

Future Conditionals and DeRose's Thesis

DAVID BARNETT

University of Colorado at Boulder

David.Barnett@Colorado.edu

Imagine that you are deciding whether to read this paper. It might seem reasonable for you to base your decision on your confidence (i) that, if you read this paper, you *will* become a better person. It might also seem reasonable for you to base your decision on your confidence (ii) that, if you were to read this paper, you *would* become a better person. Is there a difference between (i) and (ii)? If so, are you rationally required to base your decision on one of the two conditionals, rather than the other? Keith DeRose (2010) proposes a provocative set of answers to these questions. In this paper, I argue that DeRose's answers are incorrect, and I defend a rival set of answers.

A birthday cake sits before you, covered in small flames. You notice an empty box labelled 'Trick Candles' across the room. You think to yourself, 'It's very likely that, if I blow on these candles, I will look like a fool. I do not want to look like a fool on my birthday'. You refuse to blow on the candles. Here you deliberate on the basis of the following conditional:

Will that, if you blow on the candles, you *will* look like a fool.

Your high degree of confidence in **Will**, together with your desire not to look like a fool, leads you not to blow on the candles. Call this Scenario 1.

Scenario 2: A birthday cake sits before you, covered in small flames. You notice an empty box labelled 'Trick Candles' across the room. You think to yourself, 'It's very likely that, if I *were* to blow on these candles, I *would* look like a fool. I do not want to look like a fool on my birthday'. You refuse to blow on the candles. Here you deliberate on the basis of **Would**:

Would that, if you were to blow on the candles, you *would* look like a fool.

Your high degree of confidence in **Would**, together with your desire not to look like a fool, leads you not to blow on the candles.

Scenarios 1 and 2 are hard to tell apart, for they differ only with respect to **Will** and **Would**, which are themselves hard to tell apart. At least initially, it would seem that your degrees of confidence in **Will** and **Would** ought to be the same. And so it might seem that your attitudes toward **Will** and **Would** are equally good candidates for guiding your decision.

But now consider Scenario 3. A birthday cake sits before you, covered in small flames. You do not have any physical evidence that your candles are trick. However, you do know that your candles have been prepared by a

prankster who uses trick candles if, and only if, she predicts that the recipient of her candles will blow on them. You know that the prankster has a remarkably strong – though not perfect – record of predicting who will blow on their candles. Indeed, on your previous 25 birthdays she has used ordinary candles, based each instance on a correct prediction that you would not blow on them. On each occasion, she has blown out your candles after you declined to blow on them. This year you are determined to avoid this humiliation, and yet you remain equally determined not to look like a fool by blowing on trick candles. You enter the deliberative phase completely undecided as to whether you will blow: you are 50% confident that you will blow and 50% confident that you will not blow. You think to yourself:

Given the prankster's record of using ordinary candles the past 25 years, it's more likely than not that these candles are ordinary. Still, if I blow on them, it's very likely that I will look like a fool. For the prankster is good at predicting who will blow on their candles. Supposing, then, that I blow on them, that is good evidence that the prankster predicted that I would blow on them. It is thus good evidence that they are trick, and that I will look like a fool. Hence my confidence that, if I blow on these candles, I will look like a fool.

This confidence should not however dissuade me from blowing. For it merely reflects my confidence that my blowing would constitute *evidence* that the candles are trick, and I am not averse to creating *evidence* that the candles are trick. I am averse to blowing on trick candles! While my blowing might bear evidentially on whether the candles are trick, it cannot *influence* whether they are trick. Thus, I'm confident that, if I *were* to blow on these very candles, their trick status would be no different from what it actually is. Because I'm confident that their trick status is actually *not* trick, I'm confident that, if I were to blow on them, they would not be trick and I would not look like a fool (even though, in the act of blowing, I would create evidence that I was about to look like a fool).¹

So, while the evidential connection between my blowing and the trick status of these candles leads me to have a high degree of confidence that, if I blow on them, I will look like a fool, my confidence (i) that these candles are not in fact trick and (ii) that my blowing on them would in no way influence whether they are trick, leads me to maintain a low degree of confidence that, if I *were* to blow on them, I *would* look like a fool. Clearly, this low degree of confidence *should* sway me to blow on the candles.²

¹ Initially, it might seem that you can reason in similar fashion to the conclusion that, if you blow on these candles, you will *not* look like a fool: 'While my blowing might bear evidentially on whether these candles are trick, it cannot *influence* whether they are trick. Thus, I'm confident that, if I blow on them, their trick status will be no different from what it is now. Because I'm confident that their trick status is now *not* trick, I'm confident that, if I blow on them, they will not be trick and I will not look like a fool.' This is bad reasoning. Given the prankster's track record, you should be confident that, if you blow, the candles are *now* trick. Because your blowing will not influence their trick status, you should be confident that, if you blow, the candles *will be* trick.

² A puzzle arises upon deciding to blow on the candles. Your new confidence that you will blow, together with your longstanding confidence that, if you blow, the candles are trick,

You lean forward and blow on the candles. Here your attitudes toward **Will** and **Would** do not seem to be equally good candidates for guiding your decision: your low degree of confidence in **Would**, rather than your high degree of confidence in **Will**, seems to be your best guide to deciding whether to blow on the candles.

Examples of this sort motivate causal decision theory.³ When choosing among a number of options, decision theory recommends choosing the option with the greatest expected utility. The expected utility of an option depends on the probabilities and utilities of the possible outcomes of the option. But there are two divergent ways to consider the probability of a possible outcome, given an option. First, one can consider the probability that the outcome will obtain, given that the option obtains. For instance, one can consider the probability that these candles will be trick, given that one blows on them. This probability reflects an *evidential* dependence between the option and the possible outcome. Alternatively, one can consider the probability that the outcome *would obtain*, given that the option *were to obtain*. For instance, one can consider the probability that these candles would be trick, given that one were to blow on them. This probability reflects a *metaphysical* dependence—typically causal—between the option and the possible outcome. According to causal decision theory, it is the latter, and not the former, probability that should be used in calculating the expected utility of an option.

So, while **Will** and **Would** are initially hard to tell apart, on reflection there is an important difference between them, one that has significant implications for decision theory.

Keith DeRose (2010) disagrees. To see his point of disagreement, it is helpful to distinguish sentences from the ideas they express. A sincere and literal utterance of the English sentence, ‘Tofu is fun’, expresses the idea that tofu is fun. A single idea can be expressed by utterances of different sentences from different languages. Japanese, Italian, and Swahili all have sentences that can be used to express the idea that tofu is fun. Philosophers sometimes use ‘literal content’, ‘proposition literally expressed’, or ‘sense’ to designate the idea literally expressed by an utterance. In this paper, I deliberately avoid the term ‘proposition’ to avoid the presumption that what is expressed by a sentence has any of the features typically associated with propositions so-called, for instance, objective truth-conditions. Instead, I use the term ‘idea’. In our hypothetical scenario, you consider two ideas: first, that if you blow on

should breed new confidence that the candles *are* trick. In turn, this confidence should breed new confidence that, if you were to blow, the candles would be trick. It should thus breed new confidence that you have made the wrong decision. In other words, your decision is itself evidence for you that you have made the wrong decision. Of course, a decision in the opposite direction—to not blow—would also be evidence that you had made the wrong decision. Should you change your mind and not blow, based on your decision to blow? It would seem not. How best to accommodate this intuition falls within the scope of deliberation-probability dynamics, the theory of how probability evolves as one deliberates (see e.g. Skyrms 1982). One idea is that, for the purposes of deciding whether to blow, you should set aside your view on whether you will blow, by *supposing* that it is 50% likely that you will blow.

³ For general treatments of causal decision theory, see Gibbard and Harper 1978, Lewis 1981, Skyrms 1982, Joyce 1999, and Weirich 2001.

the candles, you will look like a fool; second, that if you were to blow on the candles, you would look like a fool. I named these ideas ‘**Will**’ and ‘**Would**’, respectively, and I illustrated an important difference between them.

On DeRose’s thesis, **Will** is the very same idea as **Would**. To be sure, DeRose maintains that the English language contains two importantly different *sentences* for expressing this idea:

Will_s ‘If the reader blows on the candles, the reader will look like a fool’

Would_s ‘If the reader were to blow on the candles, the reader would look like a fool’.

On DeRose’s thesis, sentences **Will_s** and **Would_s** are governed by different rules of use, despite the fact that they express the very same idea (in a fixed context). Whenever a speaker’s grounds for accepting this idea are not good grounds for deliberating on whether to realize the antecedent of the idea (whether to blow on the candles), the speaker is prohibited by the rules of English from expressing the idea with **Would_s**, but not necessarily from expressing it with **Will_s**. An utterance of **Would_s** thus indicates to speaker’s audience that the speaker’s grounds for accepting the shared idea between **Will_s** and **Would_s** are good grounds for deliberating on whether to realize the antecedent of this idea. In Scenario 3, we concluded that your grounds for accepting **Will** are not good grounds for you to deliberate on whether to blow on the candles. On DeRose’s thesis, it follows that any speaker who accepts **Will** on these same grounds is prohibited by the rules of English from uttering **Would_s**, but not necessarily from uttering **Will_s**. Perhaps this would explain why it might seem inappropriate for someone in an epistemic situation similar to yours to utter **Would_s**, but not inappropriate for such a person to utter **Will_s**.

There is much, however, that DeRose’s thesis cannot explain. For instance, it cannot explain how, in Scenario 3, you have different degrees of confidence in **Will** and **Would**—ideas that are identical on DeRose’s thesis. This alone would seem sufficient to show that the thesis is false. Nevertheless, given the initial similarity between **Will** and **Would**, and given that I have not yet examined DeRose’s reasons in favor of his thesis, it is worth taking a deeper look at the issue. That is the instrumental aim of this paper. In pursuing it, I hope to satisfy my primary aim, which is to sketch a more credible picture of the relationship between ideas like **Will** and those like **Would**.

First I formulate DeRose’s thesis in general terms and show how it entails that **Will** = **Would** (§1). Then I consider DeRose’s reasons for his thesis and argue that none is a good one (§2). Next I give five more reasons to reject DeRose’s thesis (§3). I conclude that DeRose’s thesis is false, that there is a credible alternative picture of the relationship between **Will** and **Would**, and that causal decision theorists may continue to utilize the distinction between these two types of ideas to express their view that the relevant connection between an option and its possible outcomes is causal, not evidential.

1. DeRose's Thesis

DeRose's thesis concerns the relation and nature of two types of future conditional ideas and sentences. Future conditional ideas of the first type are those, such as **Will**, that can be referred to by that-clauses of the form, 'that, if such and such happens, so and so *will* be the case'. I call these 'if-will ideas'. I call the corresponding sentences, such as **Will**s, 'if-will sentences'. Future conditional ideas of the second type are those, such as **Would**, that can be referred to by that-clauses of the form, 'that, if such and such *were* to happen, so and so *would* be the case'. I call these 'if-would ideas'. I call the corresponding sentences, such as **Would**s, 'if-would sentences'.

DeRose's thesis has two parts. The first is a claim about the *relation* between if-would sentences and their if-will counterparts. The claim is that if-would sentences have the very same meaning as their if-will counterparts, with two exceptions (p. 37). First, an if-would sentence can be used to call attention to the fact that its antecedent is false. Second, an if-would sentence should never be uttered on *deliberationally useless* grounds.⁴ DeRose maintains that grounds are deliberately useless just in case they are of the following *backtracking* sort:

Suppose that I blow on these candles. Then, most likely, the state of the world prior to my blowing is different from how, suppositions aside, I take it to be. Suppositions aside, I'm confident that the candles are ordinary. Supposing that I blow, I'm confident that they are trick. This suppositional shift in my view of the world *prior* to the supposed event of my blowing constitutes *backtracking* grounds for my confidence that, if I blow on these candles, I will look like a fool.

The second part of DeRose's thesis concerns the *nature* of if-would and if-will sentences and ideas. DeRose maintains that there are two importantly different kinds of *past* conditional sentences and ideas. He distinguishes, for instance, between the idea expressed by the sentence 'If Oswald didn't shoot Kennedy, someone else did' (an idea most of us place a high degree of confidence in) and the idea expressed by the sentence 'If Oswald hadn't shot Kennedy, someone else would have' (an idea most of us place a lower degree of confidence in). I call ideas of the former type 'if-was ideas' and ideas of the latter type 'if-would-have ideas'. I call the corresponding sentences 'if-was sentences' and 'if-would-have sentences'. The second part of DeRose's thesis is that if-would and if-will sentences and ideas are of the same kind as if-was sentences and ideas. Following standard practice, DeRose classifies if-was sentences and ideas as *indicative* and if-would-have sentences and ideas as *subjunctive*. In these terms, the second part of DeRose's thesis is that all future conditional sentences and ideas are indicative in nature. Because future conditional ideas are useful for deliberation, DeRose concludes that 'the conditionals of deliberation are indicatives' (p. 1).

⁴ DeRose makes an exception to this rule in cases where a speaker explains her grounds to her audience (fn. 31).

To see that the first part of DeRose's thesis entails that if-will ideas are identical to their if-would counterparts, suppose that if-will ideas are *not* identical to their if-would counterparts. Then, in a fixed context, if-will sentences express different ideas – that is, they have different literal content – from their if-would counterparts. Given a fixed context, a difference in literal content between two sentences entails a difference in the meanings of the two sentences (where 'meaning' is construed broadly to include features of a sentence that combine with context to determine its literal content). By the first part of DeRose's thesis, there are only two differences in the meanings of if-will and if-would sentences. Hence, one of these two differences must entail a difference in literal content. But, as I will now show, neither entails a difference in literal content. So the first part of DeRose's thesis entails that there is no difference in literal content between if-will sentences and their if-would counterparts, and thus that if-will ideas are identical to their if-would counterparts.

The first difference in meaning is that an if-would sentence can be used to call attention to the fact that its antecedent is false. This difference does not entail a difference in literal content. Two sentences can have the same literal content even though, by virtue of a difference in their respective meanings, only one of the two can be used to call attention to a given fact. For instance, 'Nevertheless, Cooper will win the race' and 'Cooper will win the race' plausibly have the same literal content, namely, that Cooper will win the race, even though only the former sentence can be used to call attention to the fact that there is a contextually salient reason to doubt that Cooper will win the race. So the first of the two differences in meaning does not entail a difference in literal content.

The second difference is that an if-would sentence should never be uttered on deliberately useless grounds. This difference does not entail a difference in literal content either. For two sentences can have the same literal content even though only one of the two is prohibited by its meaning from being used appropriately in a given context. Philosophers typically describe such a scenario by saying that the two sentences have different *assertability* conditions, despite having the same literal content. For instance, Frank Jackson maintains that each of the following pairs comprises two sentences with different assertability conditions but the same literal content: (i) 'A but B' and 'A and B'; (ii) 'Nevertheless A' and 'A'; and (iii) 'If A, then B' and 'Either not-A or B'. Jackson says,

In the widest sense of 'meaning', ['If A, then B' and 'Either not-A or B'] do not mean the same. But ... they agree in sense or literal content. The extra element is that in using ['If A, then B'] you explicitly signal the robustness of ['Either not-A or B'] with respect to ['A'], and this element affects assertion conditions without affecting [literal content]. (1979, pp. 576-577)

DeRose characterizes the second alleged difference between the meanings of if-will and if-would sentences as a difference in assertability conditions (p. 38). Like the first difference, this difference does not entail a difference in literal content.

On DeRose's thesis, there are no further differences in meaning between if-will and if-would sentences. Hence, DeRose's thesis entails that if-will sentences have the same literal content as their if-would counterparts, and thus that if-will ideas are identical to their if-would counterparts. It thus entails that **Will = Would**.

In personal communication, DeRose has proposed an alternative interpretation of his printed thesis, on which the second difference in meaning is not a *mere* difference in assertability conditions, but rather a deeper difference that manifests itself as (i) a difference in assertability conditions and (ii) a difference in literal content. On this interpretation, DeRose's thesis entails that if-would ideas are distinct from their if-will counterparts, even though both are properly classified with if-was ideas as indicative. In §3 I show that if-would ideas differ from if-will ideas in the same characteristic ways that if-would-have ideas differ from if-was ideas; I thereby show that if-would ideas are not properly classified with if-was ideas, as indicative. Ultimately, then, I show that DeRose's thesis is false on either interpretation. So the interpretive question at hand is ultimately moot.

Still, it is worth giving a few of my reasons for sticking with my original interpretation.⁵ To begin with, in presenting the second alleged difference in meaning, DeRose makes no mention of any difference other than a difference in assertability conditions. So, if his reader is expected to infer that there is an additional implicit difference that gives rise to a difference in literal content, there should be ample textual support for this inference. But nowhere in the paper does DeRose say that there is a further difference in meaning, and nowhere in the paper does he say that if-will sentences have different literal content from their if-would counterparts. To the contrary, DeRose makes a range of claims and assumptions that do not comport with the idea that if-will sentences have different literal content from their if-would counterparts. Here are four. First, DeRose assumes that the antecedent and consequent of an if-would sentence express the very same ideas as the antecedent and consequent of the counterpart if-will sentence (fn. 18; for discussion, see my fn. 10). On this assumption, contrary to the lesson of our opening example, there is no difference between supposing what is expressed by the antecedent of **Wills**, namely, *that you blow on these candles*, and supposing what is expressed by the antecedent of **Woulds**, namely, *that you were to blow on these candles*. From here it is a small step to the conclusion that **Wills** and **Woulds** have the same literal content (i.e. that **Will = Would**). Second, in various places, DeRose assumes that one's grounds for an arbitrary if-will idea must be the same as one's grounds for the counterpart if-would idea (e.g. fn. 31, pp. 37-38). He tells us that an if-would sentence should never be uttered on deliberately useless grounds, but then he identifies the conditions for having such grounds with the conditions for having deliberately useless grounds for uttering the counterpart if-will sentence (p. 38). This would make sense on the hypothesis

⁵ Various interpretative challenges arise from the fact that DeRose does not explicitly distinguish ideas from sentences. He uses 'conditional' indiscriminately, referring sometimes to sentences and sometimes to ideas. For instance, when he speaks of the conditionals of deliberation, and of believing conditionals, he is referring to ideas, since it is ideas, and not sentences, that are the objects of deliberation and belief. And when he speaks of the semantics of conditionals, he is referring to sentences, since it is sentences, and not ideas, that have semantic content.

that if-will sentences have the same content as their if-would counterparts. Third, DeRose explicitly draws the distinction between meaning and content (fn. 27), but he then formulates his thesis in terms of two differences in *meaning* that play no role in fixing content. Fourth, DeRose characterizes if-would sentences as the result of ‘souping-up’ if-will sentences for the purpose of ‘clearly marking out conditionals as based on the right sorts of grounds to be deliberationally useful’ (pp. 35-38). In line with his explicit thesis, this characterization suggests that a speaker who accepts an if-will idea on deliberationally useful grounds may ‘mark’ the idea as being so grounded by asserting it, not with the corresponding if-will sentence, but rather with the counterpart if-would sentence. There is more in DeRose’s paper that does not comport with the idea that if-will ideas are distinct from their counterpart if-would ideas, but I will not further belabour the point, given that I will ultimately argue that his thesis is false on either interpretation.⁶

2. DeRose’s Reasons for DeRose’s Thesis

In this section I argue that none of the three main reasons that DeRose gives in favor of his thesis is a good one.⁷

DeRose’s first reason for his thesis is that, in many contexts, if-will sentences and their if-would counterparts ‘can at least seem to be something like equivalent’ (p. 9). Applied to Scenarios 1 and 2, **Wills** and **Woulds** might at least seem to be something like equivalent. Should this fact compel us to reverse our verdict that **Will** ≠ **Would**? Clearly not. For there is no reason to think that we failed to respect this fact when, upon considering Scenario 3, we reached the verdict. Initially, **Will** and **Would** were hard to tell apart, but then we found a scenario that made salient an important difference between the two.

Moreover, given our conclusion about how **Will** differs from **Would**, it is easy to explain their initial resemblance. Whereas confidence in **Will** reflects a commitment to an *evidential* dependence between your option to blow on the candles and the possible outcome of your looking like a fool, confidence in **Would** reflects a commitment to a *causal* dependence between the two. Because evidence for a causal dependence between two possibilities often constitutes an evidential dependence between the two possibilities, evidence

⁶ In conversation, DeRose has suggested (i) that, when his alleged assertability conditions of if-would sentences are not met, an utterance of the if-would sentence qualifies as objectively incorrect; and (ii) that, because an utterance of the counterpart if-will sentence might, in the same context, qualify as objectively correct, and because the objective correctness/incorrectness of an utterance depends ultimately on whether its content is objectively correct/incorrect, the two sentences might in such a context have different content. In response, I do not see how having deliberationally useless grounds for asserting something could make the thing asserted – that is, the content of the assertion – objectively incorrect (unless, of course, what one asserts is that one’s grounds for making the assertion are not deliberationally useless).

⁷ I will not address DeRose’s argument that if-will sentences should be treated semantically like if-was sentences, as indicatives; I shall grant its conclusion.

for **Would** often constitutes evidence for **Will**.⁸ For illustration, consider Scenarios 1 and 2. In these scenarios, you notice an empty box labelled 'Trick Candles'. You thus have evidence that the candles are trick. Presumably, you also have evidence that blowing on trick candles tends to cause the blower to look like a fool. Hence, you have evidence for **Cause**:

Cause that your blowing on the candles would cause you to look like a fool.

Your confidence in **Would** is based on this evidence. But so is your confidence in **Will**. For suppose that you blow on the candles. In the context of this supposition, you have the very same evidence for **Cause** that you have outside the supposition. This evidence thus constitutes evidence for you – within the context of supposing that you blow – that your blowing *will* cause you to look like a fool, and thus that you *will* look like a fool. In Scenarios 1 and 2, your evidence for **Would** constitutes your evidence for **Will**. This is why you place the same degree of confidence in the two ideas, and it is also why the two ideas are so hard to tell apart.

In Scenario 3, the symmetry is broken. Here your evidence regarding the causal relation between your blowing on the candles and your looking like a fool is sensitive to the supposition that you blow on the candles. Recall that, in this scenario, the prankster's record of using ordinary candles the past 25 years constitutes some evidence for you that the candles are *not* trick, and so you have some evidence *against* **Cause**. Based on this evidence, you maintain a relatively *low* degree of confidence in **Would**. However, on the supposition that you blow on the candles, you have a new source of evidence regarding **Cause**: you now have evidence that the prankster used trick candles, and this evidence undermines your original evidence against **Cause** to the point that you now have strong evidence in favor of **Cause**. You thus have strong evidence – in the context of supposing that you blow – that your blowing *will* cause you to look like a fool, and that you therefore *will* look like a fool. Based on *this* evidence, you maintain a *high* degree of confidence in **Will**. So, in Scenario 3, your evidence regarding **Would** is significantly different from your evidence regarding **Will**. This is why you place different degrees of confidence in the two ideas, and it is also why the two ideas are no longer hard to tell apart.

So, in addition to demonstrating that **Will** ≠ **Would**, Scenario 3 illustrates a difference between the two ideas that explains their initial resemblance. I conclude that DeRose's first reason for his thesis – that if-will sentences 'can at least seem to be something like equivalent' to their if-would counterparts – is not a good one.

DeRose's second reason for his thesis is that if-will and if-would sentences typically have a feature in common with if-was sentences that is allegedly

⁸ Note that even though **Will** and **Would** are distinct ideas with distinct antecedents – supposing *that you blow on the candles* is different from supposing *that you were to blow on the candles* – their antecedents correspond to the same *possibility*, namely, the possibility of your blowing on the candles. Supposing that you blow on the candles, the possibility of your blowing *will* be realized. Supposing that you were to blow on the candles, the same possibility *would* be realized.

lacked by if-would-have sentences. The feature is *having assertability conditions determined by the following principle*:

Cert 'If A, then C' is *assertable* iff adding A as a certainty to one's belief set would put one in a warranted position to assert that C.
(p. 15)

But **Cert** does not give the assertability conditions of if-was sentences (or any type of conditional sentence, for that matter). To see that it does not, suppose that I missed your party but heard from a reliable source that the prankster used trick candles. Suppose also that I have no reason to doubt this source, and that I have good reason to think that people who blow on trick candles look like fools. Then I am warranted in sincerely uttering, 'If the reader blew on the candles, then the reader looked like a fool.' It does not follow that my adding the idea that you blew on the candles to my belief set as a certainty would warrant me in asserting that you looked like a fool. For imagine that I make the addition arbitrarily: for no reason at all, I become certain that you blew on the candles. Do I thereby gain warrant for asserting that you looked like a fool? Of course I do not. For I am not warranted in making the addition in the first place. So **Cert** does not give the assertability conditions of if-was sentences.

Perhaps I have misinterpreted **Cert**. Perhaps 'adding A as a certainty to one's belief set' means 'adding A as a *warranted* certainty to one's belief set'. On this interpretation, new problems arise for **Cert**. First, it might be impossible for me to gain warrant for becoming certain that you blew on the candles. What evidence or argument could I possibly have that would warrant my being *absolutely certain* that you blew on the candles? Plausibly, there is none, for whatever evidence or argument I might have for being confident that you blew on the candles, there will always be some sceptical scenario that prevents this confidence from rationally reaching a degree of 1. Second, even if it is possible for me to gain warrant for being certain that you blew on the candles, any such warrant might undermine my warrant for being confident that, if you blew on the candles, you looked like a fool, in which case I would *not* become warranted in asserting that you looked like a fool. So, even if 'certainty' is interpreted as 'warranted certainty', **Cert** does not provide the assertability conditions of if-was sentences.

Perhaps I have twice misinterpreted **Cert**. Perhaps 'adding A as a certainty to one's belief set' means '*hypothetically* adding A as a certainty to one's belief set', where this is not equivalent to either of the preceding two interpretations: it does not mean 'if one were to add A as a certainty to one's belief set' and it does not mean 'if one were to add A as a warranted certainty to one's belief set'. What it means, rather, is something like, 'supposing that A', or 'relative to the supposition that A'. On this interpretation, **Cert** is equivalent to **Sup**:

Sup 'If A, then C' is *assertable* iff, relative to the supposition that A, one is warranted in asserting that C.

Sup provides a plausible account of the assertability conditions of if-was sentences. Moreover, there is some evidence that DeRose intends **Cert** to be

interpreted as **Sup**: later in his paper, DeRose explicitly uses **Sup** to test whether various conditionals are assertable.⁹ So it is worth investigating whether DeRose's second reason for his thesis is a good one if **Cert** is interpreted as **Sup**. To this end, I will hereafter interpret DeRose's claims about **Cert** as claims about **Sup**. On this interpretation, DeRose's second reason is that if-will and if-would sentences typically have a feature in common with if-was sentences that is allegedly lacked by if-would-have sentences: namely, *having assertability conditions determined by Sup*.

Is this a good reason to endorse DeRose's thesis? No. For, while it is true that if-will, if-would, and if-was sentences obey **Sup**, it is false that if-would-have sentences fail to obey **Sup**. DeRose takes the following if-would-have sentence to violate **Sup**:

(K*) 'If Oswald hadn't shot Kennedy, then someone else would have'.

But in so far as any conditional sentence seems to obey **Sup**, (K*) does. Any speaker who is warranted in asserting *that, if Oswald hadn't shot Kennedy, someone else would have* (by sincerely uttering (K*)) is warranted, relative to the supposition that Oswald hadn't shot Kennedy, in asserting that someone else would have. And the converse holds as well.

So why does DeRose think that (K*) violates **Sup**? The answer is that he misapplies **Sup** to (K*). For 'A' and 'C' in the left-hand side of **Sup**, DeRose correctly substitutes the clauses embedded in the if-clause and the then-clause of (K*), respectively. But for 'A' and 'C' in the right-hand side of **Sup**, he substitutes the clauses embedded in the if-clause and the then-clause of (K):

(K) 'If Oswald didn't shoot Kennedy, then someone else did'.

Obviously, *this* substitution leads to an unacceptable biconditional: (K*) is assertable iff, relative to the supposition that Oswald didn't shoot Kennedy, one is warranted in asserting that someone else shot Kennedy. If we systematically misapply **Sup** to if-would-have sentences in this way, we arrive at DeRose's mistaken conclusion that if-would-have sentences systematically violate **Sup**.¹⁰ When properly applied, however, there is no

⁹ E.g. pp. 28-99, DeRose investigates whether various hypothetical speakers are positioned to assert various conditionals. In each case, DeRose imagines that the speaker *supposes* the antecedent of the conditional and then reasons under that supposition. For instance, DeRose writes, 'He then adds to these beliefs the supposition of the antecedent of the conditional – he supposes that Pete will play – and then reasons forward in time ...' (p. 29). According to DeRose, the result of this exercise determines whether the speaker is warranted in asserting the conditional.

¹⁰ DeRose says that he is making the 'standard (at least among philosophers) and reasonable assumption' that the antecedent of (K*) is the clause embedded in the if-clause of (K) (fn. 18). But if 'antecedent' means 'clause embedded in the if-clause of a conditional sentence', then this assumption is patently false (which is not to say that it is not standard among philosophers). But what else could 'antecedent' mean? If 'antecedent' means 'idea expressed by the clause embedded in the if-clause of a conditional' then, again, the assumption is false: the idea expressed by the clause embedded in the if-clause of (K*) is given by the nominalization of this clause – *that Oswald hadn't shot Kennedy* – and not by the nominalization of the if-clause of (K) – *that Oswald didn't shoot Kennedy*. To be sure, this idea is not the sort of idea that has a truth value, nor is it the sort of idea that could be believed, but it is the sort of idea that could be *supposed*, which is all that is needed for **Sup**. Recall from our

reason to doubt that **Sup** is a fully general principle governing the assertability conditions of *all* conditional sentences. Thus, it is false that the general principle which determines the assertability conditions of if-will, if-would, and if-was sentences fails to determine the assertability conditions of if-would-have sentences. DeRose's second reason in favor of his thesis is not a good one.

Before proceeding to his third reason, it is perhaps worth considering whether, behind DeRose's misapplication of **Sup**, there lies a legitimate reason in favor of his thesis. DeRose misapplies **Sup** to if-would sentences in a similar manner to the way he misapplies it to if-would-have sentences: for 'A' and 'C' in the left-hand side of **Sup**, he substitutes the clauses embedded in the if-clause and the then-clause of the given if-would sentence, but for 'A' and 'C' in the right-hand side of **Sup**, he substitutes the clauses embedded in the if-clause and the then-clause of the counterpart if-will sentence. In effect, then, the question that DeRose addresses – when he claims to be addressing whether if-would-have and if-would sentences obey **Sup** – is whether if-would-have and if-would sentences have the very same assertability conditions as their counterpart if-was and if-will sentences, respectively. DeRose's answer to *this* question is that, whereas if-would-have sentences do not have the very same assertability conditions as their counterparts, if-would sentences typically do have the very same assertability conditions as their counterparts.

Is this a good reason to endorse DeRose's thesis? No. First of all, we already have a plausible rival explanation of why, in many contexts, if-would sentences are assertable just in case their counterpart if-will sentences are assertable: even though if-would ideas are distinct from their if-will counterparts, evidence for an if-would idea often constitutes evidence for its if-will counterpart. In §3 I will show that the same relationship exists between if-would-have ideas and their if-was counterparts: evidence for an if-would-have idea often constitutes evidence for its if-was counterpart. So, contrary to what DeRose suggests, there is the same kind of overlap between the assertability conditions of if-would-have sentences and their if-was counterparts as there is between the assertability conditions of if-would sentences and their if-will counterparts. To be sure, for reasons to be explained in §3, the overlap may be greater between if-would and if-will sentences, but this difference is not a difference in kind. I conclude that behind DeRose's misapplication of **Sup** there lies no good reason in favor of his thesis.

It is worth emphasizing that my criticism of DeRose's second reason does not rest on the question of which interpretation of **Cert** is correct. For

opening example that the act of supposing *that you will blow on the candles* is importantly different from the act of supposing *that you were to blow on the candles*. The idea that you will blow on the candles \neq the idea that you were to blow on the candles. In any case, whatever 'antecedent' means, it is clear that substituting different clauses for 'A' throughout **Sup** is a misapplication of the rule. (Perhaps DeRose's mistake arises because he is working directly with **Cert**, which requires that the idea expressed by the antecedent of a conditional be the sort of idea that one can be certain of; because the idea expressed by the antecedent of (K*) is not the sort of idea one can be certain of – but only the sort of idea one can suppose – DeRose substitutes for this idea an idea that *is* the sort of idea one can be certain of, namely, the idea expressed by the antecedent of (K).)

DeRose's second reason ultimately rests on the premiss that *whatever* general principle governs the assertability conditions of if-was sentences, the same principle governs the assertability of if-will and if-would sentences, but not of if-would-have sentences. And we have found a general principle that governs the assertability conditions of if-was sentences, but which also governs the assertability conditions of if-will, if-would, *and* if-would-have sentences – namely, **Sup**. So, irrespective of the correct interpretation of **Cert**, a key premiss behind DeRose's second reason is false.

DeRose's third reason for his thesis is that it is supported by a version of Alan Gibbard's 'Sly Pete' story (1981). The primary role of this story, in DeRose's argument, is to highlight certain differences between if-will sentences and their if-would counterparts, and to thus provide a platform to explain these differences in a way that is consistent with DeRose's thesis. Of course, it is unlikely that a story whose intended role is to illustrate differences between if-will and if-would sentences would present us with a good reason to overturn our opening verdict that the ideas respectively expressed by these sentences are different. But I will examine the story anyway, just to rule out this unlikely possibility.

In Gibbard's original story, Sly Pete is playing cards against Mr. Stone. It is Pete's turn to play or quit. To simplify matters, I will assume that the game excludes the possibility that Pete and Mr. Stone have equally strong hands, and I will also assume that Pete *wins* just in case he either plays with the stronger hand or quits with the weaker hand. Pete's aim, then, is to win. Gibbard's henchman, Zack, peeks at Mr. Stone's hand, sees that it is very strong, and signals its content to Pete. Zack now knows (a) and (b):

- (a) that Pete knows who has the stronger hand
- (b) that Pete is determined to quit on knowing that Mr. Stone has the stronger hand.

Gibbard's second henchman, Jack, sees both hands and now knows (c):

- (c) that Mr. Stone has the stronger hand.

The room is cleared of everyone but Pete and Mr. Stone. Zack reasons from (a) and (b) to the idea expressed by his sincere utterance of (d):

- (d) 'If Pete played, he won'.

Jack reasons from (c) to the idea expressed by his sincere utterance of (e):

- (e) 'If Pete played, he didn't win'.

Because Zack and Jack reason only from true premisses, and because neither is guilty of bad reasoning, Gibbard concludes that, if (d) and (e) express ideas with objective truth-values, then (i) both express objectively true ideas and thus (ii) the two ideas expressed are compatible. Gibbard raises some doubts about (ii) and, partly on this basis, concludes that (d) and (e) do not express ideas with objective truth-values.

One view consistent with this conclusion is the suppositional view of conditionals, on which if-clauses function to express suppositions, just as the dictionary says. On this view, a conditional statement is essentially a statement of one thing relative to the supposition of another: to state *that, if A, C* is to state *that C* relative to the supposition *that A*. (On the rival categorical view, to state *that, if A, C* is to state something, not relative to any hypothesis or supposition, but categorically.) The suppositional view makes sense of Gibbard's example. Zack reasons, without any error, from true premisses to the conclusion *that, if Pete played, he won*. Jack reasons, without any error, from true premisses to the incompatible conclusion *that, if Pete played, he didn't win*. How can this be? The answer is that both conclusions are conclusions of things *relative* to a false supposition. Relative to the false supposition that Pete played, Zack draws the conclusion *that Pete won*, whereas Jack draws the incompatible conclusion *that Pete didn't win*. Combined with one set of truths, the false supposition that Pete played entails that Pete won; combined with a different set of truths, it entails that Pete did not win. Because there is no objectively correct set of truths relative to which the question of whether Pete won is to be answered, on the false supposition that he played, there is no objectively correct answer to the question relative to the supposition. Given their respective epistemic situations, Zack is justified in drawing his conclusion (on the basis of one set of truths), and Jack is justified in drawing his incompatible conclusion (on the basis of a different set of truths).

Indeed, on the suppositional view of conditionals, one thinker might be justifiably *certain* that, if A, C, while another is justifiably *certain* of the incompatible claim that, if A, not-C. For illustration, imagine that George and Saddam are chatting over a cup of tea, and that George begins to wax philosophical:

George: Suppose, as you undoubtedly already believe, that I, George, and you, Saddam, are distinct. Suppose, furthermore, that exactly one of us is conscious. Who is it?

Saddam: It must not be you, George. For I am certain that I am conscious. Hence, I am certain that, if you, George, and I, Saddam, are distinct, and exactly one of us is conscious, then you, George, are *not* conscious.

George: I must disagree. For I am certain that I am conscious. Hence, I am certain that, if I, George, and you, Saddam, are distinct, and exactly one of us is conscious, then I, George, *am* conscious.¹¹

Saddam is justifiably certain that he is conscious. Adding to this certainty the supposition that he is distinct from George and exactly one of the two is conscious, Saddam deduces that George is *not* conscious. So, relative to the given supposition, Saddam is justifiably certain that George is not conscious. Relative to the same supposition, George is justifiably certain that George *is* conscious. So, relative to a single supposition, Saddam is justifiably certain of one answer to a question and George is justifiably certain of an incompatible

¹¹ For further discussion of this example and the suppositional view, see my 2006.

answer to the same question. Hence, there is no objectively correct answer to the question, relative to the supposition. On the suppositional view of conditionals, this is just another way of saying that there is no objectively correct answer to the question of whether, if George is distinct from Saddam and exactly one of the two is conscious, George is conscious. Alternatively, it is another way of saying that the idea *that, if George is distinct from Saddam and exactly one of the two is conscious, George is conscious* has no objective truth-value; for this idea is the idea of a certain predication made relative to a certain supposition, and there is no objective fact of the matter whether this predication is correct relative to the supposition.¹²

On DeRose's thesis, if-will and if-would sentences should be treated semantically like if-was sentences, as indicative. DeRose claims to remain neutral on what this semantic treatment should be, for instance, whether it should comport with the suppositional or the rival categorical view. More specifically, DeRose claims to remain neutral on whether indicative conditional sentences express ideas with objective truth-values. And yet he describes his version of Gibbard's story in ways that are not neutral on this question. He speaks of indicative conditionals as *pieces of information*, of characters as *knowingly informing* others of indicative conditionals, and of characters as *speaking truthfully* by uttering indicative conditionals. These ways of speaking beg the question against certain theories of indicative conditionals, including the very theory that Gibbard is trying to support with his story.¹³ I will therefore recast DeRose's version of Gibbard's story in terms that do not beg the question against these theories. For instance, instead of describing a character as *knowing* a *piece of information* expressed by an indicative conditional sentence, I will describe the character as *being justifiably confident* in the *idea* expressed by the sentence. Also, at points where DeRose changes inconsequential details of Gibbard's story – for instance, where he changes the names of characters, presumably to protect their identities – I will stick with the original details.

The first part of DeRose's story is the same as Gibbard's. Zack peeks at Mr. Stone's hand, sees that it is very strong, and signals its content to Pete. Jack sees both hands and sees that Mr. Stone has the stronger hand. The room is then cleared of everyone but Pete and Mr. Stone. At this point, DeRose embellishes Gibbard's story. Before Pete announces to Mr. Stone whether he will play or quit, Zack sincerely utters (w) to Gibbard:

(w) 'If Pete plays, he will win'.

¹² To be sure, George might be justified in ascribing truth to this idea, and in self-ascribing knowledge of the idea. For, relative to the supposition, George is justifiably confident that it is true that he is conscious and that he knows that he is conscious. But this is not to say that the conditional idea is itself objectively true or that it is an objective fact that George has knowledge of the idea.

¹³ DeRose considers this worry in footnote 35. His response is that he personally endorses one of the theories against which his ways of speaking appear to beg the question, and he believes that these ways of speaking do not beg the question against the theory. Unfortunately, he provides no reason to believe this. I suspect that he is not distinguishing the question of whether one can be justified in ascribing truth to an indicative conditional, on such a theory, from the question of whether such an ascription can qualify as *objectively* correct; see my footnote 12.

And Jack sincerely utters ($\sim w$) to Gibbard:

($\sim w$) 'If Pete plays, he will not win'.

Zack and Jack are epistemically justified in uttering (w) and ($\sim w$), respectively, and their utterances are not based on any factual error.

DeRose gleans the following data from the story:

Datum 1 Given Zack's grounds for uttering (w), it is appropriate for Zack to utter (w) without explanation to Gibbard, but it would be inappropriate, in some respect, for Zack to utter (w) (or its second-person version) without explanation to *Pete*.

Datum 2 Given Zack's knowledge that Mr. Stone has a very strong hand, it would be inappropriate, in some respect, for Zack to utter the following sentence without explanation to *anyone*: (w^*) 'If Pete were to play, he would win'.

DeRose offers an explanation of Data 1 and 2 that is consistent with his thesis. Like his presentation of his story, his explanation begs the question against theories on which if-will and if-was sentences do not express ideas with objective truth-values. Unlike his presentation of the story, however, his explanation does not seem to lend itself to a neutral rephrasal.¹⁴ I therefore present it in its original terms, as follows.

From the fact that Zack and Jack's utterances of (w) and ($\sim w$) are epistemically justified and not based on any factual error, DeRose infers that the ideas expressed by (w) and ($\sim w$) are *pieces of information* that are *known* by Zack and Jack, respectively. Zack's knowledge of the information expressed by (w) depends on backtracking reasoning:

Suppose that Pete plays. Then, most likely, the relevant state of the world *prior* to his playing is different from how, suppositions aside, I take it to be. Suppositions aside, I'm confident that Pete does not have the stronger hand. Supposing that he plays, I know that he *does* have the stronger hand. Hence my knowledge that, if Pete plays, he will win.

Jack's knowledge of the information expressed by ($\sim w$) does not depend on backtracking reasoning:

Suppose that Pete plays. Then, most likely, the relevant state of world *prior* to his playing is no different from how, suppositions aside, I take it to be. Suppositions aside, I know that Pete does not have the stronger hand. Supposing that he plays, I still

¹⁴ DeRose claims that his explanation does lend itself to a neutral rephrasal (p. 25). Why, then, does he not present it in neutral terms to begin with? The answer, I suspect, is that it is more difficult than DeRose suggests to carry out such a rephrasal without losing the spirit of the original explanation.

know that he does not have the stronger hand. Hence my knowledge that, if Pete plays, he will not win.

Theoretically speaking, Zack and Jack could pass their respective pieces of knowledge onto Pete, by sincerely uttering (w) and (\sim w), respectively, to Pete. However, if Pete were to come to know the information expressed by (w) in this way, Pete's knowledge of the information expressed by (w) would depend on Zack's backtracking reasoning. Due to this dependence, the fact would be deliberately useless to Pete: it would not be a fact on which Pete could rationally rely in deciding whether to play or quit. Now, it is part of the meaning of if-will sentences that, if a speaker's knowledge of the information expressed by an if-will sentence depends on backtracking reasoning, then the sentence is not assertable for that speaker *in contexts in which an agent is deliberating on whether to realize the antecedent (in order to influence the consequent)*. This explains why it is appropriate for Zack to utter (w) without explanation to Gibbard, who is not so deliberating, but not to Pete, who is so deliberating. Thus, it explains Datum 1. Finally, it is part of the meaning of if-would sentences that, if a speaker's knowledge of the information expressed by an if-would sentence depends on backtracking reasoning, then the sentence is not assertable without explanation for that speaker *in any context*. Because if-would sentences express the same ideas as their if-will counterparts, the information expressed by (w) is identical to the information expressed by (w*): the fact that, if Pete plays, he will win = the fact that, if Pete were to play, he would win. Because Zack's knowledge of this fact depends on backtracking reasoning, he is in *no contexts* permitted to assert this fact by uttering (w*) without explanation. This explains Datum 2.

Does the preceding explanation constitute a good reason to endorse DeRose's thesis? No. For it is implausible on a number of counts, and there is a rival explanation that is both plausible and consistent with our opening verdict that if-will ideas are distinct from their if-would counterparts.¹⁵

Consider Datum 1. According to DeRose, what explains this datum is a difference in the deliberational statuses of Gibbard and Pete: although Zack could inform either Gibbard or Pete of the truth expressed by (w), (w) is not assertable for Zack to Pete because Pete is deliberating on whether to play. This explanation is implausible on several counts.

First, it is implausible that Zack could inform Pete that, if Pete plays, he will win. To see that it is implausible, consider the analogous question of whether Saddam could inform George of the idea expressed by Saddam's utterance of (\sim G):

(\sim G) 'If you, George, and I, Saddam, are distinct, and exactly one of us is conscious, then you, George, are *not* conscious'.

Saddam's utterance of (\sim G) is epistemically justified and not based on any factual error. DeRose would infer that the idea expressed by (\sim G) is a *piece of information* that is *known* by Saddam. DeRose would further infer that Saddam

¹⁵ I do not claim that the rival explanation is neutral on the correct theory of conditional sentences and ideas. I just claim that it is plausible. So much the worse for theories that are not consistent with it.

could *inform* George of this idea by sincerely uttering ($\sim G$) to George. But this is implausible. First of all, Saddam's grounds for accepting the idea are not good grounds for George to accept it, despite the fact that Saddam and George have common knowledge of these grounds. Saddam and George both know that Saddam is conscious. However, only Saddam is justified in continuing to accept this fact relative to the supposition that George and Saddam are distinct and exactly one of them is conscious. Only Saddam is justified in reasoning as follows:

Suppose that George and I are distinct and exactly one of us is conscious. Because I am certain that I am conscious, I am certain that George is not conscious. Hence my certainty that, if George and I are distinct and exactly one of us is conscious, George is not conscious.

Suppositions aside, Saddam and George are both justified in accepting the fact (i) that Saddam is conscious and the fact (ii) that George is conscious. (i) and (ii) are jointly inconsistent with the false supposition that George and Saddam are distinct and only one of the two is conscious. So, relative to the supposition, (i) or (ii) must be rejected. Whereas Saddam is justified in accepting (i) and rejecting (ii), George is justified in accepting (ii) and rejecting (i). So, the first reason it is implausible that Saddam could inform George of the idea expressed by ($\sim G$) is that Saddam's grounds for accepting this idea are not good grounds for George to accept the idea – despite the fact that George has knowledge of these grounds and has knowledge of the fact that these grounds are good grounds *for Saddam* to accept the idea. The second reason it is implausible that Saddam could inform George of the idea expressed by ($\sim G$) is that it is implausible, in the first place, that the idea expressed by ($\sim G$) is a *piece of information* – something that is objectively true. Plausibly, Saddam is justified in accepting the idea (with certainty); George is justified in rejecting the idea (with certainty); and neither is objectively correct or incorrect in his judgement – for, plausibly, the idea is not itself objectively true or false.

For similar reasons, it is implausible that Zack could inform Pete of the idea expressed by (w). First of all, Zack's grounds for accepting this idea are not good grounds for Pete to accept it, despite the fact that Zack and Pete have common knowledge of these grounds. Zack and Pete both know (a) – that Pete knows who has the stronger hand – and (b) – that Pete is determined to quit on knowing that he does not have the stronger hand. However, only Zack is justified in continuing to accept both of these facts relative to the supposition that Pete plays. Only Zack is justified in reasoning as follows:

Suppose that Pete plays. Because I signalled the content of Mr. Stone's hand to Pete, I'm confident (a) that Pete knows who has the stronger hand. Because Pete plays to win, I'm confident (b) that Pete is determined to quit on knowing that Mr. Stone has the stronger hand. Hence my confidence that, if Pete plays, he will win.

Suppositions aside, both Zack and Pete are justifiably confident in (a), (b), and (c) – that Mr. Stone has the stronger hand. These three facts are jointly inconsistent with the false supposition that Pete plays. So, relative to this supposition, at least one of the three must be rejected. Whereas Zack is justified in accepting (a) and (b) and rejecting (c), Pete is justified in accepting (a) and (c) and rejecting (b). So, the first reason it is implausible that Zack could inform Pete of the idea expressed by (w) is that Zack's grounds for accepting this idea are not good grounds for Pete to accept it – despite the fact that Pete has knowledge of these grounds and of the fact that these grounds are good grounds *for Zack* to accept the idea. The second reason it is implausible that Zack could inform Pete of the idea expressed by (w) is that it is implausible, in the first place, that the idea expressed by (w) is a *piece of information* – something that is objectively true. Plausibly, Zack is justified in accepting the idea; Pete is justified in rejecting it; and neither is objectively correct or incorrect in his judgement – for, plausibly, the idea is not itself objectively true or false. This was the very point of Gibbard's original example.

The second count on which DeRose's explanation of Datum 1 is implausible is its commitment to the idea that a difference in deliberational statuses between Gibbard and Pete plays a primary role in explaining the datum. Zack knows that his grounds for accepting the idea expressed by (w) are good grounds for Gibbard, but not for Pete, to accept the idea (at least before Jack tells Gibbard his grounds for uttering (~w)). For this reason, it would not be misleading for Zack to utter (w) to Gibbard, but it would be misleading for him to utter (w) to Pete. For it is common knowledge between Pete and Zack that Pete has seen his own cards and that Zack's grounds for accepting the idea expressed by (w) are therefore not good grounds for Pete to accept it. If, on his way out of the room, Zack were to whisper (w) into Pete's ear without explanation, he could reasonably expect to mislead Pete into reasoning as follows:

That's strange. I can see why *Zack* would accept that, if I play, I will win, but the grounds I know him to have for accepting this are not good grounds *for me* to accept it, and Zack knows this. So Zack must have *different* grounds for accepting the idea, grounds that Zack takes to be good grounds *for me* to accept it. For instance, perhaps he earlier signalled bad information to me in order to mislead someone else's henchman into thinking that Mr. Stone holds a strong hand when in fact Mr. Stone holds a weak hand. Whatever his grounds are for telling me now that, if I play, I will win, they must be good grounds *for me* to accept the idea. Otherwise he would not be telling me this. So I shall reject what he signalled to me earlier and instead accept that, if I play, I will win.

Thus, the explanation of Datum 1 has primarily to do with a difference in epistemic, and not deliberational, statuses between Gibbard and Pete: Zack knows that his own grounds for accepting the idea expressed by (w) are good grounds for Gibbard, but not for Pete, to accept the idea; given Zack's knowledge of this epistemic difference, it would be misleading – and thus

inappropriate – for Zack to utter (w) without explanation to Pete, but not to Gibbard.

To be sure, this inappropriateness may be compounded by the fact that Pete is deliberating on whether to play, and so the practical consequences of misleading him in this way are potentially severe. But the fact that Pete is so deliberating is secondary to the principal explanation. To see that this is so, imagine that, upon receiving Zack's signal, Pete sends back a signal letting Zack know that Pete has made a firm decision on whether to play. As the room is cleared, Zack knows that Pete is no longer deliberating on whether to play. Would it now be appropriate for Zack to whisper (w) into Pete's ear? Of course not. So the difference in deliberational statuses between Gibbard and Pete is at most secondary to the explanation of why it is appropriate for Zack to utter (w) without explanation to Gibbard, but not to Pete.

A third count on which DeRose's explanation of Datum 1 is implausible is its commitment to the idea that what explains Datum 1 is a semantic fact about the assertability conditions of if-will sentences. On this idea, Zack would exhibit incompetence with the meaning of (w) if he were to utter (w) to Pete. But, as a referee for DeRose's paper points out,¹⁶ this seems wrong. If, on his way out of the room, Zack were to whisper (w) into Pete's ear, he might reasonably be faulted on *pragmatic* grounds for misleading Pete in the aforementioned way, but he need not be faulted on grounds of linguistic incompetence. The third count on which DeRose's explanation of Datum 1 is implausible, then, is its commitment to the idea that what explains Datum 1 is a semantic fact about the assertability conditions of if-will sentences, rather than a pragmatic fact.

I conclude that DeRose's explanation of Datum 1 is implausible on three counts.

I turn next to DeRose's explanation of Datum 2. By it, (w*) expresses the same piece of information as (w); Zack's knowledge of this information depends on backtracking reasoning; and it is part of the meaning of if-would sentences that, if a speaker's knowledge of the information expressed by an if-would sentence depends on backtracking reasoning, then the sentence is not assertable for that speaker, in any context, without explanation. This explanation is implausible on at least two counts.

First, it is implausible that Datum 2 is explained by a semantic fact about the assertability conditions of if-would sentences. We can agree with DeRose that Zack is justifiably confident of the idea expressed by (w*). For Zack may justify this confidence as follows:

I have good reason to think (a) that Pete knows who has the stronger hand; (b) that Pete is determined to quit on knowing that Mr. Stone has the stronger hand; and (c) that Mr. Stone has the stronger hand.
Because I am confident in (a) – (c), I am confident that Pete will not play. Of course, supposing that Pete *were* to play, I am confident that at

¹⁶ DeRose says, 'A referee asks whether this aversion of ours to deliberately useless conditionals might just be a matter of Gricean pragmatics and an instance of our general attempt to say what is helpful to say' (fn. 27). The answer to the referee's question seems to me to be *yes*. DeRose responds to the question by addressing a rather different question, namely, whether differences in our use of if-will and if-would sentences might be explained by Gricean pragmatics *without* positing a difference in the content or meaning of the two sentences.

least one of (a) – (c) would be false. I doubt (a) would be false, for I am a loyal soldier: my determination to signal Mr. Stone’s hand to Pete is stable. I doubt (b) would be false, for Pete’s determination to quit on knowing that Mr. Stone has the stronger hand is also stable. Most likely, (c) would be false, for who has the stronger hand is a matter of pure chance. Most likely, then, if Pete were to play, Pete would have the stronger hand. Hence my confidence that, if Pete were to play, he would win.

If Zack were to utter (w*) in the context of making the preceding reasoning explicit, there is no reason to expect that his utterance would mislead his audience. But without explanation, there is a significant chance that it would mislead his audience into reasoning as follows:

When Zack says that, if Pete were to play, he would win, Zack probably means that, if Pete were to play *in the circumstances actually surrounding Pete*, Pete would win. After all, it is common practice, when considering the possible outcomes of a given agent’s options, to fix the circumstances external to the agent’s mind and ask what would happen if the agent were to choose a given option *in those circumstances*. For this reason, talk of what would happen if a given agent were to choose a given option is typically abbreviated talk of what would happen if the agent were to choose the option *in the circumstances actually surrounding the agent*. So what Zack probably means is that, if Pete were to play *in the circumstances actually surrounding Pete*, Pete would win. Of course, only if these circumstances include the fact that Pete has the stronger hand could Pete win by playing in them. So Zack must have reason to believe that Pete has the stronger hand.

Because Zack knows that it is common practice, when considering the possible outcomes of an agent’s options, to fix the circumstances external to the agent’s mind and ask what would happen if the agent were to choose a given option *in those circumstances*, Zack is positioned to know that, without explanation, his audience is likely to be misled by an utterance of (w*) into the preceding reasoning. For this reason, it would typically be inappropriate for Zack to utter (w*) without explanation. So, the first count on which DeRose’s explanation of Datum 2 is implausible is its commitment to the idea that what explains Datum 2 is a semantic fact about the assertability conditions of if-would sentences, rather than a pragmatic fact about the conditions under which a speaker could utter (w*) without misleading his audience.

The second count on which DeRose’s explanation is implausible is its commitment to the idea that if-would ideas are identical to their if-will counterparts. Zack’s situation with respect to the ideas expressed by (w*) and (w) resembles your situation in Scenarios 1 and 2 with respect to **Will** and **Would**: he is justified in maintaining high degrees of confidence in both ideas. But with respect to a pair of related ideas, Zack’s situation more closely resembles your situation in Scenario 3. Zack is justifiably confident *that, if Pete plays in the circumstances actually surrounding him, he will win*; but Zack is not justifiably confident *that, if Pete were to play in the circumstances actually*

surrounding him, he would win. Zack's evidence regarding the causal relation between *Pete's playing in the circumstances actually surrounding him* and *Pete's winning* is sensitive to the supposition that Pete plays. Suppositions aside, Zack has strong evidence that Pete's playing in the circumstances actually surrounding Pete would cause Pete to not win, for Zack has strong evidence that these circumstances include the fact that Mr. Stone has the stronger hand. Zack therefore maintains a *low* degree of confidence that, if Pete were to play in these circumstances, he would win.¹⁷ Supposing however that Pete plays, that is good evidence for Zack that Pete has the stronger hand. Indeed, it is such good evidence that it undermines Zack's original evidence that Mr. Stone has the stronger hand to the point where, on the supposition that Pete plays in these circumstances, Zack is justifiably confident that Pete has the stronger hand and thus that Pete's playing in these circumstances would cause Pete to *win*. Hence Zack maintains a *high* degree of confidence that, if Pete plays in these circumstances, he will win. Of course, if Zack's degrees of confidence in an if-would idea and its if-will counterpart are not identical, then the two ideas are not identical. The second count on which DeRose's explanation of Datum 2 is implausible, then, is its commitment to the idea that if-would ideas are identical to their if-will counterparts.

I conclude that DeRose's explanation of Data 1 and 2 is implausible on five counts, and that there is a rival explanation that is both plausible and consistent with our opening verdict that if-will ideas are distinct from their if-would counterparts. I therefore conclude that DeRose's third reason for his thesis – that it is supported by a version of Gibbard's Sly Pete story – is not a good one.¹⁸

I close this section by concluding that none of DeRose's three reasons for his thesis is a good one.

¹⁷ For similar reasons, Pete is confident in the idea expressed by (w), but not of the idea that, if he were to play *in the circumstances actually surrounding him*, he would win. Clearly, it is his low degree of confidence in the latter idea, and not his high degree of confidence in the former idea, that should guide him in deciding whether to play. DeRose might suggest that this is due to the fact that Pete's confidence in the former idea is grounded in backtracking reasoning. DeRose says, '...backtracking must somehow be excluded, whether you think it is counterfactuals or indicatives that are the conditionals of deliberation' (p. 31). But Pete's low degree of confidence in the latter idea is also grounded in backtracking reasoning: suppositions aside, Pete is confident (b) that he is determined to quit on knowing that Mr. Stone has the stronger hand; supposing that he were to play in the circumstances actually surrounding him, Pete is confident that (b) would be false. So, it is at most a special kind of backtracking that is bad for deliberation, namely, that with respect to how things external to the mind of the agent would be, on the given choice. And perhaps the only reason that *this* backtracking is bad for deliberation is that it is indicative of the fact that one is not deliberating on the appropriate if-would idea in the first place. For the purposes of deliberation, one should consider what the outcome would be if one were to choose a given option *in the circumstances actually surrounding one*. So, deliberating on the appropriate if-would idea in the first place is one way to 'exclude backtracking'.

¹⁸ DeRose considers three more examples in passing: Newcomb's boxes, Nozick's seminary (1969), and Edgington's prisoners (1995). In each case, DeRose claims that his thesis explains the relevant data. But the resulting explanations are implausible, for many of the same reasons that his explanation of Gibbard's example is implausible. Most importantly, in each case, his thesis fails to explain how a thinker has different degrees of confidence in a given if-will idea and its if-would counterpart.

3. Five More Reasons to Reject DeRose's Thesis

In our opening discussion, we found our first reason to reject DeRose's thesis: in Scenario 3, you have different degrees of confidence in **Will** and **Would** – ideas that are identical on DeRose's thesis. In this section, I offer five more reasons to reject DeRose's thesis.

Each of the first three reasons corresponds to a characteristic respect in which if-was ideas differ from if-would-have ideas – more generally, in which indicative conditional ideas differ from subjunctive conditional ideas. For each of the three respects, I show that if-will ideas differ in that respect from if-would ideas. I thereby give three reasons to reject both parts of DeRose's thesis, for the first part entails that if-will ideas are identical to their if-would counterparts, and the second part entails that all future conditional ideas are indicative. The three respects on which I focus are members of a larger family of respects in which indicative conditionals differ from subjunctives, and in which if-will ideas can be shown to differ from if-would ideas.¹⁹ It is only for the sake of brevity that I focus on a limited number of respects here.

The first respect in which if-was ideas differ from if-would-have ideas concerns our ability to evaluate a conditional idea in the context of supposing that the possibility corresponding to its antecedent is not realized. In the case of if-was ideas, it seems difficult, if not impossible, to carry out such an evaluation. By contrast, in the case of if-would-have ideas, the evaluation is relatively unproblematic. For illustration, consider the past-tense correlates of **Will** and **Would**:

Was that, if you blew on the candles, you looked like a fool

Would Have that, if you had blown on the candles, you would have looked like a fool.

Suppose that the possibility corresponding to the antecedents of **Was** and **Would Have** – namely, your blowing – was not realized. In other words, suppose that you did not blow on the candles. In this context, it seems difficult, if not impossible, to evaluate **Was**. After all, the natural starting place is to suppose that you blew on the candles. And this seems difficult, if not impossible, in the context of supposing that you did not blow on the candles. By contrast, the task of evaluating **Would Have** is relatively unproblematic. The natural starting place is to suppose that you *had* blown on the candles. We then consider the most likely metaphysical implications of this counterfactual difference, including whether you would have looked like a fool.

If-will ideas differ in the same respect from if-would ideas. Suppose for illustration that the possibility corresponding to the antecedents of **Will** and **Would** – namely, your blowing – will not be realized: suppose that you will not blow on the candles. In this context, it seems difficult, if not impossible, to evaluate **Will**. For the natural starting place is to suppose that you will blow on the candles, and this seems difficult, if not impossible, in the context of supposing that you will not blow on them. By contrast, the task of evaluating

¹⁹ For a sample of this family, see Gibbard 1981; Edgington 1995; and my 2006, 2009, 2010.

Would is relatively unproblematic. The natural starting place is to suppose that you *were* to blow on the candles. We then consider the most likely metaphysical implications of this counterfactual difference, including whether you would look like a fool.

So, in the first characteristic respect in which if-was ideas differ from if-would-have ideas, if-will ideas also differ from if-would ideas. Thus, we have our first additional reason to reject DeRose's thesis.

The second respect in which if-was ideas differ from if-would-have ideas concerns our evaluation of a conditional idea in the context of supposing that the possibility corresponding to its consequent is realized. In this context, we are always rationally required to be certain of an if-was idea, but we are not always rationally required to be certain of an if-would-have idea. For illustration, suppose that the possibility corresponding to the consequents of **Was** and **Would Have** – namely, your looking like a fool – was realized: suppose that you looked like a fool. In this context, we are rationally required to be certain that, if you blew on the candles, you looked like a fool. For, supposing that you looked like a fool, and supposing that you blew on the candles, we are rationally required to be certain that you looked like a fool. It is easy to see why this result generalizes to all if-was ideas.²⁰ By contrast, we need not be rationally required to be certain that, if you had blown on the candles, you would have looked like a fool. We might, for instance, have reason to believe that you did not in fact blow on the candles, in which case we have no reason to be certain that, if you *had* blown on them, things would have been the same as actual with respect to whether you looked like a fool.

If-will ideas differ in the same respect from if-would ideas. Suppose for illustration that the possibility corresponding to the consequents of **Will** and **Would** will be realized; that is, suppose that you will look like a fool. In this context, we are rationally required to be certain that, if you blow on the candles, you will look like a fool. And it is easy to see why this result generalizes to all if-will ideas. But we are not rationally required to be certain that, if you were to blow on the candles, you would look like a fool. We might, for instance, have reason to believe that you will not in fact blow on the candles, in which case we have no reason to be certain that, if you *were* to blow on them, things would be the same as actual with respect to whether you will look like a fool.

So, in the second characteristic respect in which if-was ideas differ from if-would-have ideas, if-will ideas also differ from if-would ideas. Thus, we have our second additional reason to reject DeRose's thesis.

The third respect in which if-was ideas differ from if-would-have ideas concerns what our degrees of confidence in a conditional idea purport to measure. In the case of if-was ideas, they purport to measure an *evidential* dependence between the possibility corresponding to the antecedent and the possibility corresponding to the consequent. In the case of if-would-have ideas, they purport to measure a *metaphysical* dependence – typically, causal –

²⁰ There are candidate exceptions to this generalization, for instance, if-was ideas whose antecedents are incompatible with their consequents. But such candidates are mirrored by equally good candidates in the case of if-will ideas. Hence they pose no problem for my present aim, which is to show that if-was ideas differ from if-would-have ideas in a certain respect in which if-will ideas also differ from if-would ideas.

between the two possibilities.²¹ Call this the *evidential-metaphysical difference* between if-was and if-would-have ideas. Because evidence of a metaphysical dependence between two possibilities often constitutes an evidential dependence between the two, evidence for an if-would-have idea often constitutes evidence for its if-was counterpart. For this reason, thinkers are often justified in placing the same degree of confidence in an if-would-have idea as in its if-was counterpart idea. When a thinker is justified in placing different degrees of confidence in the two ideas, that is often because the thinker's evidence regarding the metaphysical dependence between the relevant two possibilities is sensitive to the supposition of the antecedent of the if-was idea. The story is familiar: we have seen that the evidential-metaphysical difference distinguishes if-will ideas from if-would ideas. To see that it distinguishes if-was ideas from if-would-have ideas, consider the following past-tense versions of our opening example.

First consider Scenario 4, a past-tense version of Scenarios 1 and 2. Here, your friend Courtney leaves your party just before you are presented with your cake. On her way out, Courtney notices an empty box labelled 'Trick Candles'. Courtney wakes up the next morning not knowing whether you blew on the candles. Her memory of the empty box constitutes evidence for her of **Past Cause**:

Past Cause that your blowing on the candles would have caused you to look like a fool.

Based on this evidence, Courtney is confident that, if you had blown on the candles, you would have looked like a fool; in other words, she is confident in **Would Have**. Now, in the context of supposing that you blew on the candles, Courtney has the very same evidence for **Past Cause** that she has outside of the supposition. This evidence thus constitutes evidence for Courtney – within the context of supposing that you blew – that your blowing *did* cause you to look like a fool, and thus that you *did* look like a fool. Based on this evidence, Courtney is confident that, if you blew on the candles, you looked like a fool; in other words, she is confident in **Was**. In Scenario 4, Courtney's evidence for **Would Have** constitutes her evidence for **Was**. This is why she places the same degree of confidence in the two ideas.

Now consider Scenario 5, a past-tense version of Scenario 3. Without noticing any physical evidence of trick candles, Courtney leaves the party early. The next morning, she wakes up not knowing whether you blew on the candles. What she does know is (i) that your candles were prepared by a prankster who uses trick candles if, and only if, she predicts that the recipient of her candles will blow on them; (ii) that this prankster has a strong record of predicting who will blow on their candles; and (iii) that the prankster has used ordinary candles on your previous 25 birthdays. (iii) constitutes some evidence for Courtney that the candles were *not* trick. This is evidence for Courtney against **Past Cause**. Based on this evidence, Courtney maintains a

²¹ A high degree of confidence in an if-would-have idea need not purport to measure a *positive* metaphysical dependence between the two possibilities. Our confidence that, if the Big Bang had never happened, two plus two would still have equalled four reflects our confidence that two plus two's equalling four is *not* metaphysically dependent on the Big Bang.

relatively low degree of confidence in **Would Have**. However, on the supposition that you blew on the candles, Courtney has a new source of evidence regarding **Past Cause**: Courtney now has strong evidence that the prankster used trick candles, and this evidence undermines her original evidence against **Past Cause** to the point that she now has strong evidence in favor of **Past Cause**. She thus has strong evidence – in the context of supposing that you blew – that your blowing *did* cause you to look like a fool, and that you therefore *did* look like a fool. Based on this evidence, Courtney maintains a high degree of confidence in **Was**. Hence, in Scenario 5, the symmetry is broken: Courtney’s evidence regarding **Past Cause** is sensitive to the supposition that you blew on the candles; as a result, Courtney’s evidence regarding **Would Have** is significantly different from her evidence regarding **Was**.

Scenarios 4 and 5 illustrate how the evidential-metaphysical difference distinguishes if-was ideas from their if-would-have counterparts. Scenarios 1-3 illustrate how it distinguishes if-will ideas from their if-would counterparts. So, in the third characteristic respect in which if-was ideas differ from if-would-have ideas, if-will ideas also differ from if-would ideas. Thus, we have our third additional reason to reject DeRose’s thesis.

Before turning to the fourth additional reason, it is worth pausing to address the following question: why do if-will ideas seem to resemble their counterpart if-would ideas more often than if-was ideas seem to resemble their counterpart if-would-have ideas? There are a number of reasons for this asymmetry. Perhaps the most significant centers on an asymmetry between our knowledge of the past and our knowledge of the future. When we consider what *would have* happened, if a given possibility *had* obtained, we often take ourselves to know that the given possibility *did not* in fact obtain. On the supposition that the given possibility *did* obtain, we are then forced to revise our view of the historical facts, which can often force us to revise our view of the metaphysical connections between the possibility corresponding to the antecedent and the possibility corresponding to the consequent. For this reason, our evidence regarding if-would-have ideas is often different from our evidence regarding their counterpart if-was ideas. By contrast, when we consider what *would* happen, if a given possibility *were* to obtain, we often take ourselves to be ignorant of whether the given possibility will obtain. This is especially so in a context of deliberation, where it only makes sense to deliberate on whether to bring about a given option if, at the start of the deliberation, we take ourselves to be ignorant of whether the option will obtain. In such a context of ignorance, one is less likely to be forced, on the supposition that the option will obtain, to revise one’s views about the metaphysical connections between this option and its possible outcomes. For this reason, our evidence regarding if-would ideas is often the same as our evidence regarding their counterpart if-will ideas. And so we have the beginning of an explanation as to why if-will ideas seem to resemble their counterpart if-would ideas more often than if-was ideas seem to resemble their counterpart if-would-have ideas.

We also have the beginning of an explanation as to why DeRose might have been drawn to his thesis in the first place. DeRose seems to have been led to theorize about future conditionals by his interest in deliberation, which seems to have led him to focus exclusively on contexts of deliberation. In such

contexts, if-would ideas are likely to resemble their if-will counterparts, for the preceding reasons. Perhaps DeRose was led by this resemblance to identify if-would ideas with their if-will counterparts. Faced with apparent counterexamples to this identification, and still focusing on contexts of deliberation, DeRose concluded that all such apparent counterexamples can be explained away by the idea that if-would sentences are the result of 'souping-up' if-will sentences for the purpose of marking the ideas they express as deliberationally useful. Thus, DeRose's focus on contexts of deliberation seems to have led him to conclude that even apparent counterexamples to his thesis have something to do with deliberation. This brings us to our last two reasons to reject his thesis.

The fourth additional reason to reject DeRose's thesis is that there are counterexamples to it that have nothing to do with deliberation. For illustration, consider (1):

- (1) 'If the total energy of the universe changes tomorrow, the laws of nature will be different from what they will actually be'.

No rational thinker is confident of the if-will idea expressed by (1). But now consider its if-would counterpart, expressed by (1*):

- (1*) 'If the total energy of the universe were to change tomorrow, the laws of nature would be different from what they will actually be'.

Physicists have excellent reason to be confident that the total energy of the universe will not change tomorrow, for they have excellent reason to be confident that it is a law of nature that the total energy is conserved over time. On the supposition that the total energy *were* to change tomorrow, the most likely metaphysical implications of this counterfactual difference include a counterfactual difference in the laws of nature. Hence, physicists have good reason to be confident of the idea expressed by (1*). At the very least, they have good reason to be *more* confident of this idea than of the idea expressed by (1).

This counterexample to DeRose's thesis has nothing to do with deliberation. It cannot be explained away by DeRose's proposal that if-would sentences are the result of 'souping-up' if-will sentences for the purpose of marking the ideas they express as deliberationally useful. For this proposal entails that the assertability conditions of if-will sentences are more permissive than those of if-would sentences. But in the present example, it is the if-would sentence, and not its if-will counterpart, that is a candidate for being assertable. So the strategy of explaining away counterexamples by employing this alleged distinction between the assertability conditions of if-will and if-would sentences cannot even get off the ground here. The fourth additional reason to reject DeRose's thesis, then, is that there are counterexamples to it that have nothing to do with deliberation and which are not even initial candidates for being explained away by DeRose's proposal that if-would sentences are the result of 'souping-up' if-will sentences for the purpose of marking the ideas they express as deliberationally useful.

The fifth additional reason to reject DeRose's thesis is that it places an implausible constraint on what is required to understand if-would sentences. On DeRose's thesis, part of the meaning of if-would sentences is that they should not be uttered on deliberately useless grounds. But surely one could fully understand an if-would sentence, such as (I*), without having any disposition to find it inappropriate, in any sense, to utter the sentence, on deliberately useless grounds. To be sure, such a speaker might not grasp certain general rules of conversation and might even lack a sense of common decency. But such a speaker nevertheless seems possible. For, understanding (I*) seems to have nothing to do with understanding anything about deliberation. Indeed, if it is possible to understand any expression of English without having the concept of deliberation, then it seems possible to understand the if-would construction without possessing the concept of deliberation. If-would ideas can be useful tools for deliberating, but understanding why and how they are so useful does not require adding the concept of deliberation to the meanings of the sentences that express them. The fifth and final additional reason to reject DeRose's thesis is that, by adding the concept of deliberation to the meaning if-would sentences, the thesis places an implausible constraint on what is required to understand the sentences.

I conclude that, in addition to our opening reason for rejecting DeRose's thesis, there are five more reasons to reject it. Given my argument that none of DeRose's three main reasons in favor of his thesis is a good one, I conclude that DeRose's thesis is false.

In the course of our discussion, a credible alternative picture of the relationship between if-will and if-would ideas has emerged. On this picture, causal decision theorists may continue to utilize the distinction between the two types of ideas to express their view that the relevant connection between an option and its possible outcomes is causal, not evidential.²²

Department of Philosophy
University of Colorado at Boulder
Campus Box 232, Hellems 167
Boulder, Colorado
80309-232
USA
David.Barnett@Colorado.edu

REFERENCES

- Barnett, David 2006: 'Zif is If'. *Mind*, 115, pp. 519-566.
-----2009: 'The Myth of the Categorical Counterfactual'. *Philosophical Studies*, 144, pp. 281-96.
-----2010: 'Zif Would Have Been If: A Suppositional View of Counterfactuals'. *Noûs*, 44, pp. 269-304.
DeRose, Keith 2010: 'The Conditionals of Deliberation'. *Mind*, 119, pp. 1-42.
Edgington, Dorothy 1995: 'On Conditionals'. *Mind*, 104, pp. 235-329.

²² For helpful comments I am grateful to Thomas Baldwin, Keith DeRose, and Paul Weirich.

- Gibbard, Allan 1981: 'Two Recent Theories of Conditionals'. In Harper, Stalnaker, and Pearce 1981, pp. 211-47.
- Gibbard, Allan and William L. Harper 1978: 'Counterfactuals and Two Kinds of Expected Utility'. In Harper, Stalnaker, and Pearce 1981, pp. 153-90. Originally published in Hooker, Leach, and McClennen 1978.
- Harper, W. L., R. Stalnaker, and G. Pearce (eds) 1981: *Ifs: Conditionals, Belief, Decision, Chance, and Time*. Dordrecht: D. Reidel.
- Hooker, C. A., J. J. Leach, and E. F. McClennen (eds) 1978: *Foundations and Application of Decision Theory*, Vol. I. Dordrecht: D. Reidel.
- Jackson, Frank 1979: 'On Assertion and Indicative Conditionals'. *The Philosophical Review*, 88, pp. 565-589.
- 1987: *Conditionals*. Oxford: Basil Blackwell.
- Joyce, James 1999: *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
- Lewis, David 1981: 'Causal Decision Theory'. *Australasian Journal of Philosophy*, 59, pp. 5-30.
- Nozick, Robert 1969: 'Newcomb's Problem and Two Principles of Choice'. In Rescher 1969, pp. 114-46.
- Rescher, N. (ed.) 1969: *Essays in Honor of Carl G. Hempel*. Dordrecht: D. Reidel.
- Skyrms, Brian 1982: 'Causal Decision Theory'. *Journal of Philosophy*, 79, pp. 695-711.
- Weirich, Paul 2001: *Decision Space*. Cambridge: Cambridge University Press.