Open Future," *Philosophical Quarterly* 48: 67–82.

____________. (1999) "Can It Be That It Would Have Been Even Though It Might Not Have Been?" *Philosophical Perspectives* 13: 385–413.

von Fintel, Kai, and Anthony S. Gillies. (2008) "CIA Leaks." *Philosophical Review* 117: 77–98.

Hájek, Alan. (ms) "Most Counterfactuals are False."

Hawthorne, John. (2003) "Chance and Counterfactuals," *Philosophy and Phenomenological Research* 70 (2003): 396–405. Repr. in Hawthorne (2006).

____________. (2004) *Knowledge and Lotteries.* Oxford: Clarendon Press.

____________. (2006) *Metaphysical Essays.* Oxford: Clarendon Press

Lewis, David. (1973a) *Counterfactuals.* Oxford: Blackwell Publishing.

____________. (1973b) "Counterfactuals and Comparative Possibility," *Journal of Philosophical Logic* 2: 418–46. Repr. in Lewis (1986).

____________. (1979) "Counterfactual Dependence and Times Arrow," *Noûs* 13: 455–476. Repr. in (1986).

____________. (1986) *Philosophical Papers*, vol. II. Oxford: Oxford University Press.

Moss, Sarah. (ms) "Constraining Credences in Counterfactuals."

Prior, Arthur N. (1957) *Time and Modality.* Oxford: Clarendon Press.

Sosa, Ernest. (1999) "How to Defeat Opposition to Moore," *Philosophical Perspectives* 13: 142–153.

Stalnaker, Robert. (1975) "Indicative Conditionals." Repr. in Stalnaker (1999).

____________. (1984) *Inquiry.* Cambridge: MIT Press.

____________. (1999) *Context and Content.* Oxford: Oxford University Press.

Stanley, Jason. (2005) *Knowledge and Practical Interests.* Oxford: Clarendon Press.

Williams, J. Robert G. (forthcoming) "Defending Conditional Excluded Middle." *Noûs.*