Causal relevance of conditionals: semantics or pragmatics?

van Rooij, R.; Schulz, K.

Link to publication

Robert van Rooij and Katrin Schulz
Institute for Logic, Language and Computation (ILLC)
University of Amsterdam

Abstract

In this paper we argue that the antecedent of a (non-analytic) conditional is *causally* relevant to the consequent, ... at least if standard background conditions hold. Natural counterexamples to the causal relevance analysis are argued to be cases where the standardly assumed background condition(s) do not hold.

Causal Relevance: Semantics or pragmatics? Causality even for Even-if conditionals.

1 Introduction

The following sentence is inappropriate and misleading:

(1)   If it was sunny in Berlin yesterday, there are COVID casualties in Brazil today.

Why? Because this conditional sentence strongly suggests that what happened in Berlin (the antecedent) is relevant to what happens in Brazil (the consequent). The sentence is misleading because we know no such relevance relationship exists.

Although what makes this conditional sentence misleading is clear, the standard semantic analyses of conditionals by themselves do not predict that anything is wrong with it, because they do not make relevance as part of the meaning of conditionals. According to one such theory (e.g. Adams, 1975), the only thing that counts for the meaning of conditional is that the consequent is likely, or probable, given the antecedent. Given that the consequent of the above conditionals is likely, or even certain, Adams' semantic analysis does not predict that the whole conditional sentences is appropriate to use. The same prediction follows from the other popular semantic analysis of conditional sentences (e.g., Stalnaker, 1968; Kratzer, 2012), which demands, instead of relevance, that the consequent is true in all most similar/normal antecedent worlds.

We will discuss two ways to tackle this problem: according to the first **semantic** solution (e.g. Douven, 2008, 2016) conditionals are used as explanations, and relevance is built in into the meaning of the conditional. We discuss two probabilistic ways to work out such a semantic approach. On the first semantic analysis, 'If A, then C' (from now on abbreviated by $A \Rightarrow C$) is assertable if $p(C|A) - p(C|\neg A) = \Delta P_A^C \gg 0$. This analysis seems natural from a psychological point of view, because measure ΔP_A^C is used frequently to measure the learned association between A and C (cf. Shanks, 1995). Unfortunately, the use of this notion gives rise to various empirical problems. These problems motivate a second **causality-based** semantic analysis proposed by van Rooij & Schulz (2019). Unfortunately, it turns out that this analysis won't be appropriate for all (indicative) conditionals, because the analysis is problematic for some examples where causal relevance seems explicitly denied. To meet these problems, we discuss the natural **pragmatic** solution: conditional $A \Rightarrow C$ is true, or assertable, if $P(C|A) \approx 1$, but relevance comes out because of the implicature that the consequent is not believed. Although this analysis seems appealing, we will still argue in favour of a slightly adapted causal analysis after all, according to which the measure of causal power is influenced by **background conditions** as well.

2 Two relevance-based analyses of conditionals

One consequence of standard analyses of conditionals is their acceptance of an inference known as **conjunction sufficiency**, that $A, C \models A \Rightarrow C$, i.e., the assumption that the (known) truth of A and C suffices for the assertion of (indicative) conditional $A \Rightarrow C$, also without any

connection between A and C. This principle, however, seems ad odds with an appropriate use of conditionals, as exemplified by (1) and (2):

(2)   If I raise my little finger, there will be rainfall this winter.

where we should suppose that the finger is indeed raised and that there will be rainfall. Conjunction sufficiency is problematic as well if there is a connection, but where the connection is not of the right kind, for instance because the antecedent made the consequent not likely enough or led to it only through an indeterministic process:

(3)   a.   If he worked hard, he passed (he worked hard, but he passed only because he cheated, too).
         b.   If the coin was tossed, it fell heads.

Arguably, these examples are odd, because of a lack of (enough) relevance between antecedent and consequent. If one takes conditionals $A \Rightarrow C$ to express a positive relevance relation between antecedent and consequent (Douven, 2008; Spohn, 2013), what comes to mind is that instead of demanding that the conditional probability of C given A, $p(C|A)$ is high, one demands that this probability is higher than that of $p(C)$, i.e., $p(C|A) \gg p(C)$, or equivalently, $\Delta P_A^C = p(C|A) - p(C|\neg A) \gg 0$.

If this is all that is demanded, this gives rise to a number of surprising predictions. First, on this analysis *Contraposition* $(A \Rightarrow C \therefore \neg C \Rightarrow \neg A)$ and (the conditional variant of) *Denying the Antecedent* $(A \Rightarrow C \therefore \neg A \Rightarrow \neg C)$ are predicted to be valid (Suppes, 1970), in spite of many counterexamples. Furthermore, although *Transitivity* $(A \Rightarrow B, B \Rightarrow C \therefore A \Rightarrow C)$ and *Strengthening of the Antecedent* $(A \Rightarrow C \therefore (A \wedge B) \Rightarrow C)$ are correctly predicted to be invalid (or so we think, following Adams, 1975), under the above ΔP_A^C relevance-based analysis of conditionals, counterexamples to such inferences can be found very easily. For linguistic reasons, however, such counterexamples should be really exceptional, because strengthening of the antecedent, for instance, is important for licensing of NPIs in antecedents of conditionals. Indeed, for Stalnaker (1975) these are pragmatic *reasonable inferences*.

Consider the example due to Eells & Sober (1983) illustrated by the picture below. In the story that belongs to this picture, individuals either smoke, S, or not, $\neg S$, at time 1, t_1. At later time t_2 they either get heart attacks, H, or not, and still later, at t_3, they either experience heart pains, P or not. For the picture, we started with 100 representative smokers and 100 representative non-smokers. Note that $S \Rightarrow H$ and $H \Rightarrow P$ both hold according to the ΔP-based analysis: $p(H|S) = a > b = p(H|\neg S)$ and $p(P|H) = \frac{79}{110} > \frac{40}{90} = p(P|\neg H)$ (in fact, both $p(P|S \wedge H) = w > x = p(P|S \wedge \neg H)$ and $p(P|\neg S \wedge H) = y > z = p(P|\neg S \wedge \neg H)$ hold). Still it won't be the case that $S \Rightarrow P$ holds according to the ΔP analysis, because $p(P|S) = \frac{45}{100} < \frac{74}{100} = p(P|\neg S)$.
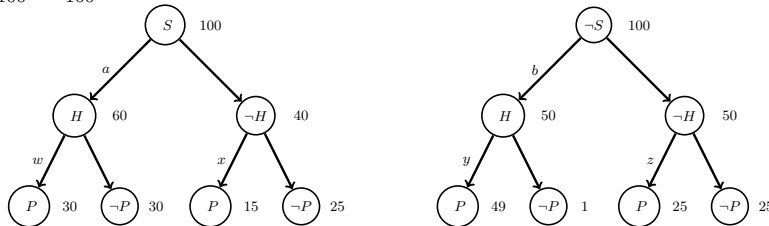
The reason that transitivity doesn't go through here is that the probability of getting heart pains depends in this story not just on whether one had a heart attach at t_2, but also on whether one smoked at t_1. In the story, the probability of pain given heart attack and smoking is less than the probability of pain given heart attack and *not* smoking. To get transitivity, we need to make three assumptions: (i) the probability of having heart pains at t_3 depends only on whether there has been a heart attack at t_2, and is thus *independent* on what happened at t_1. Thus, $p(P|H \wedge S) = p(P|H \wedge \neg S)$ and $p(P|\neg H \wedge S) = p(P|\neg H \wedge \neg S)$. In causal modelling this is known as the *Markov property*. The second assumption, (ii), is that to determine whether

$H \Rightarrow P$ holds, we should not only check whether $p(P|H) > p(P|\neg H)$, but we should also hold the other (background) factors (in this case only whether S) constant, thus we should check whether $p(P|H \wedge S) > p(P|\neg H \wedge S)$ and $p(P|H \wedge \neg S) > p(P|\neg H \wedge \neg S)$ hold. (this holds in this case). Conditions (i) and (ii) strongly suggest that the **conditionals should be given a causal relevance analysis**, because this is how causal relevance is determined. But to assure transitivity, we need a third assumption, (iii), as well: there is no other factor at t_2 caused by S that is causally relevant to what happens at t_3. If we make all these three assumptions, we can easily prove that transitivity holds, also on our suggested analysis of conditionals: $S \Rightarrow H, H \Rightarrow P \models S \Rightarrow P$. Condition (iii) is a **minimality assumption** (enforced by, e.g., the probability functions with maximal entropy): the assumption that there are no (for S, H and P relevant) other *direct* causal relevance relations than those mentioned: only $S \Rightarrow H$ and $H \Rightarrow P$. Indeed, if we added $S \Rightarrow D$ and $D \Rightarrow \neg P$ for new proposition D as premises, it wouldn't follow anymore that the *indirect* causal relation $S \Rightarrow P$ is predicted to hold.

The minimality assumption on causal relations can also explain why inferences like *Strengthening of the Antecedent* $(A \Rightarrow C \therefore (A \wedge B) \Rightarrow C)$ and *Contraposition* $(A \Rightarrow C \therefore \neg C \Rightarrow \neg A)$ – although not valid in general – go through in most cases. Strengthening holds, because by minimality $B \not\Rightarrow \neg C$. Contraposition holds because from $A \Rightarrow C$ and the assumption that C can only be caused by A due to the minimality inference, it follows that $A \Leftrightarrow C$ and we can thus conclude from $\neg C$ that $\neg A$.

Douven (2008) also argued that for a relevance-based analysis of (indicative) conditionals of form $A \Rightarrow C$ it is insufficient to just demand that $\Delta P_A^C > 0$. He did so for a rather different reason than we did above, however, and claimed that conditionals require something else as well. Douven demanded that on a proper relevance-based analysis of conditional $A \Rightarrow C$, it should not only be that (i) $p(C|A) > p(C|\neg A)$, but also that (ii) $p(C|A)$ should be high, i.e., close to 1.

Interestingly, this combination of demands means that $p(C|A) - p(C|\neg A)$ should be close to $1 - p(C|\neg A)$, which means in turn that $\frac{p(C|A) - p(C|\neg A)}{1 - p(C|\neg A)} = \Delta^* P_A^C$ should be high. But this later measure is exactly how Cheng's (1997) notion of **causal power** can be estimated (under certain conditions), and van Rooij & Schulz (2019) show that on a causal power analysis, many conditionals, including diagnostic ones, can be handled naturally. Moreover, under the proposed causal power analysis of conditionals and the above mentioned minimality assumption (due to the attested fact (cf. Mill, 1843; Brem & Rips, 2000) that hearers typically ignore alternative causes of C on interpreting $A \Rightarrow C$), it immediately follows that assertability of the conditional normally 'goes by' the corresponding conditional probability, $p(C|A)$. A good analysis of conditionals should also account for 'analytic' conditionals like 'If x is a bachelor, x is a man'. Fortunately, our analysis in terms of $\Delta^* P_A^C$ can, because if $A \models C$ and $p(C) \neq 1$, it follows immediately that $\Delta^* P_A^C = 1$, its maximal value.

3 A pragmatic analysis of relevance

According to the semantic analyses of conditionals discussed in the previous section, the antecedent should be positively relevant to the consequent. But it is clear that this is not always the case: there are examples where positive relevance is not required:

(4)   a.   If Mary leaves the party early, Bill will be unhappy,
        b.   but if Mary doesn't leave the party early, Bill will *still* be unhappy.

This suggests that the presumed positive relevance of the antecedent should be due to a pragmatic cancellable implicature. A pragmatic strategy presupposes a semantic analysis of indicative conditionals. What should the basic semantic analysis be? Although they have something to say for themselves, we strongly believe that the material or strict conditional account of indicative conditionals won't do, for one thing because such analyses predict that indicative conditionals always allow for inferences like 'strengthening of the antecedent', 'transitivity' and 'contraposition', which seems false. For another thing, with such semantic analyses it is

hard to capture the intuition that the probability of the conditional 'If I pick an ace, it is going to be the ace of clubs' intuitively 'goes with' the corresponding conditional probability.

The most direct route to account for the above intuition is that conditional 'If A, then C' (abbreviated from now by $A \Rightarrow C$) just expresses the speakers' conditional probability of C given A, i.e., $p(C|A)$. But we don't want to assume that the **assertability value** of $A \Rightarrow C$, i.e. $\mathcal{AV}(A \Rightarrow C)$ is simply $p(C|A)$. The reason is that for assertability we have to take into account pragmatic presuppositions and implicatures as well (cf. Skyrms, 1980). Instead, let us assume with Skyrms that the **basic** assertability value of 'If A, then C', $\mathcal{BAV}(A \Rightarrow C)$, is $p(C|A)$. The assertabiility value of the indicative conditional $A \Rightarrow C$, i.e., $\mathcal{AV}(A \Rightarrow C)$, is then based on its basic assertability value plus the (uncancelleble) appropriateness condition that A is really possible, i.e., $0 \ll p(A)$, and some cancellable pragmatic implicatures.

It only seems natural that the conditional is asserted because the alternative assertion that the consequence is believed/known C could not yet be appropriately made. Thus, $A \Rightarrow C$ conversationally implies that $p(C) \ll 1$. Let's see how this works out:

1. $\mathcal{BAV}(A \Rightarrow C) \approx 1$, and thus $p(C|A) \approx 1$[1]   by assertion.
2. $p(C) \ll 1$ by implicature.
3. $p(C|A) - p(C) \gg 0$ by 1 and 2.
4. $p(C|A) \gg p(C|\neg A)$,   because $p(C|A) \gg p(C)$ iff $p(C|A) \gg p(C|\neg A)$.
5. Thus $\mathcal{AV}(A \Rightarrow C) = \frac{p(C|A) - p(C|\neg A)}{1 - p(C|\neg A)} \approx 1$ because $p(C|A) \approx 1$ (from 1).

Thus, the demand that relevance measure $\Delta^* P_A^C$ should be high discussed in section 2 falls out as a pragmatic implicature! The idea is that (i) a standard **semantic** analysis suffices according to which $\mathcal{BAV}(A \Rightarrow C) = p(C|A) \approx 1$, and (ii) that causal relevance follows as a **pragmatic** cancellable conversational implicature. The implicature of (4-a) that Bill's unhappiness would be caused by Mary's leaving early is cancelled by (4-b).

Also for **counterfactual**, or subjunctive **conditionals** the antecedent is normally causally relevant to the consequent. Indeed, such conditionals have much in common with indicative conditionals, although in contrast to indicative conditionals, for counterfactuals it is typically assumed that $p(A) = 0$, if p is the current subjective probability function (cf. Stalnaker, 1975). Indeed, Adams (1975) proposed that for counterfactuals we should look at the *prior* probability state, where A was still possible. This gives rise to the correct predictions for many examples, but not all. To see this, look at the following examples due to Morgenbesser and Edgington, respectively:

(5)   a.   If John had bet on heads, he would have won.
       b.   If John had caught the flight, he would be in Paris now.

Intuitively, $\mathcal{BAV}(A \Rightarrow C)$ would be $\frac{1}{2}$ for (5-a), because we assume that the past chance of winning conditional on betting heads was $\frac{1}{2}$. Similarly, we may be almost certain about (5-b), given that we take the past chance of being in Paris conditional on taking the flight to Paris to be very high. However, suppose for (5-a) that we know that the coin came up heads. We then can be certain that John would have won, because how the coin would have landed would not have been influenced by my betting on the outcome. An analogous situation can arise with respect to (5-b) as described by Edgington (1995). Assume that we know that the plane crashed because the on-board computer broke down. Given this, we should have a probability close to 0 in (5-b), because if John had caught the flight, he would be dead now. Thus, Adams' prior probability account is not completely correct:[2] we have to take into account facts that occurred later, but that were *causally independent* of the antecedent.

In the above cases the antecedent was causally relevant to the consequent. But it seems that that doesn't always have to be the case.[3] Consider the following variant of Tichy's (1976)

---

[1]On a Kratzerian analysis, it is demanded that $p(C|A) = 1$ and the same reasoning goes through.

[2]Skyrms (1980) argued for a prior *propensity* analysis to account for other counterexamples.

[3]Consider also an example by Stalnaker (1968), 'If the Chinese enter the Vietnam conflict, the United States will use nuclear weapons,' uttered by someone who believed (in the 1960s) that the US was going to use nuclear

well-known example due to Frank Veltman:

> Suppose that Jones always flips a coin before he opens the curtains to see what the weather is like. Heads means he is going to wear his hat in case the weather is fine, whereas tails means he is not going to wear his hat in that case. Like above, bad weather invariably makes him wear his hat. Now suppose that today heads came up when he flipped the coin, and that it is raining. So, Jones is wearing his hat.

Now the question is whether the following sentence is acceptable:

(6)   If the weather had been fine, Jones would (still) have been wearing his hat.

Intuitively, the answer is 'yes'. This means that we conclude that the antecedent doesn't have any causal effect on the consequent. It seems that the use of 'still', like in (6), indicates that causal relevance doesn't play a role here, suggesting that the inference to causal relevance is indeed a cancellable conversational implicature.

4   Some doubts on the semantic/pragmatic picture

First, there are **doubts** about the **semantic** analysis that $\mathcal{BAV}(A \Rightarrow C) = p(C|A)$. Rothschild (2013) argues by example that at least sometimes $p(A \Rightarrow C) \neq p(C|A)$. As it turns out,[4] to account for Rothschild's example, $\mathcal{BAV}(A \Rightarrow C)$ should not be equated with $p(C|A)$, but rather with Pearl's (2000) causal measure $p(C|do(A))$ that he used to analyse counterfactual conditionals (see also Schulz (2011)). Pearl's $p(C|do(A))$ can be estimated by $\sum_{b_i \in \mathbf{B}} p(b_i) \times p(C|A, b_i)$, a measure already proposed by Skryms (1984) to account also for indicative conditionals. Here $\mathbf{B} = \{b_i\}$ is the partition of the set of worlds still possible into causally relevant background conditions for C.[5] Notice that the new proposal that $\mathcal{BAV}(A \Rightarrow C) = p(C|do(A))$ is still compatible with the pragmatic analysis: by a similar reasoning as above it now follows that $\mathcal{AV}(A \Rightarrow C) = \frac{p(C|do(A)) - p(C|do(\neg A))}{1 - p(C|do(\neg A))} \approx 1$, due to the cancellable implicature that C is not believed.[6]

Second, Krzyzanowska (2019) and Skovgaard-Olsen et al. (2019) argue on experimental grounds **against** the above **pragmatic** analysis: the relevance effect doesn't seem to be cancellable, which according to Griceans is a defining feature of conversational implicatures. They point out, for instance, that an attempt to cancel the relevance effect by an explicit denial is rated by participants as completely inappropriate:

(7)   *If Mary left the party early, Bill was unhappy, though these things have nothing to do with each other.

The inappropriateness of (7) suggests that it is impossible to make a conditional claim if antecedent and consequent have nothing to do with each other. However, that doesn't mean that the relevance effect of the antecedent is thus not due to an implicature. It just means, or so we will argue, that the inference to a relevance relation comes about in another way.

---

weapons in any case. Douven (2016) proposes that Stalnaker's example is not a counterexample to a relevance-based analysis because this is a **concessive** conditional for which the analysis does not apply. We think such an 'ambiguity'-analysis is an all too easy way out.

[4]Due to space limitations we cannot explain this here, unfortunately.

[5]If each element of the partition $\mathbf{B}$ is probabilistically independent of A, $p(C|do(A)) = p(C|A)$. According to Kaufmann (2004), for the determination of $p(A \Rightarrow C) = \sum_{b_i \in \mathbf{B}} p(b_i) \times p(C|A, b_i)$, it doesn't have to be that $\mathbf{B}$ is a partition of *causally* relevant background factors.

[6]Note, though, that for Rothschild's example also $\frac{p(C|do(A)) - p(C|do(\neg A))}{1 - p(C|do(\neg A))} = p(C|do(A) = p(\neg D)$, thus the example does not show anything about a preference for an analysis of $p(A \Rightarrow C)$ as $p(C|do(A))$ compared to the causal relevance-based analysis.

5 Causal relevance w.r.t. background assumptions

To work towards our proposal for how the implicature comes about, let us look again at (6), on the assumption that we interpret the conditional in terms of an intervention. Given that (6) is appropriate, we can learn something about the underlying causal model. First, $\mathcal{BAV}(A \Rightarrow C)$ can only be high in case $p(C|do(A))$ is high. If we know, or strongly believe, that $p(C) \approx 0$, this means that A has a positive causal effect on C, and thus that the causal picture is either $\mathcal{A} \rightarrow^+ \mathcal{C}$, or something of which it is is part. In case C is already known, however, $\mathcal{BAV}(A \Rightarrow C)$ can only be high in case the causal model must be such that A doesn't make any difference anymore. This can be if A is causally completely unrelated to C. However, based on the inappropriateness of examples like (7), that seems unnatural. Another possibility is that now A is only a contributing cause of C, and the causal model is an AND-gate of the form $\mathcal{A} \rightarrow^+ \mathcal{C} \ ^+\!\!\leftarrow \mathcal{B}$: A and B are both causally necessary conditions for C to hold. The latter seems to hold in the example involving Jones and his hat.

Before we take background, or enabling, assumptions into account, let us first see how the measure $\Delta^* P_A^C$ that we discussed in section 2 for the analysis of conditionals follows from Cheng's (1997) causal analysis. To do so, assume with Cheng (1997) that events of type a have unobservable **causal powers** to produce events of type c denoted by p_{ac}. Causal power p_{ac} is taken to be a *local* property of a, and thus very different from $p(c|a) = \frac{p(a \wedge c)}{p(a)}$, which is only a global property. We assume for now that events of type c are either due to events of type a, or due to other events of type o, thus the relevant causal structure is an OR-gate like $\mathcal{A} \rightarrow^+ \mathcal{C} \ ^+\!\!\leftarrow \mathcal{O}$, and $p(c|\neg a, \neg o) = 0$. It follows that $p(c)$ can be determined as follows:

(8) $\quad p(c) \quad = \quad p(a) \times p_{ac} + p(o) \times p_{oc} - (p(a \wedge o) \times p_{ac} \times p_{oc})$.

From this we immediately derive p_{ac}, the **causal power of a to generate** c. This is nothing else but the probability of c, conditional on a and $\neg o$:

(9) $\quad p_{ac} \quad = \quad p(c|a, \neg o)$.[7]

One problem with this notion is that it depends on o, and this is not always observable. Fortunately, if we assume that a and o are, or are believed to be, independent, Cheng (1997) shows that we can estimate p_{ac} after all. The estimation of p_{ac} given the above OR-gate is exactly the probabilistic relevance notion that we mentioned in section 2:

(10) $\quad p_{ac} \quad = \quad \frac{p(c|a) - p(c|\neg a)}{1 - p(c|\neg a)} \quad = \quad \Delta^* P_a^c$.

So far, to determine p_{ac} it was assumed that a *by itself* can cause c. Of course, this is a simplification for almost all cases of causal attributions. Striking a match, for instance, does not by itself cause it to light. Certain **background**, or **enabling conditions** have to be in place: there must be oxygen in the environment, the match must be dry, etc.. In fact, for deterministic causation we can think of $\Delta^* P_a^c = p_{ac}$ as modelling the probability of the background conditions. Suppose that a can interact with b to cause c. The causal power of conjunctive cause ab to produce c is then

(11) $\quad p_{ab,c} \quad = \quad p(c|a, \neg o, b) \quad = \quad \frac{P(c|a,b) - p(c|\neg a, b)}{1 - p(c|\neg a, b)} \quad = \quad \Delta_b^* P_a^c$.

If b is necessary for c and a and b jointly sufficient, $p(c|\neg b) = 0$ & $p_{ab,c} = 1$, $p_{ac} = p(b)$.

With this machinery we can now tackle examples where the antecedent seems (causally) irrelevant to the consequent. How is it possible for party-lover Bill that on the one hand (12) and (4-a)-(4-b), repeated here as (13-a)-(13-b), are appropriate,

(12) $\quad$ *Even* if Mary leaves the party early, Bill will be happy.

---

[7]More generally, we should use $p(c|do(a), \neg o)$ using causal intervention, but we ignore that here.

(13) a. If Mary leaves the party early, Bill will be unhappy,
      b. but If Mary doesn't leave the party early, Bill will *still/nevertheless* be unhappy.

although on the other hand (7) is inappropriate, suggesting that Mary's leaving early is taken to be causally related with Bill's unhappiness? To tackle this question, let's say that a stands for 'Mary leaves the party early', c for 'Bill will be unhappy', and alternative o for, say, 'Sue leaves the party early.' Now consider the following conditional probability tables (where $p(c|a, o, b) = 1 - (p_{ab,\neg c} \times p_{ob,\neg c})$):

b	a	o	$p(c)$	$p(\neg c)$
1	1	1	0.97	0.03
1	1	0	0.9	0.1
1	0	1	0.7	0.3
1	0	0	0	1
0	–	–	0	1

Table 1

b	a	o	$p(c)$	$p(\neg c)$
1	1	1	0.97	0.03
1	1	0	0.9	0.1
1	0	1	0.7	0.3
1	0	0	0	1
0	–	–	1	0

Table 2

The tables differ only in the boxed entries. Notice that according to both tables,

(14) $p_{ab,c} \quad = \quad p(c|a, \neg o, b) \quad = \quad 0.9 \quad = \quad \frac{p(c|a,b) - p(c|\neg a,b)}{1 - p(c|\neg a,b)} \quad = \quad \Delta_b^* P_a^c.$

We will assume that a conditional of the form $a \Rightarrow c$ is appropriate iff p_{ac} is high, where p_{ac} is now determined as follows, taking background b into account as well:

(15) $p_{ac} \quad = \quad p(c|a, \neg o, b) \times p(b) + p(c|a, \neg o, \neg b) \times p(\neg b).$

How can we account for the appropriateness of (12) represented by $a \Rightarrow \neg c$? Let's assume that Bill's supposedly unhappiness if Mary leaves the party early is based on the background assumption that, say, Bill not only loves parties but is even more desperately in love with Mary. So, what (12) asserts is that $p_{a\neg c}$ is high, and it implies that $p_{ab,c}$ is high, higher than $p_{ob,c}$. What is asserted and what is implied are not in contrast with each other according to table 1, if the speaker has good reason to believe that in contrast to the standard assumption, background b is false. Similarly, Table 2 can explain why conditionals (13-a) and (13-b) are both appropriate. These sentences claim that p_{ac} and $p_{\neg ac}$ are high, respectively, and according to table 2 this can only be if the background condition b is taken to be false. On the other hand, it is still the case that a and c (are assumed to) have something to do with each other, for indeed a has high causal power to produce c in case the assumed background condition b is in place. Thus, or so we argue, examples like (13-a)-(13-b) do not falsify a causal relevance analysis of conditionals. However, this causal relevance is dependent on the standard background assumption, an assumption the speaker implies to be false by (13-b). As it turns out, the background assumption doesn't even have to be taken to be standard, it can also be on issue, as shown by the following example due to Reinhard Muskens:

(16) a. Will Bill be unhappy, if Mary leaves the party early?
      b. No, even if Mary leaves early, Bill will be happy.

6 Conclusion

We propose that with a conditional of the form $A \Rightarrow C$ it always has to be the case that A is causally relevant to C. This implies that antecedent A normally makes a difference to C, but that this doesn't always have to be the case. It is not the case if the standardly, or on issue, assumed background condition does not hold. One can say that with the use of conditional $A \Rightarrow C$ a speaker **asserts** (and thus not just implies) that on the relevant background condition, A is causally relevant to C (where this relevance can be negative, as

in (12)), that (s)he conversationally implies that this background condition holds, but that a speaker can make clear – e.g. by using markers like 'still' and 'even' – that the background condition does not hold.

References

Adams, E. (1975), *The Logic of Conditionals*, D. Reidel Publishing Company, Dordrecht.

Brem, S. and L. Rips (2000), 'Explanation and evidence in informal arguments', *Cognitive Science*, **24**: 573-604.

Cheng, P. (1997), 'From Covariation to Causation: A Causal Power Theory', *Psychological Review* **104**: 367-405.

Douven, I. (2008), 'The evidential support theory of conditionals', *Synthese*, **164**: 19-44.

Douven, I. (2016), *The Epistemology of Indicative Conditionals*, Cambridge University Press, Cambridge.

Edgington, D. (1995), 'On Conditionals', *Mind*, **104**: 235-329.

Eells, E. and E. Sober (1983), 'Probabilistic causality and the question of transitiviity', *Philosophy of Science*, **50**: 35-57.

Kaufmann, S. (2004), 'Conditioning against the grain: abduction and indicative conditionals', *Journal of Philosophical Logic*, **33**: 583-606.

Kratzer (2012), *Modals and Conditionals*, Oxford University Press.

Krzyzanowska, K. (2019), 'What is wrong with false-link conditionals?', *Linguistics Vanguard*, 5.

Mill, J.S. (1983), *A System of Logic*, John W. Parker, West Strand, London.

Pearl, J. (2000), *Causality: Models, Reasoning and Inference*, Cambridge University Press.

Rooij, R. van and K. Schulz (2019), 'Conditionals, causality and conditional probability', *Journal of Logic, Language and Information*, **28**: 55-71.

Rothschild, D. (2013), 'Do indicative conditionals express propositions?', Noûs, **47**: 49-68.

Schulz, K. (2011), 'If you'd wiggled A, then B would've changed. Causality and Counterfactual Conditionals, *Synthese*, **179**: 239-251.

Shanks (1995), *The Psychology of Associative Learning*, Cambridge University Press, Cambridge.

Skovgaard-Olsen, N., Collins, P., Krzyzanowska, K., Hahn, U., & Klauer, K. C. (2019), 'Cancellation, negation, and rejection', *Cognitive Psychology,* 108, 42-71.

Skyrms, B. (1980), *Causal Necessity*, Yale University Press, New Haven.

Skyrms, B. (1984), *Pragmatics and Empiricism*, Yale University Press, New Haven.

Spohn, W. (2013), 'A ranking-theoretic approach to conditionals', *Cognitive Science*, **37**: 1074-1106.

Stalnaker, R. (1968), 'A theory of conditionals', in J. W. Comman et al. (eds.), *Studies in Logical Theory: essays*, Blackwell, Oxford, pp. 98-112.

Stalnaker, R. (1975), 'Indicative conditionals', *Philosophia*, **5**: 269-286.

Suppes, P. (1970), *A Probabilistic Theory of Causality*, North-Holland, Amsterdam.

Tichy, P. (1976), 'A counterexample to the Stalnaker-Lewis analysis of counterfactuals', *Philosophical Studies*, **29**: 271-273.