

Capturing the Relationship between Conditionals and Conditional Probability with a Trivalent Semantics

Daniel Rothschild*

1 An Old Problem

Consider these two sentences:

- (1) a. It's likely that Georgia will join the E.U.
- b. It's likely that if Russia needs foreign aid, Georgia will join the E.U.

Intuitively, it seems that both sentences are of the form: *it's likely that x*, where, in the first case, $x = \textit{Georgia will join the E.U.}$ and, in the second case, $x = \textit{if Russia needs foreign aid, Georgia will join the E.U.}$.¹ Understanding how these

*I am grateful to Chris Barker, Paul Égré, Linda Rothschild, James Shaw, Robbie Williams, and Seth Yalcin for pointing out some errors and giving helpful comments on earlier drafts. Also thanks to Kai von Fintel and Robbie Williams for giving me some crucial references.

¹Kratzer (1981, 1986) denies this syntactic parsing, claiming that the function of “if”-clauses is to restrict higher up modal quantifiers in the sentence, in the case of (1-b) the natural choice is the probability operator *it's likely that*. While there is much to be said for her view, I think that, generally speaking, her strategy will not easily explain all the recalcitrant facts about probability and conditionals that the proposal presented below aims to. For example, many judge (i) to have the same truth conditions as (1-b):

- (i) It's likely to be true that if Russian needs foreign aid, Georgia will join the E.U.

sentences get their meaning as a function of the meaning of both the embedding construction, *it's likely that ...*, and the embedded sentence, x , would seem an easy task. The natural way would be to assume i) that in each case the particular sentence x expresses a proposition, and ii) that sentences of the form *it's likely that x* are true iff x expresses a proposition that has a probability greater than one-half. In other words the in general *it's likely that x* is true iff $p(x) > .4$. This simple proposal leaves various questions unanswered. For instance, what does it mean for x to have a probability greater than one-half? Is it that the speaker has a credal state that assigns a subjective probability of more than one-half to x , or is that there is a more objective or intersubjective probability at stake? We will leave questions along these lines aside, since a more basic problem confronts this proposal.

The problem goes as follows: Suppose we understand (1-b) to be true iff $p(x) > .5$ where x is *if Russia needs foreign aid, Georgia will join the E.U.*. It seems that as a matter of fact (1-b) is judged true iff the conditional probability that Georgia will join the E.U given that Russia needs foreign aid is high. To see this think about two sorts of cases: in the first, you think it's very likely that Russia will need foreign aid, but you think it's very unlikely that it's both the case that Russia will need foreign aid and Georgia will join the E.U. Then your conditional probability of Georgia joining the E.U given that Russia will need foreign aid cannot be very high, and correspondingly you must judge (1-b) as false. By contrast, if you think it's likely Russia will need foreign aid and only slightly less likely that Russia will need foreign aid *and* Georgia will join the E.U. then you would seem to need to judge (1-b) as true. Generalizing, *it's likely that if a then c* is true iff the conditional probability of c given a is greater than one half. Our prior assumptions about the meaning of sentences of the form *it's likely that x* along with this last observation should lead use to the conclusion that conditionals express propositions whose probability is just the conditional probability of the consequent given the antecedent. This hypothesis is what is often called *the equation*: the probability of "if a then c " is equal to $p(c|a)$ (see Edgington, 1995, and references therein). The problem with this conclusion, is that, in a certain sense, *there is no proposition that satisfies the equation*.

However, in this case it seems rather implausible that the antecedent of the conditional embedded under *it's true that* pops out to restrict *it's likely*. Other cases can be found of this sort: such as cases where there is no linguistically present modal at all and speakers or audiences merely make *judgments* of the probability of various conditional sentences without expressing them in sentence. But I will leave further discussion of these points to another occasion, so that this discussion may be considered as presenting an *alternative* to Kratzer's account of (1-a) and (1-b). See also von Stechow (2007) for further discussion of the limitation of Kratzer's approach to conditionals.

Here’s the sense:² suppose we assume that there’s a set of possible worlds W , and any given sentence, x , expresses a proposition by picking out a subset of W , i.e. the worlds where x is true. Let’s then say that any given credal state is a probability function defined over the power-set of W , i.e. a function that tells you for any given set of possible worlds how much credence you have that the actual world is one of those. Now let $a \rightarrow c$ be *if Russia needs foreign aid, then Georgia will join the E.U.*, and a and c , be the antecedent and consequent respectively. Since $a \rightarrow c$, a , c are all sentences that we can have credences in, they each pick out a subset of W . Suppose that on someone’s credal state, p , $p((a \rightarrow c) \wedge a) > 0$, and $p(c|a) \neq p(c) \neq p(a)$, and none of those values are 0 or 1. Further suppose, as suggested, that $p(a \rightarrow c) = p(c|a)$. Now, it is easy to show that there will exist another probability function p_1 such that $p_1(a \rightarrow c) \neq p_1(c|a)$.³

What’s the problem with this? Well, what we’ve just shown is there isn’t any subset of W to which every probability function assigns a probability equal to the conditional probability of c given a . This suggests that our semantic theory will not be able to assign a *general* meaning to $a \rightarrow c$ which 1) applies across different credal states *and* 2) fits into the natural account of the semantics for sentences of the form *it’s likely that x*. This is not a happy situation, since $a \rightarrow c$, intuitively, has some sort of uniform meaning.⁴

In the next section, I will give an extension of probability theory that covers trivalent expressions, and I will show that on this theory we can treat conditionals as expressing (trivalent) propositions whose probability is equal to the conditional probability of the consequent given the antecedent. Using this trivalent probability theory we can give a natural, compositional semantics of expressions of the form *it’s likely that x* for conditional and non-conditional x ’s. I will then conclude, in the third and last section, with some remarks on this system and its overall plausibility.

²This is a variation on Lewis’s first “triviality theorem” (Lewis, 1976).

³One can construct p_1 for instance by making it the result of conditionalizing p on $\neg(\neg a \wedge (a \rightarrow c))$. It will then follow that $p_1(a \rightarrow c) < p(c|a)$, since $p_1(a \rightarrow c) < p(a \rightarrow c)$ but $p_1(c|a) = p(c|a)$.

⁴Some have argued that the proposition expressed by a conditional sentence varies with the epistemic state of the speaker—this is, for instance, a direct consequence of Kratzer’s theory of epistemic modals and indicative conditionals. In this case, the argument I gave in the previous paragraph would have no force. There are, in fact, further problems with maintaining “the equation” even if one allows conditionals to express different propositions relative to different credal states, but I won’t discuss them here. See Edgington (1995) and Bennett (2003) for discussion and citations to the major results.

2 A Trivalent Extension of Probability Theory

We'll here give an interpretation of probability functions and conditionals that allows us to assign propositions to conditionals. The view I will explore is almost exactly the same as that put forward by de Finetti (1936) in a short lecture.⁵ I'll define a trivalent language, L , and describe how we can assign values analogous to probabilities to it. The language, L formed in the usual way, with an extra operator \parallel :

- $A, B \dots$ are atomic sentences.
- if α and β are sentences so are $\alpha \wedge \beta, \alpha \vee \beta, \alpha \parallel \beta$ and $\neg\alpha$.

Again, let W equal a set of possible worlds. We'll think of each sentence in L as having as its meaning a function from W to the set of values $\{T, F, U\}$, where we think of T as false, F as true, and U as undefined, so each sentence is true, false or undefined at each "possible world". Each atomic sentence is *bivalent* in the sense that it takes W into $\{T, F\}$, in other words, it is true or false at every possible world. We get trivalence through the extra binary operator \parallel , whose semantics is defined as follows: $\alpha \parallel \beta$ is undefined for every member of W where β is not T (i.e. where β is false or undefined) and otherwise has the same truth-value as α . Graphically, its truth table is as follows:⁶

A	B	A \parallel B
T	T	T
F	T	F
T	F	U
F	F	U
T/F	U	U
U	T/F	U

We will understand the trivalent values of complex formulas to be, otherwise, de-

⁵Thanks to Guy Politzer and Jean Baratgin, via Paul Égré, for this reference. See Milne (1997) for discussion of de Finetti's theory. The idea of using a trivalent semantics for conditionals generally put forward by Belnap (1970). Von Fintel(2007) suggests Belnap's idea as a possible solution to problems with conditionals and probabilities, thus putting forward de Finetti-style approach to conditionals an probability. Another working out of Belnap's idea can be found in Huitink (2008, 2010).

⁶This is the same trivalent truth table for conditionals given by McDermott (1996). On his account, however, the probability of a conditional is *not* its conditional probability (p. 20), rather it is simply the probability it is true *simpliciter*.

terminated by the strong-Kleene truth tables⁷

For any given probability function, p , defined over the powerset of W , we will call p' its *extended probability*.⁸ The definition of p' is straight-forward: p' is a function that goes from sentences in L to real numbers in the interval $[0,1]$ such that for any sentence α in L , if R is the set of worlds where α is T and S is the set of worlds where α is T or F , $p'(\alpha) = p(R)/p(S)$ where $p(S) > 0$ and is otherwise undefined.

Some notation: For any formula α in L , let $p(\alpha)$ be equal to the probability of all the worlds in W is assigned T by α (we will generally only discuss p as applied to bivalent formulas). Likewise $p(\alpha|\beta)$ is defined as $\frac{p(\alpha\wedge\beta)}{p(\beta)}$ where $p(\beta) > 0$. By contrast, $p'(\alpha||\beta)$ is just the extended probability function p' applied to the formula $\alpha||\beta$. So, for α and $\beta \in L$, $p(\alpha|\beta)$ and $p'(\alpha||\beta)$, despite the superficial similarity, have a different syntax since the former uses a conditional probability function applied to two formulas, α and β whereas the latter just applies an extended probability function to the formula $\alpha||\beta$. To not multiply possible expressions too much let us suppose there is no analogue of conditional probability for p' so that $p'(\alpha|\beta)$ is not well-formed. Luckily, there are some very close connections between $p(\alpha|\beta)$ and $p'(\alpha||\beta)$, just as there are between $p(\alpha)$ and $p'(\alpha)$:

Fact 1. *For any bivalent formula α , $p(\alpha) = p'(\alpha)$.*

Proof. This is trivial: since α is true or false in all worlds in W , $p'(\alpha) = \frac{p(\alpha)}{p(W)} = \frac{p(\alpha)}{1}$. □

Fact 2. *For any bivalent formulas, α and β , $p(\alpha|\beta) = p'(\alpha||\beta)$, if $p(\beta) > 0$.*

Proof. Since $\alpha||\beta$ is only defined over the worlds in W in which β is true and is only true in the worlds in which α and β are true: $p'(\alpha||\beta) = \frac{p(\alpha\wedge\beta)}{p(\beta)}$. □

⁷As follows:

- $\alpha \wedge \beta$ is true in all worlds in which α and β are true, false in all worlds in which one of α or β is false, and undefined in all other worlds.
- $\alpha \vee \beta$ is false in all worlds in which α and β are false, true in all in which where one of α or β is true, and undefined in all other worlds
- $\neg\alpha$ is true in all worlds where α is false, false in all worlds where α is true, and undefined in all other worlds.

Since we will mostly talk about conjoining bivalent formulas with \wedge and \vee we can usually just think about these connectives as classical.

⁸Cantwell (2006) independently suggests this means of extending probability functions to trivalent formula, though he does not use it for the same purposes as I have done here.

Extended probability functions allow us to interpret conditional probabilities in a different way than usual: we just think of conditional probabilities as being the extended probabilities of certain trivalent formulas formed out of the bivalent ones. So we don't need a specially defined conditional probability operator for extended probability functions: the \parallel operator will do the work for us. Extended probabilities can do more than this: after all, *any* trivalent formula will have an extended probability. So we can also discuss the probability of formulas that have no analogue in standard probability theory, such as $(\alpha\parallel\beta)\parallel\gamma$.

To return to our original discussion of sentences, (1-a) and (1-b). We will understand sentence of the form *it's likely that x* to be true iff $p'(x) > .5$. We will also understand sentences of the form *if a then c* (where *a* and *b* express bivalent propositions) to express the proposition $c\parallel a$. It then follows that a sentence of the form *it's likely that if a then c* is true iff $p'(c\parallel a) > .5$ iff $p(c\parallel a) > .5$, the desired result.

I'll go through a concrete example to illustrate how this system works (taken from Grice, 1967/1989, lecture IV).

Chess Yog and Zog play chess according to normal rules, but with the special condition that Yog has white 9 out of 10 times and there are no draws. Up to now there have been a hundred games. When Yog had white, he won 80 out of 90. When he had black he lost 10 out of 10. Suppose we are talking of a randomly chosen game of these hundred. It's then true to say of this game:

- (2) There's a 50/50 chance that (if Yog lost, then he drew white).

Let W = Yog had white and L = Yog lost. The embedded sentence in (2), "If Yog lost, then he drew white" can be formalized as $W\parallel L$. Let us suppose that there are 100 possible "worlds" each one corresponding to the selection of one of the 100 games and the probability that we are in each of these worlds is .01. Then L is true at 20 of these worlds and W is true at 80 of them. $W\parallel L$ is true or false at 20 of these worlds, and of those it is true at 10 and false at 10. So $p'(W\parallel L) = .5$, which explains why we accept sentence (2).

3 Comments

Syntax and Semantics of Conditional Probability The trivalent system is a way of thinking about conditional probabilities as the probabilities of certain

trivalent formulas, rather than as the ratio of probabilities of two formulas. This is simply not possible in standard probability theory in which the formula $\alpha|\beta$ as it appears in $p(\alpha|\beta)$ has no independent status. Since we utter complete sentences, i.e. conditionals, that bear some relation to $\alpha|\beta$, as it appears in standard probability theory, it is perhaps worthwhile describing a system on which we do give something akin to $\alpha|\beta$ a free-standing interpretation.

Relation to Natural Language Conditionals Let us assume that the indicative conditional operator \rightarrow is just a backwards \parallel since we have seen evidence that $p(\alpha|\beta) = p(\beta \rightarrow \alpha)$ and we showed that $p'(\alpha|\beta) = p(\alpha|\beta)$, for bivalent α and β .

Logical We validate various plausible logical inferences involving conditionals using our trivalent truth-table definition. Consider $\gamma \rightarrow (\beta \rightarrow \alpha)$, which some (e.g. Lewis, 1976) propose is equivalent to $(\gamma \wedge \beta) \rightarrow \alpha$. On our interpretation of \parallel and \wedge this equivalence obtains (and thus its probabilistic analogue does as well).

Of course, I have not yet discussed which particular consequence relation we are using for this trivalent logic. The choice of this relation will obviously have a large effect on how we understand the proposal.⁹ Huitink (2008) suggests Strawson entailment: entailment on the assumption that the every formula is either true or false. We could consider other alternatives, though.

Semantics of Conditionals The trivalent system treats conditional sentences as having the same type of semantic value as non-conditional sentences while still maintaining the link between conditional sentences and conditional probability. It do so by treating all sentences as potentially trivalent and treating probability as a feature of trivalent sentence. By this trick we eliminate the need for a sharp distinction between conditional and non-conditional sentences such as that posited by Edgington (1995) and Bennett (2003). Admittedly we do partially affirm their conclusion that conditionals lack truth values, in the sense that we do make conditionals neither true nor false when their antecedent is false.

Trivalence We use trivalence in a different way than it is used to model either vagueness or presuppositions. Soames (1989) observes that the use of triva-

⁹This point was emphasized to me by Paul Egge, James Shaw, and Robbie Williams.

lence for vagueness and presupposition are clearly distinct: one cannot treat the lack of truth values in vague expressions in the same way one treats them in presuppositional expressions. The trivalence in the system here seems like it cannot be assimilated to either of these two, different uses. There are many reasons for this. Here are two: 1) conditionals do not generate presuppositions that their antecedent is true,¹⁰ and 2) the probabilities of vague or presupposition-bearing sentences do not behave as in the account above. Trivalence, on one level, is a technical trick that allows sentences to encode a three-fold partition of possibilities rather than the usual two-fold one. Not all uses of this trick to characterize various linguistic phenomena fit together.

Plausibility Just how independently plausible are these truth-conditions for conditionals? Do they correspond to our *judgments* of the truth-value of conditionals? Not obviously, as many indicative conditionals when uttered seem to have clear truth-values. Suppose my friend Bob is made entirely of combustible paper. There is a fire in his office building, but, unbeknownst to me, he is vacationing in the Mexican Riviera at the time. Seeing the fire but not knowing if Bob is in his office building or not, I might utter:

(3) If Bob is in his office, he's going to go up in flames.

In this situation, suitably elaborated, it seems plausible that my utterance of (3) is true, despite the fact that the antecedent is false. Many more cases along these lines can be constructed. This should make us somewhat skeptical of semantic legislation that makes indicative conditionals with false antecedents truth-valueless.¹¹ So the kind of truth-valuelessness that conditionals might have, is not as with presupposition failure, something akin to falsity, nor is it, as with borderline instances of vague predicates, something akin to semantic underdetermination. What is it then?

Assertion Conditions Let us accept for the moment that one can only assert an indicative conditional $\alpha \rightarrow \beta$ if one is uncertain about the truth value of α . Many authors, such as Kratzer (1986) and Yalcin (2007), propose that the further condition for asserting an indicative conditional, $\alpha \rightarrow \beta$ is essentially

¹⁰von Stechow (2007) observes this point.

¹¹McDermott (1996, p. 2) argues in a different case to the opposite conclusion. While I agree that in *some* cases the trivalent theory is more intuitive, my point is just that often it is not at all intuitive.

that one can rule out the possibility that any α world is a $\neg\beta$ world. These are exactly the conditions under which one can assert the trivalent conditional with the knowledge that it is not false (with the material conditional by contrast one can assert it with knowledge that it is true in this circumstance). So, if we are to take the trivalent account on board, the rules governing its assertion can be stated as follows:

1. Only assert $\alpha \rightarrow \beta$ if you are uncertain about the truth-value of α .
2. Only assert $\alpha \rightarrow \beta$ if you are sure it is not false.

These rules are not so implausible: 1. is something that almost every account of conditionals needs to stipulate one way or the other (or sketch some more or less ad hoc Gricean explanation of), while 2. is natural enough in the sense that in the bivalent case it simply amounts to an injunction to say what you are sure is true.

References

- Belnap, Nuel (1970). Conditional assertion and restricted quantification. *Noûs*, 4:1–12.
- Bennett, Jonathan (2003). *A Philosophical Guide to Conditionals*. Oxford.
- Cantwell, John (2006). The laws of non-bivalent probability. *Logic and Logical Philosophy*, 15(2):163–171.
- Edgington, Dorothy (1995). On conditionals. *Mind*, 104(414):235–329.
- de Finetti, Bruno (1936). La logique de la probabilitié. *Actes du congrès international de philosophie scientifique*, Fasc. IV:31–9.
- von Fintel, Kai (2007). If: The biggest little word. Slides from talk at Georgetown University Roundtable.
- Grice, Paul (1967/1989). Logic and conversation. In *Studies in the Ways of Words*. Harvard University Press.
- Huitink, Janneke (2008). *Modals, Conditionals and Compositionality*. Ph.D. thesis, Radboud Universiteit Nijmegen.

- (2010). Quantified conditionals and compositionality. *Language and Linguistics Compass*, 4:42–53.
- Kratzer, Angelika (1981). The notional category of modality. In H.-J Eikmeyer and H. Reiser (eds.), *Words, Worlds, and Contexts*, pages 38–74. Walter de Gruyter.
- (1986). Conditionals. *Chicago Linguistics Society*, 22(2):1–15.
- Lewis, David (1976). Probabilities of conditional and conditional probabilities. *Philosophical Review*, 8:297–315.
- McDermott, Michael (1996). On the truth conditions of certain if-sentences. *The Philosophical Review*, 105(1):1–37.
- Milne, Peter (1997). Bruno de finetti and the logic of conditional events. *The British Journal for the Philosophy of Science*, 48(2):195–232.
- Soames, Scott (1989). Presuppositions. In D. Gabbay and F. Guenther (eds.), *Handbook of Philosophical Logic*, volume IV, pages 553–616. Dordrecht.
- Yalcin, Seth (2007). Epistemic modals. *Mind*, 116:983–1026.