NOÛS 33:3 (1999) 405–420

Conditional Assertions and "Biscuit" Conditionals

KEITH DEROSE
Yale University

RICHARD E. GRANDY
Rice University

1. "Biscuit" Conditionals

(1) There are biscuits on the sideboard if you want some
(2) If you're interested, there's a good documentary on PBS tonight

and

(3) Oswald shot Kennedy, if that's what you're asking me,

as they'd typically be used, are examples of what are often called "biscuit" conditionals, after J.L. Austin's example ((1), above; see Austin, 1970, p. 212). A mark of such a conditional is that, after it has been uttered, it can only be a certain kind of joke to ask what is the case if the antecedent is false—"And where are the biscuits if I *don't* want any?", "And what's on PBS if I'm *not* interested?", "And who shot Kennedy if that's *not* what I'm asking?". With normal indicative conditionals like,

(4) There are biscuits on the sideboard if Bill hasn't moved them
(5) If the TV Guide is accurate, there's a good documentary on PBS tonight

and

(6) Oswald shot Kennedy, if there hasn't been an enormous conspiracy,

though the answer may well be "I don't know," it seems to make good sense to seriously ask what is the case if the antecedent is false. An intuitive first shot at characterizing the difference is that, with the normal conditionals, it's the truth or believability of the consequent that is somehow being made contingent on the antecedent, while with the biscuit conditionals, it's the conversational relevance of the consequent that seems to be contingent on the antecedent. The joke, then, seems to consist in the joker's pretending to misunderstand the likes of (1)–(3) as a normal indicative conditional, where it's the truth or believability of the consequent that's being said to be contingent on the truth of the antecedent, so that it's a real question what is the case or what should be thought if the antecedent is instead false, while, in fact, it's only the *relevance* of the consequent that a utterer of (1)–(3) would typically intend to be making contingent on the antecedent.

We should emphasize that "typically". (1)–(3) *can* express normal conditionals, and there are unusual contexts in which they would naturally interpreted in such a way (which helps to explain why the joker can pretend to misunderstand them as such). If your local PBS affiliate is known to be monitoring your interests and adjusting their broadcasting schedule accordingly on a daily basis, then (2) may well express a normal conditional. And in a weird evidential situation in which your asking who killed Kennedy would serve as excellent evidence to someone that it was Oswald who did the deed, that person might use (3) to express a normal conditional while talking to you. It's only in the more ordinary situations in which you'd imagine them typically being used that (1)–(3) express abnormal, biscuit conditionals.

Another mark of biscuit conditionals is that they are contraposition-resistant. Now, at least some of our very strongest intuitions about the validity/invalidity of arguments involving indicative conditionals are wrong.[1] Still, we can use how inferences involving conditionals intuitively strike us as a classifying device. With normal indicative conditionals like (4)–(6), though it's not one of the most powerful intuitions we have, it certainly seems reasonable to suppose that, for instance, (6) implies

(6C) There has been an enormous conspiracy if Oswald didn't shoot Kennedy,

By contrast, contraposition produces intuitive clunks when it's applied to biscuit conditionals; it seems crazy to suppose that

(3C) You're not asking me about that if Oswald didn't shoot Kennedy,

follows from (3), as (3) would typically be used. Note that the same unusual circumstances in which it would not be a joke to respond to (1)–(3) by asking what is the case if their antecedents are false (e.g., where your asking about who shot Kennedy would serve as excellent evidence that it was Oswald) are also circumstances in which the likes of (1)–(3) can at least seem with some plausibility to support contraposition.

(1)–(3) certainly aren't "subjunctive" or "counterfactual" conditionals; if their "biscuit" uses are to be handled at all by any leading theory of conditionals, it will have to be by an account of "indicative" conditionals. But the leading theories of indicative conditionals, we believe, are not up to the task.

But should that be a mark against those theories? Should they even try to account for these "abnormal" conditionals? A defender of one of the leading theories might plausibly argue that they are trying to account for normal indicative conditionals, and since the abnormal, biscuit conditionals are quite different from the normal ones, any attempt to accommodate the latter in their theory would be a mistake. Better to cordon off the abnormal cases for separate treatment. Though we won't take the space to argue it here, we believe that that's the best the leading theories can do by way of "handling" biscuit conditionals.

In reply, it is true that biscuit conditionals are different from normal indicative conditionals, so it would indeed be a mistake for a theory to blur the differences between them. But the best case scenario, of course, is for there to be a unified theory which can account for both kinds of conditionals, while also respecting, and even explaining, the significant differences between them.

Fortunately, the general theory of indicative conditionals that we endorse, and would endorse even independently of its ability to handle biscuit conditionals, does just that.

In section 2, then, we present the general theory we endorse. In section 3, we apply it to the issue of biscuit conditionals. Finally, in section 4, we reply to an objection to the general theory that utilizes a biscuit conditional as a counter-example.

2. The Theory: Indicative Conditionals As Devices of Conditional Assertion

Many these days are tempted by the thought that indicative conditionals, which we'll here schematize as "A→C", don't assert anything that can be true or false. On the view we'll here promote, such thoughts are about half-right. That is, about half the time, according to us, "assertions" of "A→C" don't assert anything that is or can be true or false. But also half-wrong: The other half of the time a truth-evaluable assertion is made by means of 'A→C'. What is then asserted? C. That is, whenever stating 'A→C' does result in a truth-evaluable assertion, that assertion is true iff C is true. When does 'A→C' result in an assertion of C? When A is true. What, then, is asserted where A is false? Nothing. 'A→C' is used to *conditionally assert* that C.

On the theory being developed, then, conditionals are radically unlike conjunctions and disjunctions in semantic function. Conjunctions and disjunctions take two propositions and form a new proposition from them. But on our theory, no new proposition is formed by an indicative conditional; rather, such a conditional takes two propositions and produces a conditional assertion of the second proposition, where the first proposition expresses the condition under which the second proposition is asserted.

While we advocate this view of indicative conditionals, and so it is appropriate for us, in that sense, to call it "our" theory, as we will for convenience, we certainly can't claim to have come up with this theory. Indeed, as early as 1950, Quine mentioned it, rather favorably, in *Methods of Logic*.[2] In 1970, Nuel Belnap, though not himself endorsing it, pursued the view, writing that "the idea has not been taken with the high seriousness it might deserve."[3] Now, almost a half-century after Quine's remark and over a quarter of a century after Belnap's remark, it might still be true that the view is not considered one of the leading theories of indicative conditionals, but it has a gained a recent important advocate in Dorothy Edgington,[4] and it has something of a history of prominent backers and near-backers.[5] Perhaps it now *is* being taken seriously.

Why has the view been largely neglected? Well, there certainly are problems with it. We believe these are mostly problems with understanding the view in the first place—understanding just what conditional assertions are, how they operate, and how they interact with the common machinery of current philosophy of language—machinery that was not set up with such a speech act in view. Given that some of this machinery doesn't apply smoothly to the case of this speech act we're positing, it can seem unclear just exactly what a conditional assertion is supposed to be, much less to see that indicative conditionals are devices for performing such a speech act. It can also seem problematic what kind of logic conditionals could have if they were devices of conditional assertion; what is going on when, as we're wont to say, we *believe* such a conditional; and how conditionals could embed themselves within other conditionals or within the scope of sentential connectives like 'or'.

But even on a quite preliminary understanding of what a conditional assertion would be, it becomes clear that conditional assertions, if there were such things, would have certain properties. We believe that these properties that conditional assertions would have match very well the properties that statements of indicative conditionals do have, and so constitute good reasons for thinking that indicative conditionals are devices of conditional assertion.

For a very good start on answering the apparent problems of logic, belief and embedding, we recommend (Edgington 1995), which, on the issue of logic, adapts the work of Ernest Adams to the view of conditionals as devices of conditional assertion.[6] Our focus here will be on the view's handling of biscuit conditionals. But, in the interest of convincing the reader that the view is at least worth taking seriously, we will in this section note a couple of features that, it seems, conditional assertions would have and which match up well with the properties that statements of indicative conditionals do have.

For instance, it seems that, as with other forms of assertion, some conditional assertions would be warranted and others not. The warrant for making a conditional assertion would not be warrant for a conditional proposition, but rather warrant for asserting one proposition should another proposition be true. When would it be appropriate to conditionally assert C (on the condition A)? Well, when is it appropriate to make a *simple* assertion—i.e. an assertion that does not result from a conditional assertion? The common understanding is that one should simply assert that P only when P's probability is very high—quite close to 1. Given that common account of simple assertion, then, given even a preliminary understanding of conditional assertions, one would expect that one would be in a position to conditionally assert that C (on condition A) when the *conditional* probability of C on A is very high—quite close to 1.[7]

And this constitutes evidence in favor of our theory. But in explicitly stating why this constitutes a recommendation, we already run into a problem of terminology which will force us to extend a key term used in current philosophy of language. It has been largely agreed in discussions of normal, non-biscuit indicative conditionals that their assertability goes by the conditional probability of their consequents on their antecedents. It's seen as a desideratum of a theory that it account for this fact. Our theory provides the needed explanation, but forces us to redescribe the fact being accounted for. We can't call 'A→C' assertable or unassertable, because we think there is no new proposition corresponding to the construction 'A→C' to be assertable or not. We need a broader term, so we introduce "statable" for this purpose. In the case of simple (non-conditional) assertions, to say that a sentence is statable is to say that it's appropriately assertable. But also, by stipulation, if the function of a sentence is to conditionally assert a proposition (on the condition of another proposition), then that sentence is statable when it is appropriate to make that conditional assertion. (Note that if we're wrong about the function of indicative conditionals, then this stipulated sense of "statable", as it's applied to indicative conditionals, just collapses into the old notion of "assertable", and if there are *no* sentences whose job it is to make conditional assertions, then "statable" generally collapses into "assertable".) We can now state the datum to be explained in a way that doesn't beg the question of the function of "A→C": Whether "A→C" generates a new proposition or whether its function is to conditionally assert that C, it is statable where the conditional probability of C on A is high.[8] One virtue of our view is that it predicts that result.

Very quickly, here's another virtue of our theory. The leading theories of indicative conditionals take seriously our intuitions about the truth or falsity of various uses of the relevant sentences—whether that takes the form of verifying those intuitions or of explaining them away. However, they seem fairly unconcerned about our intuitions that certain indicative conditionals are neither true nor false—intuitions that are often elicited where the antecedent is false.[9,10] If one takes these intuitions seriously, as one should, that, it seems, will favor either our view, or the view mentioned in a footnote of (Stalnaker 1975; see p. 137, fn. 2) and later taken up in (McDermott 1996), according to which "A→C" generates a proposition which has the same truth-value as the consequent if the antecedent is true, but which goes truth-valueless where the antecedent is false. But, as Stalnaker pointed out in that same footnote, this latter view comes to grief over the fact that we're often very happy to assert an indicative conditional where we think the antecedent is quite probably false, so long as we think that the conditional probability of the consequent on the antecedent is high. It seems strange that we would assert a proposition which has a very low probability of being true, even if it also has a low probability of being false. Our view avoids this problem, for on it, should the antecedent prove false, you haven't asserted a proposition which fails to be true, for you've asserted nothing.

Again in the spirit of getting our view on the table as something to be taken seriously, we should at least *mention* one other very important virtue we believe the theory has, even though this is not the place to develop this aspect of the theory. For it may strike the reader that even if the theory fits the biscuit conditionals, this is a slim reed with which to support a theory that declares conditionals to be radically unlike conjunctions and disjunctions, and radically unlike what generations of logicians and philosophers of language have thought.

As for going against the current grain, we're disappointed in the results of each of the choices on the usual menu of options, and think that trying a different direction is certainly in order.

As for positing a new speech act (conditional assertions) and making indicative conditionals radically unlike conjunctions and disjunctions, we think we can repay that cost. For, when we broaden our view beyond the realm of assertions, we believe our theory of conditional assertions will fit nicely into a general account of conditional speech acts in ways that theories of indicative conditionals that make them like conjunctions and disjunctions cannot, as Edgington points out.[11] Given our view of indicative conditionals, on which the making of an assertion is conditional on the truth of a sentence, one naturally wonders why one couldn't conditionally perform speech acts other than assertion by means of "if" sentences with indicative "if" clauses, but where the other clause, taken by itself, would execute the other speech act in question. And, when we look to natural language, we find constructions which seem to do just that. To take just one example, we can issue conditional warnings:

(7) If you go to New York, watch out for the taxi drivers.

That we can issue conditional warnings, questions, and commands in such a way is unsurprising given our view. And in support of our claim that conditionals are unlike conjunctions and disjunctions, note that neither of the latter can combine indicatives with warnings:

(8) ?You are going to New York and watch out for the taxi drivers.
(9) ??You are not going to New York or watch out for the taxi drivers.

Much more needs to be said, of course, about conditional speech acts, and we won't say it here. We hope the above, though, gives enough of a sense of promise to our view that the reader can now take it as a serious contender. Toward that end, we should mention that, in addition to Edgington and ourselves, who actually endorse the view, others certainly believe that the view is a serious contender. In addition to Belnap, who we've already mentioned, it's worth noting that Robert

Stalnaker took the view so seriously that he listed it as the most plausible contender to his own theory of indicative conditionals (Stalnaker 1975, p. 137, fn. 2).

3. Applying the Theory to Explain the Two Types of Conditionals

Continuing the fruitful project of trying to discern what properties conditional assertions would have, if there were such creatures, we do well to ask ourselves: What would conditional assertions be good for? Under what circumstances would it be handy to conditionally assert C (on the condition that A), as opposed to simply asserting that C or making no assertion or statement at all? In the previous section, we saw the makings of one answer to this question. Sometimes C might not be probable enough to assert, but the conditional probability of C on A might be high enough to warrant conditional assertion.

But would there be other reasons for making a conditional, rather than a simple, assertion of C? What if C itself is probable enough to warrant simple assertion, and the conditional probability of C on A is not appreciably higher than is the probability of C? Could there *still* be reasons for only conditionally asserting C? Yes. In addition to the potential problem of C's not being probable enough to be asserted, there is a very different reason why C might fail to be simply assertable: It might not be conversationally relevant.[12] Now, assuming C's probability is sufficient for assertion, if you know that C is also conversationally relevant, you would not bother to conditionally assert it: Simple assertion would be a better choice in such a case. And if you're sure that C is irrelevant, then stating nothing seems a good choice. But what if you don't know whether C is relevant? And what if you knew C would be relevant if a certain condition, say, A, were fulfilled? Then it would be handy to conditionally assert C—assert it on the condition that A. Such a conversational maneuver would allow you to make a warranted assertion should it be called for, while at the same time shielding you from the danger of committing the conversational misdeed of making an irrelevant assertion.

"But wait! Whatever device we might use to make a conditional assertion would have to specify both the proposition being conditionally asserted and the proposition on which the assertion is conditioned. It is thus bound to be significantly more long-winded than a simple assertion of C. Would we really go to all that trouble just to avoid the possibility of bothering our listeners with an irrelevant assertion? If your solution to the problem of potentially bothering us with an irrelevant assertion is to bother us even more with a long-winded conditional assertion, we prefer the disease to the cure!"

But there are dangers lurking about other than merely bothering a listener with irrelevancy—dangers of a Gricean sort. For when you assert an irrelevancy, you not only assert an irrelevant truth, but, much more problematically, you implicate a falsehood. Since there is a general conversational rule to the effect that you should not assert the irrelevant, when you assert that C, you conversationally implicate that C is relevant.[13] And, then, of course, if C isn't relevant, you've implicated a falsehood. That, by itself, is sufficient to explain why we wouldn't simply assert that C where we think C might be irrelevant, given our general reluctance to implicate falsehoods. (In current philosophy of language, that making an assertion under certain circumstances would generate a false implicature—or, from the potential asserter's point of view would likely generate such a false implicature—is generally considered a sufficient explanation for why we wouldn't make the assertion under those circumstances.[14])

Still, it might be illuminating to explore what's problematic about generating a false implicature in the special case where the false implicature is that the assertion is relevant. Of course, there's the potential problem of misleading your listener. If you falsely implicate that C is relevant, then your listener may begin to wonder whether you know something that she doesn't that would make C relevant, and where you don't know whether C is relevant, you have no such knowledge. And, besides the possibility of misleading one's listener about the relevance of C, there's another potential problem that could be produced by generating an implicature to the effect that C is relevant. Often, the relevance of C will turn crucially on your listener's interests and desires. Now, your listener is unlikely to be misled by you about what her interests and desires are. But, in such a situation, where you don't know whether your listener has the interests or desires to render C relevant, it can be presumptuous of you to implicate that C is relevant. You run the risk of indicating that the listener should be interested in C or should have the desires that would make C relevant, when in fact you think it's perfectly OK that your listener should be completely uninterested. Being conversationally uncool in this presumptuous way goes beyond the mere matter of bothering your listener with an irrelevancy.

For whatever reason one might want to avoid generating a false implicature to the effect that C is relevant—and we seem reluctant to implicate a falsehood regardless of what further problems doing so might cause—a conditional assertion of C seems to be just what the conversational doctor ordered. If you don't know whether C is relevant, but do know that it would be relevant under condition A, then being able to make a conditional assertion of C on the condition that A would be very helpful. By making your assertion of C conditional on A, you avoid generating an implicature to the effect that C is relevant, an implicature which might well be false, from your point of view. But if A is true, and C *is* relevant, then you will have asserted C. So, in conditionally asserting, you have to be long-winded, but, in return, you get the benefit of making the assertion of C should it prove relevant, without running the risk of falsely implicating that C is relevant should it prove to be irrelevant. Sounds like a good bargain to us.

Well, hopefully the cat's been out of the bag for quite a while now. Given that there are two very different reasons why one might refrain from simply asserting C—reasons involving C's probability on the one hand, and its conversational relevance on the other—our theory of indicative conditionals as devices of conditional assertion would lead one to suspect that there should be two very different uses of indicative conditionals: (1) cases where "A→C" is stated rather than "C" because it's only the conditional probability of C on A, and not the simple probability of C, that's high enough to warrant assertion[15] and (2) cases where the speaker isn't sufficiently sure that C is conversationally relevant, but does know that C is relevant if A is true. The "normal" indicative conditionals of section 1, then, will be generated in cases of type (1), while "biscuit" or "relevance" conditionals will be generated in type-(2) cases.

(1) There are biscuits on the sideboard if you want some,

then, in its typical use, would be a biscuit conditional, because one would typically state it rather than

(10) There are biscuits on the sideboard

in cases where it's the relevance, and not the probability, of (10) that's problematic. By contrast,

(4) There are biscuits on the sideboard if Bill hasn't moved them,

because it would likely be asserted where it's the probability of (10) that's problematic, is, in its typical use, a normal conditional. Because there can be two very different reasons why one might settle for a conditional, rather than a simple, assertion, our single theory of indicative conditionals as conditional assertions can handle these two very different uses of indicative conditionals. Thus, as promised, ours is a unified theory which accounts for both kinds of conditionals, while also respecting, and even explaining, the significant differences between them.

4. Implicatures and Belnap's Counter-Example

Still continuing our inquiry into what properties conditional assertions would have, we do well to ask what implicatures would be generated by making a conditional assertion.[16] Given the results of the previous section, and supposing it's obvious in context whether a conditional assertion is being made for reasons of relevance or for reasons of probability, we would expect two very different types of implicatures to be generated, depending on why a conditional assertion is being made.

If it's clear that C is being only conditionally asserted for reasons of probability, then we should expect that conditional assertion to generate the implicature that C is conversationally relevant. This because a non-conditional assertion of C would generate such an implicature, and we're supposing that it's clear that the reason you are settling for a conditional assertion has nothing to do with C's relevance but rather has to with C's probability. Your listener should then be able to calculate that, if you're to be trusted, C is conversationally relevant. Of course, here you generate no simple commitment to the truth of C;[17] after all, it was precisely in order to generate only a conditional commitment to the truth of C that you settled for a conditional assertion of C.

By contrast, where it's clear that C is being only conditionally asserted due to relevance concerns, that situation is reversed: You're only conditionally committed to C's relevance, but are simply committed to the truth of C. First, you here generate no simple commitment to the *relevance* of C. After all, it was precisely because of doubts about C's relevance that you resorted to a conditional assertion. But you do in this case generate a simple commitment to the *truth* of C, for you generate such a commitment by non-conditionally asserting that C, and we're supposing it's clear that the only reason you are not making such a simple assertion is that you're concerned over potential problems about C's conversational relevance. Thus, in cases dominated by relevance concerns, even where the proposition on which your assertion of C is conditioned proves false, and you end up not asserting C, you do implicate that C and generate a simple commitment to the truth of C.

Thus, on our theory of indicative conditionals as devices of conditional assertion, we would expect the use of biscuit conditionals like (1)–(3) to generate implicatures of their consequents, while no such implicature is generated by the use of normal indicative conditionals like (4)–(6). (Stating a normal conditional will instead generated an implicature that C is conversationally relevant.)

Note that we should expect the implicature of C's truth generated by a "biscuit" use of "A→C" to be uncancellable.[18] Why? If C were unconditionally asserted, then, of course, one would not just implicate, but actually state, that C, and one's commitment to C would be uncancellable. Now we're supposing that you've only conditionally asserted that C. But since, as we're also supposing, it's clear that the reason you've only conditionally asserted that C has to do with C's relevance, and not at all with its probability of truth, there's nothing here in the conditional nature of your assertion to render your commitment to the truth of C any more cancellable than it is where C is unconditionally asserted. And our expectation that this implicature (the implicature that C is true that's generated by a "relevance" or "biscuit" use of "A→C") will be uncancellable is not disappointed: "There are biscuits on the sideboard if you want some—but don't get me wrong, I don't mean to be implying that there are biscuits on the sideboard" does not come off smoothly, to say the least!

Some, we know, will be bothered by our use of such uncancellable implicatures, for the implicatures we allege seem most naturally classified, in Gricean terminology, as conversational, as opposed to conventional, implicatures, and cancellability is taken to be one of the key marks of conversational implicatures. In addressing such worries, it's important to remember the rationale behind the cancellability test in the first place: Given what conversational implicatures are, we should expect them typically to be cancellable; thus, where a theory posits the presence of a conversational implicature, but the alleged implicature proves to be uncancellable, that's typically a substantial mark against the theory. But, in the present case, for the reasons given in the above paragraph, we should expect, on our theory, that the relevant implicatures will be uncancellable. Thus, of course, that they indeed prove to be uncancellable is no mark against the theory.[19]

That a "biscuit" use of "A→C" generates an implicature that C is true, whether or not it results in an assertion of C, explains an important feature of relevance conditionals like (1)–(3). Among those who grant themselves the luxury of cordoning off such biscuit conditionals for separate treatment, a favorite view seems to be that they're simply ways of asserting their consequents, with the "if" clause apparently being some sort of semantic side-order that has no effect on what's asserted. On these views, then, even if A proves false, you've asserted that C.

Such a view responds to the sense that we generate a non-conditional (simple) commitment to C when we use "A→C" in a "biscuit" way. But, as we've just seen, our view, on which biscuit conditionals are integrated into a unified account of indicative conditionals, explains why that commitment to the truth of C is generated even where C ends up not being asserted. Thus, we don't have to give biscuit conditionals a separate treatment.

This may help explain away apparent counter-examples to our theory. We have appealed to our theory's ability to explain the use of biscuit conditionals as a source of support for that theory. Ironically, though, some of the best apparent counter-examples to the theory are biscuit conditionals. Thus, when Nuel Belnap reaches for an example of a conditional that is not a conditional assertion, it's just such a conditional he appeals to.[20] In fact, he uses Austin's example (our (1)):

> But I do know that "There are biscuits on the sideboard if you want some" is *not* generally used as a conditional assertion; for if there are no biscuits, even if you don't want any, it is plain false, not nonassertive. (Belnap 1970, p. 11)

While we find Belnap's verdict intuitively plausible, we don't find the intuition all that strong. To the extent that one shares Belnap's intuition that in the circumstances described (there are no biscuits on the sideboard, and your listener doesn't want any), you assert a falsehood by uttering (1), then one should take this to be a prima facie difficulty for our view. But we've seen how our view can handle this apparent problem. Given the fact that an utterance can seem false to us when it generates a false implicature, we should expect, on our theory, that there would be some intuitive pull toward thinking that one asserts something false in uttering "A→C" in a "biscuit" way where A and C are both false.

The theory of indicative conditionals as devices of conditional assertion, then, integrates biscuit conditionals into a unified account of indicative conditionals, while explaining the difference between them and normal indicative conditionals partly in terms of the importantly different implicatures that each use generates. And the theory explains the data that might lead you to adopt a different theory of biscuit conditionals if you gave yourself the luxury of providing a separate treatment of them. The interesting phenomenon of biscuit conditionals, then, generates significant support for taking indicative conditionals to be devices of conditional assertion.[21]

Notes

[1]For instance, for many combinations of P and Q, it will seem that (a) *P or Q* entails that *not-P→Q* but that (b) *P* does *not* entail *not-P→Q*. Given that (c) *P* entails *P or Q* and given the transitivity of entailment, at least one of these two powerful intuitions must be wrong. (It can't be that *Not-P→Q* is entailed by the weaker *P or Q* but not entailed by the stronger *P*.) Of course, under this much intuitive pressure, one might consider abandoning (c) instead of (a) or (b). This is the guts of what Frank Jackson has labeled the "paradox of indicative conditionals" (Jackson 1979).

[2]Quine doesn't endorse the view, but his attitude toward it strikes us as unclear. He writes: "Now under what circumstances is a conditional true? Even to raise this question is to depart from everyday attitudes. An affirmation of the form 'if p, then q' is commonly felt less as an affirmation of a conditional than as a conditional affirmation of the consequent." (At this point, he drops a footnote in which he writes, "I am indebted here to Dr. Philip Rhinelander.") He then elaborates: "If, after we have made such an affirmation, the antecedent turns out true, then we consider ourselves committed to the consequent, and are ready to acknowledge error if it proves false. If on the other hand the antecedent turns out to have been false, our conditional affirmation is as if it had never been made." Quine continues: "Departing from this usual attitude, however, let us think of conditionals simply as compound statements which, like conjunctions and alternations, admit as wholes of truth and falsity." But he doesn't let on—or at least doesn't clearly let on—that the reason he departs from the usual attitude is that he thinks it is in any way wrong about the conditionals of ordinary language (1950, pp. 12–13).

[3](Belnap 1970), p. 1. That Belnap doesn't there endorse the thesis that indicative conditionals are conditional assertions comes out most clearly in this passage: "But I do not know if there are in English any quite clear cases of if-thens used as conditional assertions" (p. 11). He then immediately goes on to give the counter-example that we deal with in section 4, below. An expanded version of Belnap's paper appeared as (Belnap 1973). We cite the earlier paper because, in addition to containing all the material we're directly concerned with here, it is easily accessible to more readers. But those who are interested in Belnap's views should know that the later paper contains significant material not in the earlier version.

[4](Edgington 1995); see especially pp. 287–291.

[5]For a good history of the idea, see (Milne 1997), pp. 197–212, which also collects, in its References, many of the relevant sources.

[6]See (Adams 1975). For an account of Adams's own relation to the conditional assertion view, see (Milne 1997), pp. 199–201.

[7]One of us (DeRose) believes that the common understanding that one should simply assert that P only when P's probability is very high is mistaken and that the real rule for simple assertion is that one should assert only what one *knows*. Still, the more popular "assert the highly probable" is close enough to getting things right for our present purposes. If one, like DeRose, believes that simple assertability goes by knowledge, then one will naturally think that conditional assertability goes by conditional knowledge. That is, one should assert that C on the condition that A only if, given the assumption of A, one knows that C. It is worth noting that the same lottery cases that cause problems for the probability account of simple assertion cause similar problems for the probability account of conditional assertion. Even where the probability of having lost the lottery is *very* close to 1, it still seems that you cannot assert "I've lost," in the case where you've already bought your ticket. Similarly, where you're considering buying a ticket, then, even if the conditional probability of your losing if you buy the ticket is *very* close to 1, it seems you still cannot state, "If I buy a ticket, I will lose." For defense of the knowledge account of proper assertion, see (Williamson 1996) as well as some of the sources referenced there. For an application of that account to lottery problems, see (DeRose 1996).

[8]The generalization (often called "Adams's Generalization," after E.W. Adams) that "A→C" is statable where the conditional probability of C on A is high, of course, only tells part of the story of when "A→C" can be appropriately stated. As with other assertions, more than *being well-enough positioned to state* is involved; conditional assertions, like other assertions, must also be *conversa-*

tionally relevant. But Adams's Generalization does at least a *fairly* good job of expressing the *being well-enough positioned* constraint on the statability of indicative conditionals. (But see note 7 above for doubts over whether a "high probability" account really captures when one is well enough positioned to assert, and a corresponding doubt about whether high conditional probability suffices for conditional assertability.)

[9]Here's a typical example of an indicative conditional to illustrate this point. The local zoo has obtained a new animal, and Fred is wondering what the animal's name is. You say, "If it's a gorilla, its name is Magilla." Now, if it turns out that the animal is a gorilla, and its name is indeed Magilla, it seems to most that you've asserted a truth. And if it's a gorilla, but its name is Mary, most will respond that you've asserted something false. But if the animal turns out to be an elephant, most will think you've not said anything that's either true or false. (Those who take polls of their students will find this to be especially true of students who have not yet had a symbolic logic class, and so have not yet encountered '⊃'. A poll I recently took of students in a large introductory philosophy class returned very strong results indeed.) All bets are off, according to the intuitions of most people. Indeed, if you *have* made a bet, you putting money down on "If it's a gorilla, its name is Magilla," and Fred betting against that, most think that if the animal turns out to be a non-gorilla, then *literally* all bets are off: Neither of you has to pay up. Insofar as the matter of who has to pay whom in a bet is to be correlated with the truth of the statement that's the object of the bet, that's further evidence that the conditional expresses nothing that's either true or false (evidence of a type that McDermott utilizes in his (1996)). Of course, there may be other explanations for why nobody should have to pay up here. But the direct intuition that the statement is neither true nor false is pretty strong in its own right. Similarly for many, in fact, most, indicative conditionals with false antecedents. But not all. (See the below note.)

[10]A more sustained development of this virtue of our theory would here both take up the fact that there are some exceptions to the generalization that we think an indicative conditional is truth-valueless where the antecedent is false, and would consider ways by which supporters of the leading theories might try to explain away our intuitions to the effect that indicative conditionals with false antecedents are truth-valueless. But to develop this virtue sufficiently would take more space than we can here give it in the context of trying to fairly quickly give the reader reason for taking our view seriously.

[11]Edgington argues, in support of the account of indicative conditionals as devices of conditional assertion, that speech acts other than assertion can be performed conditionally (1995, pp. 287–288), and argues (1995, p. 288) that the leading theories of indicative conditionals, both Stalnaker's account and theories according to which indicative conditionals are equivalent to material conditionals, cannot be generalized to cover cases in which the main clause is, for instance, a command.

[12]Prompted by Grice's initial comments, a considerable amount of sophisticated research has refined our understanding of relevance, especially the work of Sperber and Wilson (1995), Blakemore (1987, 1992), Carston (1998) and others. However, the details of these developments are not significant for our main point.

[13]Though we count this as an "implicature", Grice would not have, for, in a (rather underexplained) paragraph of "Logic and Conversation" (1989, pp. 41–42), he defines the term "implicature" so as to exclude commitments that are *trivially* generated by the supposition that the speaker is observing the Cooperative Principle and its associated conversational maxims. Grice's example: "On my account, it will not be true that when I say that p, I conversationally implicate that I believe that p; for to suppose that I believe that p (or rather think of myself as believing that p) is just to suppose that I am observing the first maxim of Quality on this occasion" (p. 42). But similarly, the supposition that the speaker thinks that p is relevant is just the supposition that she is observing the maxim of Relation, ("Be relevant"). Thus, that p is relevant would not be counted by Grice as an implicature generated by asserting that p.

Why this restriction? Grice writes that "it is not a natural use of language to describe one who has said that p as having, for example, 'implied,' 'indicated', or 'suggested' that he believes that p; the natural thing to say is that he has expressed (or at least purported to express) the belief that p." But this rationale is dubious. Though it is fine to say the speaker has expressed his belief, it seems *also* quite acceptable under many circumstances to say that he has implied that he believes that p. Consider the following:

A: His level of culpability depends, not just on whether his action was in fact dangerous, but also on what he believed. Did he ever say that he believed that his action would put his company at risk?

B: What he actually said was that it *would* put his company at risk. And by saying this, of course, he certainly implied that he believed that it would.

B's use of "implied" seems quite acceptable here. Likewise, when a speaker asserts that P, it can seem natural enough to say that she has implied that it's relevant that P.

Still, we think that Grice is on to an important distinction here. So, though we use "implicature" more broadly to include commitments Grice would exclude, we recognize that there are important differences between these conversational implicatures, which we subclassify as "trivial" implicatures, and the other ("non-trivial") conversational implicatures that Grice recognizes. Most importantly, perhaps, it seems that it will be more problematic to attempt to *cancel* a trivial conversational implicature than a non-trivial one. It's especially hard to cancel a trivial implicature simply by denying it. Consider, "P, but I don't believe that P." This does not come off smoothly, to say the least. It's paradoxical. In fact, of course, it's Moore's Paradox, which is the focus of Grice's cryptic paragraph on pp. 41–42 of (1989). Likewise, "P, but P isn't conversationally relevant" is quite puzzling, to say the least. The problem with attempts to cancel such trivial implicatures is that it becomes difficult to see just why in the world the speaker would be asserting that P if even these trivial commitments of asserting that P are supposed to be vitiated: Just what is the speaker doing asserting that P if he doesn't believe it or doesn't think it's relevant? (Note, however, that the implicature of P's relevance, though difficult to cancel, is not as hard to cancel as is the implicature that the speaker believes that P. We *can* make sense of canceling the implicature of relevance by means of a "but I don't mean to imply" construction: "P, but I don't mean to imply that P is relevant" can be sensibly taken to indicate that the speaker is unsure whether P is relevant. "P, but I don't mean to imply that I believe that P" is strange precisely because we presume that the speaker isn't uncertain about whether she herself believes that P. And the implicature of P's relevance can sometimes at least seem to be successfully canceled even by means of a simple denial of it. "P, but P isn't relevant" *can* work if, for instance, one has just been asked whether or not it's the case that P. But here, it seems, you do think that P is conversationally relevant in the relevant sense required for assertion, precisely because it answers the question you've just been asked. Your point in such a case is that P is irrelevant in some other sense, and that in the sense in which my asking you about P *makes* P relevant, P *shouldn't* be relevant—that for current purposes the question shouldn't have been asked, since its answer doesn't matter to those purposes. Rather than pursuing the details here, let's just leave it at this: the (trivial) implicature that P is relevant is somehow difficult to cancel, though not as difficult as is the (trivial) implicature that the speaker believes that P.) By contrast, to use one of Grice's examples, if you say to a motorist, "There is a garage round the corner" (1989, p. 32), the supposition that you're observing the maxim of Relation ("Be relevant") may well generate the implicature that the garage is open. But since this isn't a trivial implicature, we can make sense of an attempt to cancel it, even to cancel it by simply denying it. If you say, "There's a garage round the corner, but it isn't open," we're free to suppose that, for instance, it may be open later, and so that the garage is there is relevant after all. However, when you flatly deny that P is relevant, or that you believe it, it becomes very difficult for the listener to avoid the conclusion that you should have just kept quiet about it.

[14]This point would seem to hold for trivial implicatures (see the above note) as well as for the non-trivial ones. This is one of the reasons why we classify such trivial and difficult-to-cancel commitments as "implicatures."

[15]Sometimes, though the simple probability of C is plenty high enough for assertion, we might still resort to a conditional assertion of C if the conditional probability of C on A is appreciably higher still. This, though, would still be resorting to conditional assertion for reasons involving the probability of C's truth, and so would also be cases of type (1).

[16]Here, we use "implicature" in the "broad" way described above in note 13, a way broader than Grice's own use of the term. We shall have to be alert to the possibility that in this broader use, there will be implicatures that, though they're conversational (as opposed to conventional), are still uncancellable. See note 19, below, as well as the paragraph of text to which note 19 is attached and the paragraph preceding that one.

[17]Our use of "simple commitment" here parallels our use of "simple assertion" in section 2. There are two types of situation in which one has asserted that C: either one has simply asserted it, or one has conditionally asserted it and the condition is satisfied. Likewise, there are two types of situation in which one is committed to C: either one is simply committed to C, or one is conditionally committed to C and the condition is met. A simple commitment, then, is a commitment that does not result from a conditional commitment, just as a simple assertion is an assertion that does not result from a conditional assertion.

[18]What of the implicature of C's conversational relevance that is generated by a normal use of "A→C"? Our theory leads one to suspect that this implicature will be difficult to cancel in the same way that it's difficult to cancel the implicature that C is relevant that would be generated by an unconditional assertion of C. Why? If C were unconditionally asserted, then for the reasons of note 13 above, and in the way pointed to in that note, the implicature that C is relevant would be difficult to cancel. Now we're supposing that C is only conditionally asserted. But since, as we're also supposing in the case of a normal conditional, it's clear that the reason you've only conditionally that C has to do with C's probability of truth, and has nothing to do with doubts about C's relevance, there's nothing here in the conditional nature of your assertion to render your commitment to the relevance of C any easier to cancel than it is where C is unconditionally asserted.

[19]There is still room here for terminological disputes. Some may take it to be *definitional* of conversational implicatures that they be cancellable. And some, for reasons similar and parallel to those of Grice (see note 13, above), might insist that the commitments our theory predicts shouldn't even be counted as "implicatures" at all. We've already cited our reasons for preferring our own, broader use of "implicature". And, given that the relevant commitments that our theory predicts will be generated by the assertion of indicative conditionals should be counted as "implicatures" in the first place, we think it's best to count them *conversational* implicatures because we think the best construal of what *makes* an implicature conversational is a matter of how it is generated. That conversational implicatures are typically uncancellable, then, is, for us, just a mark of conversational implicatures that we should expect them to frequently display—but would not expect them to display in the cases at hand. The implicatures we should expect, on our theory, to be generated by statements of "A→C" should be counted as conversational because their generation can be calculated by Gricean-style conversational reasoning, given the semantic function of "A→C" (to conditionally assert that C (on the condition that A)), and so they needn't be posited as an extra semantic feature of indicative conditionals, in the way conventional implicatures have to be posited. But such terminological matters shouldn't affect our argument here. You may follow us in calling the commitments our theory predicts will be generated by statements of indicative conditionals "implicatures", or you may insist on calling them something else. Either way, we can still do both of the following. First, we can explain an important difference between biscuit and normal indicative conditionals in terms of the importantly different simple commitments generated by the two types of statements. Second, in the argument that's about to follow in the text, we can use the commitment to C that our theory predicts will be generated, even where A is false, by a biscuit-style statement of "A→C" to explain why a biscuit-style statement of "A→C" can seem to assert something false where A and C are both false and in fact nothing is asserted at all.

[20]It's not clear to us whether Belnap reaches for such a conditional because he thinks that biscuit conditionals are particularly *bad* candidates for a "conditional assertion" reading, and so provide the best and clearest counter-examples, or because he thinks such conditionals are, are thought by some to be, or can appear to be, particularly *good* candidates for a "conditional assertion" reading, so the counter-example shows that that reading doesn't even generally work where it might be thought to have its best chance.

[21]Thanks to Kent Bach, Graeme Forbes, Mitchell Green, Roy Sorensen, Timothy Williamson, and the editors of *Noûs* for helpful comments.

References

Adams, E.W. (1975) *The Logic of Conditionals* (Dordrecht: Reidel).

Austin, J.L. (1970) "Ifs and Cans", in Austin, *Philosophical Papers*, second ed. (London: Oxford University Press, 1970): 205–232.

Belnap, Nuel D., Jr. (1970) "Conditional Assertion and Restricted Quantification," *Noûs* 1: 1–12.

Belnap, Nuel D., Jr. (1973) "Restricted Quantification and Conditional Assertion," in H. Leblane, ed., *Truth, Modality and Syntax* (Amsterdam: North-Holland): 48–75.

Blakemore, Diane. (1987) *Semantic Constraints on Relevance* (Oxford: Blackwell).

Blakemore, Diane. (1992) *Understanding Utterance: An Introduction to Pragmatics* (Oxford: Blackwell).

Carston, Robyn and Seijji Uchida. (1998) *Relevance Theory: Applications and Implications* (Amsterdam: John Benjamins Publishing Co.).

DeRose, Keith. (1996) "Knowledge, Assertion and Lotteries," *Australasian Journal of Philosophy* 74: 568–580.

Edgington, Dorothy. (1995) "On Conditionals," *Mind* 104: 235–329.

Grice, Paul. (1989) *Studies in the Way of Words* (Cambridge, Mass.: Harvard University Press).

Jackson, Frank. (1979) "On Assertion and Indicative Conditionals," *Philosophical Review* 88: 565–589.

McDermott, Michael. (1996) "On the Truth Conditions of Certain 'If'-Sentences," *Philosophical Review* 105: 1–37.

Milne, Peter. (1997) "Bruno de Finetti and the Logic of Conditional Events," *British Journal for the Philosophy of Science* 48: 195–232.

Quine, W.V.O. (1950) *Methods of Logic* (New York: Henry Holt and Company).

Sperber, Dan and Deirdre Wilson. (1995) *Relevance: Communication and Cognition*, second ed. (Oxford: Blackwell).

Stalnaker, Robert. (1975) "Indicative Conditionals," *Philosophia* 5: 269–286; reprinted in Jackson, ed., *Conditionals* (Oxford: Clarendon Press, 1991): 136–154. Page reference to (1991).

Williamson, Timothy. (1996) "Knowing and Asserting," *Philosophical Review* 105: 489–523.