

Unspecific antecedents*

Wolfgang Schwarz

1 Introduction

I want to comment on an old objection to the “similarity analysis” of counterfactuals, and on a more recent, but related, argument for counterfactual skepticism.

According to the similarity analysis, a counterfactual $A > C$ is true iff C is true at all A worlds that are most similar, in certain respects, to the actual world.¹ The old objection that I have in mind is that the similarity analysis fails to validate *Simplification of Disjunctive Antecedents* (SDA), the inference from $(A \vee B) > C$ to $A > C$ and $B > C$.² Imagine someone utters (1a) on a hot summer day.

- (1) a. If it had rained or snowed, the match would have been cancelled.
- b. \leadsto If it had rained, the match would have been cancelled.
- c. \leadsto If it had snowed, the match would have been cancelled.

In this context, worlds where it snows are presumably more “remote” (less similar to the actual world) than worlds where it rains. The closest worlds where it rains *or* snows are then all worlds where it rains. By the similarity analysis, (1a) should imply nothing about what would have happened if it had snowed. Yet (1a) seems to imply (1c).

Much has been written about this objection. Some have argued that it calls for a different approach to the semantics of counterfactuals.³ Others have suggested that the word ‘or’ in the antecedent does not express Boolean disjunction.⁴ Yet others have argued that the apparent entailment in cases like (1) is merely an implicature.⁵ I will eventually side with this last hypothesis. But first I want to explore a more direct line of response.

I claimed that on a hot summer day, snow worlds are more remote than rain worlds. Friends of the

* Thanks to Alan Hájek for helpful comments on an earlier version and for many great conversations through the years.

1 The analysis is commonly associated with [Stalnaker 1968] and [Lewis 1973]. My formulation assumes the “Limit Assumption” but not the “Uniqueness Assumption”.

2 This objection to the Stalnaker-Lewis account was raised in [Fine 1975] and [Nute 1975].

3 See, for example, [Warmbrod 1981] or [Fine 2012].

4 See, for example, [Lewis 1977] or [Alonso-Ovalle 2008].

5 See, for example, [Loewer 1976] or [Franke 2011].

similarity analysis, however, have emphasized that the similarity standards that figure in the analysis need not match our intuitive judgements of overall similarity between worlds.⁶ If rain worlds and snow worlds are equally “close”, the similarity analysis has no trouble with the inference in (1).

Any non-trivial version of the similarity analysis will still render SDA invalid. Unless all worlds are equally close, there will be instances of SDA in which one disjunct is more remote than the other, and these instances are predicted to be invalid. But this may not be a problem. After all, there appear to be counterexamples to SDA, such as (2), from [McKay and Van Inwagen 1977].

- (2) a. If Spain had fought with the Allies or the Axis [in World War II], it would have fought with the Axis.
- b. \rightsquigarrow If Spain had fought with the Allies, it would have fought with the Axis.

Can the similarity analysis be made to fit the data, if only we plug in the right similarity measure? This hypothesis (although not expressed in these words) has recently been put forward by Alan Hájek. It is worth exploring.

Hájek also notes that SDA-type effects don’t just arise with disjunctive antecedents. Consider the following examples.

- (3) a. If it had rained, the match would have been cancelled.
 - b. \rightsquigarrow If it had rained heavily, the match would have been cancelled.
-
- (4) a. If Bob had wine, his medication would cause a severe reaction.
 - b. \rightsquigarrow If Bob had red wine, his medication would cause a severe reaction.

In each of (1), (3), and (4), the ‘a’ conditional has an *unspecific* antecedent, leaving open, respectively, whether it rains or snows, how strongly it rains, and what kind of wine Bob consumes. The conditional with the unspecific antecedent appears to entail conditionals in which the antecedent is made more specific.

This brings me to my second topic. Hájek has invoked the generalized form of SDA illustrated by (3) and (4) to support his “counterfactual skepticism” – the claim that most ordinary counterfactuals are false. (See [Hájek 2021b], [Hájek 2021a], [Hájek 2023])

Take (4a). Let’s assume that Bob is on a medication that is guaranteed to cause a severe reaction if he consumes alcohol. In this context, (4a) seems true: his medication would cause a severe reaction if he had wine. But wait. What about de-alcoholised wine? That wouldn’t trigger a reaction (let’s assume). Since (4a) implies that Bob would have an allergic reaction no matter what type of wine he had, it is false. Or so Hájek would argue.

In general, ordinary counterfactuals often have unspecific antecedents. For them to be true, Hájek argues, they would have to remain true on every way of resolving the unspecificity. But we can often find unusual resolutions that would render them false.

I will return to this argument in the final section of this paper. First I want to explain how Hájek’s version of the similarity analysis can account for inferences like (1), (3), and (4), as well as some other

⁶ See, for example, [Lewis 1979: 42–48], [Stalnaker 1984: 127ff.], and [Bennett 2003: 195ff.].

phenomena. I will then point at yet other facts – some well known, some new – that raise problems for Hájek’s proposal, and indeed for any purely semantic treatment of SDA-type effects. I will sketch a (somewhat new) pragmatic derivation of these effects in section 6, before returning to the argument for counterfactual skepticism in section 7.

2 Chance-based similarity

I have stated the similarity analysis in abstract generality. We can be more concrete if we focus on conditionals whose antecedent describes a particular event, as in all the examples from the previous section.

It has often been observed that we seem to evaluate such counterfactuals by a “Rewind, Revise, Re-Run” heuristic: to evaluate whether $A > C$ is true, we consider worlds that are just like the actual world until shortly before the time of the antecedent (“rewind”), at which point the antecedent becomes true (“revise”) and history unfolds again in accordance with the general laws of nature (“re-run”). The conditional is true iff C is true in all such worlds. (These are the “closest” A worlds to the actual world, in terms of the similarity analysis.)

How, exactly, does the antecedent become true, if it is not already true in the actual world? In a deterministic world, this requires either a “miracle”, as argued in [Lewis 1979], or a divergence across the entire past, as argued in [Dorr 2016]. Let’s avoid this difficult choice by assuming that the world is thoroughly indeterministic, so that a counterfactual antecedent describes a non-actual outcome of an actual chance setup.

Assume, then, that the antecedent A specifies an outcome of a chance setup S at a suitable *fork time* t .⁷ Following the “Rewind, Revise, Re-Run” heuristic, we might now say that $A > C$ is true iff the chancy laws of nature guarantee C on the condition that S yields A . More concisely: $A > C$ is true iff the chance of C given A at t is 1; that is, iff $Ch_t(C/A) = 1$, where Ch_t is the (actual) chance function at the fork time t .⁸ Variants of this analysis have been defended by several authors throughout the years – including Skyrms [1980], Loewer [2007], Pearl [2000], Leitgeb [2012], and, most recently, Hájek [2023].

The chance analysis looks like a straightforward instance of the similarity analysis. The similarity ordering would rank worlds that conform to the actual laws before worlds that don’t, and sort worlds of the first kind by the time until they first diverge from the actual world.⁹

One would probably want to finesse the analysis to balance lateness and smoothness of the divergence. Consider an example from [Bennett 2003: 211]: ‘If the German army had reached Moscow on 15 August 1941, ...’. In the actual world, the Germans were far from Moscow on the morning of 15

⁷ See [Bennett 2003: ch.13] for more on the notion of a fork.

⁸ According to orthodox, Archimedian probability theory, events with chance 1 are not actually “guaranteed” to obtain. For any point in time, a radium atom’s chance of not decaying at that point is 1. The present account seems to predict that if (say) a radium atom had been created today then it would not have decayed at x , for every future point in time x . Let’s ignore this problem.

⁹ With this ordering, the standard similarity analysis implies that counterfactuals with true antecedents are true whenever they have a true consequent. If such pseudo-counterfactuals are to be evaluated by conditional chance, one probably needs an “antecedent-relative” similarity semantics.

August. There was, however, a tiny chance that they would quantum tunnel to the Moscow outskirts. Since there were few Russian defenders in the city at the time, they would then almost certainly have encountered little resistance. But we don't want to say that if the Germans had reached Moscow on 15 August then they would (probably) have encountered little resistance, or that they would have been baffled how they got there. To make the right predictions, we need an earlier fork time that allows the Germans to reach Moscow in a less dramatic way.¹⁰

As presented above, the chance analysis also only covers counterfactuals in which the antecedent describes a single, specific event. What about counterfactuals of the form $(A \vee B) > C$,¹¹ where A and B describe different events? Some have proposed to extend the analysis by stipulating that $(A \vee B) > C$ is true iff $A > C$ and $B > C$ are both true.¹² This would validate SDA, and make the account incompatible with a similarity analysis. Hájek suggests that no special stipulation is needed: $(A \vee B) > C$ is true iff $Ch_t(C/A \vee B) = 1$, where t is the fork time for $A \vee B$.¹³ What does this predict for the inference from $(A \vee B) > C$ to $A > C$ and $B > C$?

Suppose, to begin, that A and B are *fork-mates*, meaning that they have the same fork time t . (Perhaps they describe alternative outcomes of the same chance setup.) Plausibly, $A \vee B$ then has that same fork time. By the probability calculus, $Ch_t(C / A \vee B) = 1$ entails $Ch_t(C / A) = 1$ and $Ch_t(C / B) = 1$, provided $Ch_t(A) \neq 0$ and $Ch_t(B) \neq 0$. That is, the SDA inference is truth-preserving whenever the disjuncts (and their disjunction) are fork-mates and both have a positive chance at the fork time.

In the examples from section 1, the disjuncts are plausibly fork-mates. Recall the “match” example:

- (1) a. If it had rained or snowed, the match would have been cancelled.
- b. \rightsquigarrow If it had rained, the match would have been cancelled.
- c. \rightsquigarrow If it had snowed, the match would have been cancelled.

Physics tells us that the weather is thoroughly chancy. Even on a hot summer day, there is a small chance that it will start to snow. Both *rain* and *snow* therefore have positive chance, in the context of the example. On the chance account, this is enough to vindicate the inference in (1).

What about the case of (2), where we are not tempted to apply SDA?

- (2) a. If Spain had fought with the Allies or the Axis, it would have fought with the Axis.
- b. \rightsquigarrow If Spain had fought with the Allies, it would have fought with the Axis.

¹⁰ Friends of chance-based or “interventionist” accounts sometimes suggest that their analysis is conceptually simpler and clearer than standard similarity semantics. (See, e.g., [Pearl 2000: 239f.]) But the analysis is only simple and clear once we have identified the right causal model or the right chance function. Bennett’s example illustrates that this requires careful judgment.

¹¹ I use ‘ \vee ’ for classical disjunction. I generally assume that ‘or’ expresses \vee , setting aside the contrary view of, for example, [Lewis 1977].

¹² Such a stipulation is often considered for the interventionist approach. See, for example, [Briggs 2012], [Günther 2017], [Ciardelli et al. 2018].

¹³ This assumes that we can find a suitable chance function Ch_t that is defined for $A \vee B$. Pearl’s models do not generally provide such a function. I am not convinced that the chances of quantum physics could fill the gap. Let’s (again) ignore this problem.

A speaker who utters (2a) arguably believes that there was no chance that Spain would have fought with the Allies. Indeed, (2a) seems to convey just this belief. The chance account explains how: if there had been a chance that Spain fights with the Allies then (2a) would be false.

The chance account also seems to make good predictions about cases like (3) and (4), where an antecedent is unspecific but not disjunctive. Suppose A^+ is a more specific alternative to A . Whatever this means, it plausibly requires that A^+ entails A . By the probability calculus, $Ch_t(C / A) = 1$ then entails $Ch_t(C / A^+) = 1$, provided $Ch_t(A) \neq 0$. By the chance account, $A > C$ therefore entails $A^+ > C$ whenever A and A^+ are fork-mates and A^+ has a positive chance at the fork time. (SDA is a special case, where A is a disjunction and A^+ one of its disjuncts.)

With more than two resolutions of the relevant unspecificity, we can construct cases that combine elements of (1) and (2). Consider (5).

(5) If I had gone for a run today, I would feel less tired.

Suppose I usually go for a run either in the morning or in the early afternoon. Today I did not go for a run, and I feel tired. In this context, (5) seems to imply that I would feel less tired if I had gone for a run in the morning, and also if I had gone for a run in the early afternoon. It does not suggest that I would feel less tired if I had gone for a run at, say, 3 am. The unspecific counterfactual does not seem to distribute over all resolutions of the unspecific antecedent, but only over the “realistic” ones.

The chance account might explain this effect. The “realistic” resolutions are those with positive chance at the fork time. Since the antecedent of (5) describes an interval that spans all of today, we may assume that the relevant fork time is at around the start of that day. When I utter (5), however, I arguably take for granted that there is no chance that I would have gone for a run at 3 am.¹⁴

(5) draws attention to cases in which different resolutions of an unspecific antecedent have different fork times. Let’s look at a disjunctive example:

- (6) a. If I had gone for a run at 7am or 2pm, I would feel less tired.
- b. \rightsquigarrow If I had gone for a run at 7am, I would feel less tired.
- c. \rightsquigarrow If I had gone for a run at 2pm, I would feel less tired.

The inference looks good, even though the fork time for (6b) is arguably different from that for (6c). We can still vindicate the inference if we assume, as above, that the fork time t for (6a) lies at or before the interval covered by the antecedent, so that it coincides with the fork time for (6c). Let t' be the later fork time for (6b). Assuming that $Ch_t(C / A)$ and $Ch_{t'}(C / B)$ are not zero, $Ch_t(C / A \vee B) = 1$ entails $Ch_t(C / A) = 1$ and $Ch_{t'}(C / B) = 1$. Under plausible assumptions about the dynamics of chance, $Ch_t(C / A) = 1$, in turn, implies $Ch_{t'}(C / A) = 1$.¹⁵

I have one more piece of good news, before I bring in the bad news.

¹⁴ The case is similar to that of Pollock’s coat, discussed in [Nute 1980b: 104]. Either case appears to challenge the similarity criteria proposed in [Lewis 1979], which would favour later fork times even within the interval of time described by the antecedent. In fact, these cases appear to cast doubt on any version of the similarity analysis, as I’ll explain at the end of the next section.

¹⁵ We need to assume that the chance of an event conditional on another cannot decrease from 1. Informally, if the state of the world at t entails that the only possible outcome of 7 am at t' is C then this state of the world can’t evolve into

It is often said that alleged counterexamples to SDA like (2a) feel unnatural or marked. As [Lassiter 2018] points out, they become much more natural if we insert a probability operator:

- (7) a. If Spain had fought with the Allies or the Axis, it would probably have fought with the Axis.
b. \rightsquigarrow If Spain had fought with the Allies, it would probably have fought with the Axis.

This is a challenge for accounts that treat SDA as unrestrictedly valid. It is not a problem for the chance account.

To apply the chance account, we first need to know how it extends to ‘would probably’ counterfactuals. A natural idea is that ‘if A , would probably C ’ is true iff the relevant chance of C conditional on A exceeds some contextual threshold.¹⁶ That (7a) does not entail (7b) then follows from the fact that $Ch_t(C / A \vee B) > \theta$ does not entail $Ch_t(C / A) > \theta$, even if both disjuncts have a positive chance at t .

In sum, Hájek’s chance account seems to reconcile the similarity analysis with our judgments about counterfactuals with unspecific antecedents.

Alas, things are not so simple. Let’s have a look at some more data.

3 Context and content

Return once more to the ‘match’ case (1). According to the chance account, the inference from (1a) to (1c) is licensed by the fact that even on a hot summer day there is a positive chance of snow. But does the inference really rest on this non-obvious fact about chance? I suspect that hearers would be inclined to infer (1c) from (1a) even if they don’t think the weather is so thoroughly chancy. Counterfactuals like (1a) seem to support SDA even if one of the disjuncts is assumed to have chance zero.

For a more controlled example, imagine I’m about to draw a ball from an urn. You bet that I’ll draw a red ball. The ball I draw is green. In this context, (8a) appears to imply (8b).

- (8) a. If I had drawn a red ball or a yellow ball, you would have won.
b. \rightsquigarrow If I had drawn a yellow ball, you would have won.

This apparent entailment remains intact, I think, even after we have both inspected the urn’s content and seen that it does not contain any yellow balls, so that (as I hereby stipulate) there was no chance of drawing a yellow ball.

In this kind of case, one can, however, also utter a “Spain” type conditional. Having established that there were no yellow balls in the urn, I might assert (9).

one where there is a possible outcome of $\neq am$ other than C at t' .

¹⁶ One way to derive this interpretation, drawing on the classical theory of [Kratzer 2012], assumes that the ‘if’ clause restricts a modal: in ordinary counterfactuals, that modal (possibly lexicalized by ‘would’) expresses chance 1, in ‘would probably’ counterfactuals, it expresses high chance. Assuming that probability modals get “restricted” by conditionalization (compare [Yalcin 2010]), it follows that ‘if A , would C ’ expresses $Ch_t(C/A) = 1$, while ‘if A , would probably C ’ expresses that $Ch_t(C/A)$ is high.

(9) If I had drawn a red ball or a yellow ball, I would have drawn a red ball.

This sounds acceptable to my ears. Curiously, after (9) has been accepted, (8a) becomes more acceptable: it no longer seems to imply (8b).

This curious effect has been noted in [Nute 1980a]. Nute observes that we would normally reject (10), as it suggests that Hitler would have been pleased no matter on which side Spain had fought:

(10) If Spain had fought with the Allies or the Axis, Hitler would have been pleased.

(10) becomes more acceptable, however, if it follows (2a).

Needless to say, the chance account does not explain this context-sensitivity. If we wanted to account for the present data in terms of the similarity order, we would have to assume that possibilities can become “close” simply by being mentioned, even if they have zero chance, and that this effect can be blocked by utterances like (9) or (2a). ([Nute 1980a] proposes something like this.)

Here is a somewhat different, but related problem. Compare (10) with (11).

(11) If Spain had joined the war, Hitler would have been pleased.

Modulo some facts that are common ground, the two antecedents are equivalent. Yet (11) is more acceptable than (10). Unlike (10), it does not suggest that Hitler would have been pleased no matter on which side Spain had fought. What explains the difference?

Hájek might argue that (10) and (11) are both false, and that (11) looks more acceptable because the falsity is easier to miss. But take a different example.

(12) If Trump had won in 2020, he would have won a second consecutive term.

This is plausibly true, even by Hájek’s skeptical standards. Trump’s winning in 2020 would have *constituted* winning a second term. No unlikely chance events could have intervened between the antecedent and the consequent, for they describe the same event. But they do so only on the background of what happened in 2016. Thus (13) is false:¹⁷

(13) If Trump had lost in 2016 and won in 2020, he would have won a second consecutive term.

Now consider (14).

(14) If Trump had won in 2020 after either winning or losing in 2016, he would have won a second consecutive term.

This also seems false, as it seems to imply (13). But its antecedent is (practically) equivalent to that of (12).¹⁸

¹⁷ We have here a counterexample to Antecedent Strengthening. Do we also have a counterexample to the generalized form of SDA? I wouldn’t say so. The hypothesis that Trump won in 2020 – the antecedent of (12) – is not *unspecific* with respect to the 2016 election. It isn’t about the 2016 election at all.

¹⁸ One might try to account for the difference between (12) and (14) in terms of the relevant fork time: perhaps the fork time for (14) is before 2016, simply because that year is mentioned in the antecedent, while the fork time for (12) is

Another problem. In the previous section, I argued that the chance account can predict the SDA effect in (6), where the disjuncts have different fork times.

- (6) a. If I had gone for a run at 7am or 2pm, I would feel less tired.
 b. \rightsquigarrow If I had gone for a run at 7am, I would feel less tired.
 c. \rightsquigarrow If I had gone for a run at 2pm, I would feel less tired.

The explanation assumed that the fork time for (6a) coincides with that of (6b), and not with the later fork time of (6c). This is a natural assumption. But it has an unwelcome consequence. We now can't predict the converse inference from (6b) and (6c) to (6a). No independently plausible assumptions about chance would ensure that $Ch_t(C/A) = 1$ and $Ch_{t'}(C/B) = 1$, with $t < t'$, together entail $Ch_t(C/A \vee B) = 1$. Yet the inference from (6b) and (6c) to (6a) looks at least as safe as the SDA inference from (6a) to (6b) and (6c).

This problem does not just arise for the chance account. Any attempt to generalize the “Rewind, Revise, Re-Run” approach to temporally unspecific antecedents like (5) and (6) faces a dilemma. To predict the SDA-type effect these cases seem to display, one must assume that the fork time for a temporally unspecific antecedent can be as early as the fork time for its earliest resolution. If later resolutions can have later fork times, however, this invalidates the converse of SDA.¹⁹

At this point, the prospects for reconciling SDA-type effects with the similarity analysis no longer look as rosy as they did in the previous section. The really bad news is yet to come.

4 Beyond ‘would’ counterfactuals

So far, we have looked at ‘would’ counterfactuals with disjunctive or otherwise unspecific antecedents. But SDA-type effects also arise for ‘might’ counterfactuals (as [Alonso-Ovalle 2006] notes):

- (15) a. If it had rained or snowed, the match might have been cancelled.
 b. \rightsquigarrow If it had rained, the match might have been cancelled.
 c. \rightsquigarrow If it had snowed, the match might have been cancelled.

Let's see if the chance account can explain this. On the chance account, it's natural to assume that a ‘might’ counterfactual $A >_m C$ is true iff $Ch_t(C/A) > 0$, where t is the relevant fork time. This makes ‘might’ and ‘would’ conditionals duals of one another.²⁰ But it does not predict the inference in (15): $Ch_t(C / A \vee B)$ can easily be positive even though $Ch_t(C / B)$ is zero.

There's no easy way out. Suppose we accept SDA for both ‘would’ and ‘might’ counterfactuals, perhaps conditional on certain facts about chance or whatever. Assuming duality, we can then derive

later. I doubt that this gets to the heart of the matter. See [Champollion et al. 2016] and [Romoli et al. 2022] for other cases in which the formulation of the antecedent seems to make a difference.

¹⁹ The converse of SDA is valid in standard similarity semantics such as [Lewis 1973] and [Stalnaker 1968]. The natural assumption that temporally unspecific antecedents can have fork times that precede the fork time of their latest resolutions is therefore incompatible with standard similarity semantics. It leads to an “antecedent relative” similarity semantics. (See, e.g., [Stalnaker 1984: 129ff].)

²⁰ Strictly speaking, we only predict that $\neg(A > C)$ and $A >_m \neg C$ are equivalent if $Ch_t(A) \neq 0$.

$B > C$ from $A > C$:²¹

1. $A > C$ (Assumption)
2. $\neg(A >_m \neg C)$ (1, Duality)
3. $\neg((A \vee B) >_m \neg C)$ (2, SDA for 'might' counterfactuals)
4. $(A \vee B) > C$ (3, Duality)
5. $B > C$ (4, SDA for 'would' counterfactuals)

We can similarly derive $B >_m C$ from $A >_m C$:

1. $A >_m C$ (Assumption)
2. $\neg(A > \neg C)$ (1, Duality)
3. $\neg((A \vee B) > \neg C)$ (2, SDA for 'would' counterfactuals)
4. $(A \vee B) >_m C$ (3, Duality)
5. $B >_m C$ (4, SDA for 'might' counterfactuals)

We could avoid these triviality results by dropping the duality assumption. According to a popular alternative, 'might' counterfactuals should be understood as 'would' counterfactual embedded under 'might'.²² This is not an attractive option for the chance account: 'if the coin were tossed, it might land heads', for example, can surely be true even if the coin is certain to be fair, so that the conditional chance of heads is definitely not 1. To yield sensible verdicts about such cases, the analysis of $A >_m C$ as $\diamond(A > C)$ is usually combined with a view of 'would' counterfactuals that validates "Conditional Excluded Middle":

(CEM) $(A > C) \vee (A > \neg C)$.

This alternative approach, however, is also incompatible with the validity of SDA:

1. $A > C$ (Assumption)
2. $\neg(A > \neg C)$ (1, CNC)
3. $\neg((A \vee B) > \neg C)$ (2, SDA)
4. $(A \vee B) > C$ (3, CEM)
5. $B > C$ (4, SDA)

'CNC' in the second step stands for 'Conditional Non-Contradiction'. It assumes that $A > C$ and $A > \neg C$ cannot both be true. This might fail if A is itself impossible, in which case the possibility of A might be needed as an additional premise.²³

It gets worse. Yet other types of conditionals also give rise to SDA-like effects. Of special interest are deontic conditionals, as in (16).

- (16) a. If you consume alcohol or cocaine, you ought not to drive.

²¹ See [Cariani and Goldstein 2020: 11f.] for a similar result.

²² See, for example, [Stalnaker 1980], [DeRose 1999], [Williams 2010].

²³ Again, see [Cariani and Goldstein 2020] for a series of similar results.

- b. \leadsto If you consume alcohol, you ought not to drive.
- c. \leadsto If you consume cocaine, you ought not to drive.

According to an attractive and popular analysis of deontic conditionals, ‘if A , ought C ’ is true (roughly, and on its most salient reading) iff the best of the circumstantially accessible A worlds are C worlds, where “circumstantial accessibility” is determined by relevant facts and “best” is evaluated in accordance with relevant rules or norms. (See, e.g., [Feldman 1986: ch.4], [von Fintel and Heim 2021]). This is structurally analogous to the similarity analysis of ‘would’ counterfactuals, and it does not validate SDA.

Concretely, suppose consuming alcohol is less bad than consuming cocaine, by the relevant rules or norms, and both possibilities are circumstantially possible. Then the best (accessible) worlds where you consume alcohol or cocaine are worlds where you consume alcohol. In this context, (16a) should imply nothing about what you should do if you consume cocaine. But it does: it seems to imply (16c).²⁴

An SDA-type effect can also be observed with plural descriptions – which are often thought to be structurally analogous to counterfactuals.²⁵

- (17)
- a. The poorest countries once colonized by France or Spain are all in Africa.
 - b. \leadsto The poorest countries once colonized by France are all in Africa.
 - c. \leadsto The poorest countries once colonized by Spain are all in Africa.

This inference should be invalid: if we list the countries once colonized by France or Spain, we might find that the poorest ones have all been colonized by France. (In fact, this seems to be the case.) The fact that these countries are all in Africa should then tell us nothing about the location of the poorest countries once colonized by Spain. Why, then, does the inference in (17) look OK?

Finally, it has often been noted²⁶ that the SDA effect closely resembles the “Free Choice” effect that arises when disjunctions are embedded under possibility modals, as in (18).

- (18)
- a. You may have tea or coffee.
 - b. \leadsto You may have tea.
 - c. \leadsto You may have coffee.

According to the standard semantics of possibility modals, (18a) is true as long as at least one of tea and coffee is allowed. It should not entail (18b) or (18c).

There appears to a general mechanism by which constructions of the form $\phi(A \text{ or } B)$ can be used to

24 As in the case of ‘would’ counterfactuals, this SDA effect seems to generalise beyond disjunction; it arises with ‘may’ as well as ‘ought’; and it can be blocked by moves analogous to that in (2). For example, suppose I am determined to consume alcohol or cocaine before going on a drive. You’d rather that I consume neither. If you think that cocaine is even worse than alcohol, you might still say: ‘if you’re going to consume alcohol or cocaine, you ought to consume alcohol’. This does not suggest that if I consume cocaine then I should (in addition) consume alcohol. (Compare the “gentle murder paradox” of [Forrester 1984].)

25 See, for example, [Schlenker 2004].

26 for example, in [van Rooij 2006], [Klinedinst 2007], and [Franke 2011]

convey $\phi(A)$ and $\phi(B)$, even though this is not supported by the literal meaning of these constructions. An attractive hypothesis is that this mechanism is a kind of implicature.

It does not follow, of course, that constructions of the form $\phi(A \vee B)$ never genuinely entail $\phi(A)$ and $\phi(B)$. Sometimes they do. (Let ϕ be negation.) For the cases above, however, including the case of ‘would’ counterfactuals, there are positive reasons to think that we are dealing with a scalar implicature.

5 The implicature hypothesis

Scalar implicatures arise when a speaker utters a comparatively weak sentence A when they could have chosen a stronger alternative A' . Assuming that the speaker is cooperative and well-informed, one can infer that the alternative is false. The hallmark of scalar implicatures is that they disappear (and even reverse) in negative (downward-entailing) environments, since what is stronger in normal environments here becomes weaker. We can confirm this for SDA:

- (19) a. I don’t think the match would have been cancelled if it had rained.
 b. \rightsquigarrow I don’t think the match would have been cancelled if it had rained or snowed.

Because plain negations of counterfactuals are generally unnatural, I here exploit the “neg-raising” effect of ‘think’, by which ‘I don’t think’ comes to be interpreted as ‘I think not’.²⁷

Now suppose $(A \vee B) > C$ entails $A > C$, in line with SDA. Then $\neg(A > C)$ should entail $\neg((A \vee B) > C)$, and thinking that $\neg(A > C)$ should incur a commitment to $\neg((A \vee B) > C)$. The inference from (19a) to (19b) should sound good. But it does not. (19b) conveys a belief that the match would have gone ahead both in case of rain *and* in case of snow. That is, while $(A \vee B) > C$ appears to entail $A > C$ and $B > C$, $\neg((A \vee B) > C)$ appears to entail $\neg(A > C)$ and $\neg(B > C)$.

I have exploited the reversal of SDA under negation in the triviality arguments from section 4. It is not a coincidence that all my triviality arguments involved the contraposition of SDA – the inference from $\neg(A > C)$ to $\neg((A \vee B) > C)$. While SDA has intuitive force, its contraposition does not.

All the constructions from the previous section (‘might’ counterfactuals, deontic conditionals, etc.) display an SDA-type effect in positive environments, and a reversed effect in negative environments.

Looking at the wider range of constructions offers another argument for an implicature account. I said that there are many constructions ϕ for which $\phi(A \vee B)$ tends to convey both $\phi(A)$ and $\phi(B)$, even though this is not predicted by the standard semantics of these constructions. (We can now add $\neg((A \vee B) > C)$ to the list.) But there are also many constructions for which $\phi(A \text{ or } B)$ does *not* tend

²⁷ If you prefer, you may substitute ‘doubt’ for ‘I don’t think’. The reversal can also be witnessed with quantified constructions:

- (20) a. No animal would go to sleep if threatened or attacked.
 b. \rightsquigarrow No animal would go to sleep if threatened.
 c. \rightsquigarrow No animal would go to sleep if attacked.

The validity of SDA would predict an opposite inference from, for example, (20b) alone to (20a): by SDA, $\neg(Tx > Sx)$ entails $\neg((Tx \vee Ax) > Sx)$. So $[\text{Every } x : Ax]\neg(Tx > Sx)$ should entail $[\text{Every } x : Ax]\neg((Tx \vee Ax) > Sx)$.

to convey $\phi(A)$ and $\phi(B)$. ‘I believe that A or B ’, for example, does not suggest that I believe both A and B .

What makes the difference? Danny Fox has conjectured that the effect – whatever its exact nature – depends on the unavailability of a conjunctive alternative:²⁸ $\phi(A \text{ or } B)$ tends to convey $\phi(A)$ and $\phi(B)$ iff there is no alternative to $\phi(A \vee B)$ that expresses the conjunction of $\phi(A)$ and $\phi(B)$.

To illustrate, let $\phi(A \text{ or } B)$ be ‘You may have tea or coffee’. There are, of course, sentences that express the conjunction of $\phi(A)$ and $\phi(B)$; notably, ‘You may have tea and you may have coffee’. But this is (normally) not an alternative to ‘You may have tea or coffee’, in the semi-technical sense of ‘alternative’ that figures in theories of focus and scalar implicatures.²⁹ Loosely speaking, ‘You may have tea and you may have coffee’ is considerably more complex than ‘You may have tea or coffee’, and this disqualifies it as an alternative. ‘I believe that A or B ’, by contrast, has a conjunctive alternative: ‘I believe that A and B ’. This is no more complex than ‘I believe that A or B ’, and it is plausibly equivalent to ‘I believe that A and I believe that B ’. Fox’s conjecture therefore predicts – correctly – that we get an SDA-type effect for ‘May($A \vee B$)’, but not for ‘Believe($A \vee B$)’. It also predicts that the effect should arise for counterfactuals, since there is no relevant conjunctive alternative to $(A \vee B) > C$.

If Fox’s conjecture is on the right track, the mechanism that gives rise to the conjunctive reading of ‘ $\Phi(A \vee B)$ ’ depends on the availability or unavailability of a conjunctive scalar alternative. This is just what we would expect if the mechanism is (or involves) a scalar implicature.

For the case of counterfactuals, an implicature account also promises to shed light on the recalcitrant data from section 3. Recall, for example, the difference between (21a) and (21b).

- (21) a. If Spain had joined the war, Hitler would have been pleased.
b. If Spain had joined the war on the side of the Allies or the Axis, Hitler would have been pleased.

(21b) suggests that Hitler would have been pleased even if Spain had joined the war on the Allied side; (21a) does not. The two antecedents are practically equivalent, but they have different alternatives: ‘Spain joined the war on the side of the Allies’ is an obvious alternative to the antecedent of (21b), but not to that of (21a).

Let’s see, then, if we can derive the SDA effect as a scalar implicature. There is no shortage of proposals. (See, among others, [Loewer 1976], [Nute 1980a], [Bennett 2003], [Klinedinst 2007], [van Rooij 2010], [Franke 2011], [Bar-Lev and Fox 2020], [Schwarz 2021]). I’ll throw another one into the mix.³⁰

6 A derivation

I’m going to assume, for the sake of argument, that the similarity analysis is correct. We want to explain why an utterance of $(A \vee B) > C$ tends to convey that $A > C$ and $B > C$ are both true, even

28 See [Fox 2007], [Singh et al. 2016], [Bar-Lev and Fox 2020], [Fox and Katzir 2021].

29 See, for example, [Fox and Katzir 2011] and [Breheny et al. 2018].

30 I still have some hope that the account I sketched at the end of [Schwarz 2021] can be made to work. It is not obvious,

if one of A and B is more “remote” than the other.

The standard neo-Gricean recipe for deriving scalar implicatures compares the uttered sentence A with its scalar alternatives $\text{Alt}(A)$. If a cooperative and well-informed speaker uttered A rather than some stronger alternative, we are supposed to infer that the alternatives must be false. Unfortunately, the most obvious scalar alternatives to $(A \vee B) > C$, $A > C$ and $B > C$, are not stronger than $(A \vee B) > C$. Besides, we want to infer that they are true, not false. We have a little more work to do.

The key insight that helps us move further goes back to [Kratzer and Shimoyama 2002]. Kratzer and Shimoyama point out that a well-informed and cooperative speaker might prefer one utterance over another not only if the alternative is false, but also if it would merely have a false implicature.

Suppose we are interested in whether $A > C$ and whether $B > C$. In this context, an utterance of $A > C$ tends to implicate $\neg(B > C)$:

- (22) X: Would the party have been fun if Alice or Bob had come?
Y: It would have been fun if Alice had come.
 \rightsquigarrow Not: It would have been fun if Bob had come.

This is probably an instance of a more general mechanism whereby propositions that are not affirmed are assumed to be denied:³¹

- (23) X: Did Alice or Bob leave early?
Y: Alice left early.
 \rightsquigarrow Not: Bob left early.

The exact mechanism behind this “exhaustification” effect is, to my knowledge, an open question.³² Let’s simply take for granted that in a context where a speaker has uttered $(A \vee B) > C$, each of $A > C$ and $B > C$ tends to implicate the negation of the other.

Now imagine you know that $A > C$ and $B > C$ are both true, and you want to communicate this to your addressee. More precisely, suppose you want to increase your addressee’s degree of belief in $(A > C) \wedge (B > C)$. You could utter $(A > C) \wedge (B > C)$. But suppose you also have a penchant for the pithy. You hate laborious expressions. You could utter $A > C$, or $B > C$. These are nice and short, but they would not lead to accurate beliefs, as either implicates the negation of the other. (Or so we assume.) The optimal choice might be $(A \vee B) > C$. This is still reasonably short, and it should increase the addressee’s confidence in $(A > C) \wedge (B > C)$, if they believe what you say.³³

Now let’s switch perspectives. Imagine you’re the addressee and you receive the message $(A \vee B) > C$. You believe the speaker to be rational, cooperative, and well-informed about $A > C$ and $B > C$.

however, how it can be combined with an ordering semantics for conditionals, how it would extend to ‘might’ conditionals, and why the effect would disappear in negative environments.

31 The implicature from $A > C$ to $\neg(B > C)$ is also closely related to the phenomenon of “Conditional Perfection” ([Geis and Zwicky 1971]), whereby $A > C$ tends to implicate $\neg(\neg A > C)$. This common characterization seems to underdescribe the phenomenon: In a context where three (exclusive) possibilities A_1 , A_2 , and A_3 are under discussion, $A_1 > C$ implicates $\neg(A_2 > C)$ and $\neg(A_3 > C)$. These are not entailed by $\neg(\neg A > C)$.

32 See, for example, [Schulz and Van Rooij 2006] and [Cremers et al. 2023].

33 This follows from the fact that $(A > C) \wedge (B > C)$ entails $(A \vee B) > C$ (according to the similarity analysis and almost all of its rivals). By probability theory, $P(H/E) > P(H)$ whenever H entails E , provided that $P(H) < 1$ and

What information might they have? They can't know that $A > C$ and $B > C$ are both false, as this would contradict the literal content of their utterance. Could they know that $A > C$ is true and $B > C$ false? It seems not. In that case, $A > C$ would have been shorter and more informative (due to the exhaustification effect). By the same reasoning, the speaker can't know that $B > C$ is true and $A > C$ false. The only remaining possibility is that $A > C$ and $B > C$ are both true. As we've just seen, a speaker who knows this may indeed utter $(A \vee B) > C$, if they have a strong enough preference for brevity.

At this stage, we can already predict the SDA inference. But we also predict that an utterance of $(A \vee B) > C$ conveys a strong preference for brevity. We can get rid of this unwanted second prediction by switching perspectives once more.

Imagine again that you are the speaker and you'd like to convey your knowledge of $A > C$ and $B > C$. You wonder what your addressee will infer from an utterance of $(A \vee B) > C$. If they will go through the line of reasoning just described, they will infer that $A > C$ and $B > C$ are both true. This is just what you want to achieve. It follows that you may utter $(A \vee B) > C$ even if you do not have a strong preference for brevity.

Finally, if you are an addressee who believes that the speaker has gone through *this* line of reasoning, you will take an utterance of $(A \vee B) > C$ to signal knowledge of $A > C$ and $B > C$, and you will no longer infer that they have a strong preference for brevity.

The iterated switch of perspectives can't be found in [Grice 1975], but it has a motivation. Grice assumes that speakers tend to be rational and cooperative, and that hearers can draw on this assumption to interpret an utterance. Grice did not realize that this leads to a regress.³⁴ In order to choose an utterance, a speaker needs to anticipate how the hearer would respond. The hearer's response, however, turns on their reconstruction of the speaker's decision problem. The speaker needs a model of the hearer that includes a model of the speaker, and so on.

Since our cognitive resources are limited, we can expect these models within models to be increasingly unsophisticated, bottoming out in a "literal" hearer or a "literal" speaker who does not engage in strategic reasoning.

You may doubt that we engage in *any* kind of strategic reasoning when we interpret simple utterances like (1a). This may be right. My suggestion is only that we interpret utterances *as if* we were engaged in such reasoning. In practice, we may well have learned to draw the relevant inferences automatically, at least if nothing alerts us to the falsity of their premises.

The model of communication I have outlined is formalized in the "Rational Speech Act" (RSA) framework (See, e.g., [Franke 2017], [Scontras et al. 2018].) With a bit of tidying, the above derivation can be spelled out as an RSA model. I won't descend into the formal details here.³⁵ But I want to flag one point that needs tidying.

I claimed that a well-informed speaker with a sufficiently strong preference for brevity would choose $(A \vee B) > C$ only if they know that $A > C$ and $B > C$ are both true. A speaker who knows that $A > C$ is true and $B > C$ false, for example, would prefer $A > C$ over $(A \vee B) > C$, as it conveys more information about $A > C$ and $B > C$. This is correct. By the similarity analysis,

$0 < P(E) < 1$.

³⁴ [Lewis 1969] saw it. See, esp., pp.27ff.

³⁵ My derivation loosely resembles the derivation of Free Choice in [Champollion et al. 2019].

however, $(A \vee B) > C$ carries information not just about $A > C$ and $B > C$, but also about the relative closeness of A and B . The speaker might prefer $(A \vee B) > C$ over $A > C$ because they want to convey this extra information. Concretely, imagine $(A > C) \wedge \neg(B > C)$ is already common ground, but it is not yet settled which of A and B is closer to the actual world. A speaker might then use $(A \vee B) > C$ to convey that the closest A worlds are closer than the closest B worlds.

But this seems to be unusual. In normal situation, when a speaker utters $(A \vee B) > C$, they address whether $A > C$ and whether $B > C$; the relative closeness of A and B is not of interest. We can use this datum to fix the derivation. The norms of conversation only require speakers to be informative with respect to questions that are of interest (“under discussion”). In normal contexts, the fact that $(A \vee B) > C$ carries information about the relative closeness of A and B is therefore no reason to prefer it.

The question-sensitivity of the derivation might shed light on the puzzle that (24b) becomes more acceptable if it follows (24a), as noted by Nute:

- (24) a. If Spain had fought with the Allies or the Axis, it would [probably] have fought with the Axis.
 b. If Spain had fought with the Allies or the Axis, Hitler would [probably] have been pleased.

(24a) is most naturally interpreted as conveying information about the relative closeness of the *Allies* and the *Axis* possibilities.³⁶ Once this question has been raised, it might still be salient when we interpret (24b).

Looking back at other data from the previous sections, the derivation I have outlined promises to explain why SDA effects are sensitive to substitution of equivalents and why they disappear under negation.³⁷ Since it is not based on doctoring with the meanings of ‘>’ and ‘or’, the derivation smoothly generalizes to other constructions. For example, to explain why $(A \vee B) >_m C$ tends to convey $(A >_m C) \wedge (B >_m C)$, we would assume that $A >_m C$ is normally exhausted to implicate $\neg(B >_m C)$. A cooperative and well-informed, but not overly sophisticated speaker would therefore choose $(A \vee B) >_m C$ only if they know that $A >_m C$ and $B >_m C$ are both true, and if they have a strong preference for brevity. As before, a more sophisticated speaker might use $(A \vee B) >_m C$ even without a strong preference for brevity. (The triviality arguments from section 4 fail because the derivations do not go through under negation.)

What about unspecific antecedents that aren’t (binary) disjunctions? Suppose an antecedent A divides into two sub-possibilities, A_1 and A_2 .³⁸ Unlike the disjuncts of a disjunction, these do not

36 Our derivation of SDA does not go through for (24a) because *Axis > Allied* can hardly be exhausted to imply the negation of *Allied > Allied*. More generally, it is implausible that a speaker of (24a) wants to convey information about $?(\textit{Axis} > \textit{Allied})$ and $?(\textit{Allied} > \textit{Allied})$.

37 Can we explain that the effect *reverses* under negation? Note that $\neg((A \vee B) > C)$ semantically entails $\neg(A > C) \vee \neg(B > C)$, just as $(A \vee B) > C$ entails $(A > C) \vee (B > C)$. Its literal content therefore provides some information about whether $A > C$ and whether $B > C$. We can explain the inference to $\neg(A > C) \wedge \neg(B > C)$ if we assume that the alternative $\neg(A > C)$ would get exhausted to $\neg(A > C) \wedge (B > C)$. This is not obvious if the question under discussion is whether $A > C$ and whether $B > C$. (See, e.g., [Spector 2007], [Schulz and Van Rooij 2006].) But perhaps the question that is addressed in cases where we find reversal is which of $A > C$ and $B > C$ is *not* the case.

38 That is, assume that A is (contextually) equivalent to $A_1 \vee A_2$, and that $A_1 \wedge A_2$ is (contextually) ruled out. When I talk of “possibilities”, I equivocate, for the sake of simplicity, between expressions and their meanings.

automatically qualify as scalar alternatives to A and to one another. But they plausibly do so if they are sufficiently salient in the conversational context.³⁹ In that case, one can run essentially the same derivation as above to show that $A > C$ implicates $A_1 > C$ and $A_2 > C$.

What if the antecedent divides into more than two salient sub-possibilities A_1, A_2, \dots, A_n ? If all conjunctions (unions) of these possibilities are also salient, we can predict that $A > C$ implicates $A_i > C$ for all A_i .⁴⁰ If no intermediate possibilities are salient, the account may only predict an inference that $A_i > C$ is true for *multiple* A_i .⁴¹ In [Schwarz 2021], I argued that this “partial” or “existential” prediction is correct for Free Choice. ‘Fred might be in France’, for example, tends to convey that there are multiple locations in France where Fred might be, but not that every location in France is a possibility. Whether this prediction is also true for counterfactuals is not clear to me.⁴²

As I said, I don’t want to go too far into the details of the proposed derivation. I have introduced it only as a proof of concept, as an illustration of how SDA effects might be explained as a kind of implicature. All I want to suggest is that this direction looks more promising than trying to derive the effects by tweaking the semantics of counterfactuals.⁴³⁴⁴

7 The case for skepticism

Let’s now return to the case for counterfactual skepticism. The argument from unspecific antecedents, you may recall, starts with the observation that ordinary counterfactuals often have unspecific antecedents. It assumes that a counterfactual with an unspecific antecedent is true only if it remains true on every way of resolving the unspecificity. Since one can often find resolutions that make the counterfactual false, the argument concludes that most ordinary counterfactuals are false.

I have cast doubt on the central premise: that a counterfactual with an unspecific antecedent is true only if it remains true on every way of resolving the unspecificity. This assumes that a generalized form of SDA is semantically valid. I have argued that it is not. Granted, counterfactuals with unspecific antecedents often *appear* to entail corresponding counterfactuals with more specific antecedents. But the appearance is misleading. The effect is better explained as an implicature.⁴⁵

39 This is a common assumption in the literature; see, for example, [Katzir 2007] and [Breheny et al. 2018].

40 We need to assume that if A' is an alternative that excludes some A_j , then $A' > C$ would be exhausted to imply $\neg(A_j > C)$. A speaker who knows $\bigwedge_i A_i > C$ would therefore prefer $A > C$ over $A' > C$. A speaker who knows that $A' > C$ and $\neg(A_j > C)$, by contrast, would prefer $A' > C$. So $A > C$ conveys that $A_i > C$ for all A_i .

41 In this case, $A > C$ might also be the optimal choice for a speaker who knows, say, $A_1 > C \wedge \dots \wedge A_{n-1} > C \wedge \neg A_n > C$. It would not, however, be optimal for a speaker who knows that a single $A_i > C$ is true and the rest are false, since $A_i > C$ would then convey more relevant information.

42 If the sub-possibilities A_1, \dots, A_n are equally close, similarity semantics alone implies that $A > C$ can only be true if each $A_i > C$ is true. A test case would have to involve sub-possibilities with different degrees of remoteness. We also have to ensure that no intermediate possibilities are salient. Both of these are somewhat hard to assess independently.

43 [Fox and Katzir 2021] complain that RSA-type accounts of Free Choice and SDA are implausibly sensitive to the agents’ priors, and that they fail to account for the effect with more than two disjuncts. This is true for the account of [Franke 2011] and the toy model constructed by Fox and Katzir, but not of the model I have sketched (nor of that in [Champollion et al. 2019]).

44 Yet other proposals could, of course, be explored. An important third direction expands on the idea that ‘or’ may not simply express Boolean disjunction. See, for example, [Alonso-Ovalle 2008], [Ciardelli et al. 2018], [Romoli et al. 2022].

45 The same kind of response is available to [Loewenstein 2021]’s argument from reverse Sobel sequences.

Let's have another look at the example from section 1. Bob is on a medication that would cause an allergic reaction if he consumed alcohol. In this context, (25) seems true.

(25) If Bob had wine, his medication would cause an allergic reaction.

The skeptic now raises the possibility of de-alcoholized wine: if Bob had de-alcoholized wine, his medication would not cause a reaction. Once this possibility is raised, (25) looks problematic.

The implicature account explains why (25) might look problematic, even if it is true. 'Bob has wine' divides into 'Bob has regular wine' and 'Bob has de-alcoholized wine'. If both of these are contextually salient, (25) tends to implicate that in either case, Bob's medication would cause an allergic reaction.

The explanation is different from more familiar contextualist moves. A contextualist might say that de-alcoholized wine possibilities *become* non-remote by being made salient (compare [Lewis 2016]). Alternatively, they might say that a counterfactual is true as long as all the *relevant* closest *A* worlds are *C* worlds, and that de-alcoholized wine worlds are normally irrelevant (compare [Sandgren and Steele 2021]). I'm suggesting that it does not matter if de-alcoholized wine worlds are remote or close. (I suspect this depends on the details of the scenario.) Either way, attending to these worlds will tend to associate (25) with a false implicature.

There are, of course, other arguments for counterfactual skepticism. There is the argument from chance (see [Hájek 2021a] or [Hájek 2023]). (26), for example, seems true.

(26) If I had put an ice cube in my tea, it would have melted.

Yet quantum physics tells us that a state in which an ice cube has been put into my tea has a positive chance of evolving into a state in which the entire liquid is frozen. This suggests that the ice cube *might not* have melted. And doesn't this, in turn, suggest that (26) is false?

The argument from chance does not rely on the unspecificity of the antecedent, so the implicature account has little to say about it. It does, however, have something to say about a version of the argument that appeals to (classical) *statistical* mechanics rather than quantum mechanics. Classical statistical mechanics also assigns a positive probability to the tea freezing, but not because it posits an indeterministic dynamics. According to statistical mechanics, the melting of an ice cube in a hot liquid depends on the exact configuration of all relevant molecules. Some configurations would lead to a freezing of the liquid. If these configurations are no more remote than any others, (26) is predicted to be false. (See, e.g., [Bennett 2003: 247].)

Why should we think that the peculiar configurations are not remote? Remember that remoteness is a technical concept, not to be judged by intuitions about distance or similarity. It may *seem* that the peculiar configurations are no more remote than others because attending to them makes (26) seem false. The implicature account can explain away this appearance. It predicts that attending to the peculiar configurations can make (26) seem false, even if it is true.

References

Luis Alonso-Ovalle [2006]: *Disjunction in Alternative Semantics*. PhD Thesis

-
- [2008]: “Alternatives in the Disjunctive Antecedents Problem”. In *Proceedings of the 26th West Coast Conference on Formal Linguistics*, 42–50
- Moshe E. Bar-Lev and Danny Fox [2020]: “Free Choice, Simplification, and Innocent Inclusion”. *Natural Language Semantics*, 28(3): 175–223
- Jonathan Bennett [2003]: *A Philosophical Guide to Conditionals*. New York: Oxford University Press
- Richard Breheny, Nathan Klinedinst, Jacopo Romoli and Yasutada Sudo [2018]: “The Symmetry Problem: Current Theories and Prospects”. *Natural Language Semantics*, 26(2): 85–110
- Rachael Briggs [2012]: “Interventionist Counterfactuals”. *Philosophical Studies*, 160(1): 139–166
- Fabrizio Cariani and Simon Goldstein [2020]: “Conditional Heresies”. *Philosophy and Phenomenological Research*, 101(2): 251–282
- Lucas Champollion, Anna Alsop and Ioana Grosu [2019]: “Free Choice Disjunction as a Rational Speech Act”. *Semantics and Linguistic Theory*: 238–257
- Lucas Champollion, Ivano Ciardelli and Linmin Zhang [2016]: “Breaking de Morgan’s Law in Counterfactual Antecedents”. In *Semantics and Linguistic Theory*, vol 26. 304–324
- Ivano Ciardelli, Linmin Zhang and Lucas Champollion [2018]: “Two Switches in the Theory of Counterfactuals”. *Linguistics and Philosophy*, 41(6): 577–621
- Alexandre Cremers, Ethan G. Wilcox and Benjamin Spector [2023]: “Exhaustivity and Anti-Exhaustivity in the RSA Framework: Testing the Effect of Prior Beliefs”. *Cognitive Science*, 47(5): e13286
- Keith DeRose [1999]: “Can It Be That It Would Have Been Even Though It Might Not Have Been?”. *Philosophical Perspectives*, 13: 385–413
- Cian Dorr [2016]: “Against Counterfactual Miracles”. *Philosophical Review*, 125(2): 241–286
- Fred Feldman [1986]: *Doing the Best We Can: An Essay in Informal Deontic Logic*. Dordrecht: Reidel
- Kit Fine [1975]: “Review of Lewis, *Counterfactuals*”. *Mind*, 84: 451–458
- [2012]: “Counterfactuals without Possible Worlds”. *Journal of Philosophy*, 109: 221–246
- James William Forrester [1984]: “Gentle Murder, or the Adverbial Samaritan”. *The Journal of Philosophy*, 81(4): 193–197
- Danny Fox [2007]: “Free Choice and the Theory of Scalar Implicatures”. In U. Sauerland and P. Staveva (Eds.) *Presupposition and Implicature in Compositional Semantics*, Basingstoke: Palgrave Macmillan, 71–120

-
- Danny Fox and Roni Katzir [2011]: “On the Characterization of Alternatives”. *Natural Language Semantics*, 19: 87–107
- [2021]: “Notes on Iterated Rationality Models of Scalar Implicatures”. *Journal of Semantics*, 38(4): 571–600
- Michael Franke [2011]: “Quantity Implicatures, Exhaustive Interpretation, and Rational Conversation”. *Semantics and Pragmatics*, 4: 1:1–82
- [2017]: “Game Theory in Pragmatics: Evolution, Rationality, and Reasoning”. In *Oxford Research Encyclopedia of Linguistics*,
- Michael L. Geis and Arnold M. Zwicky [1971]: “On Invited Inferences”. *Linguistic Inquiry*, 2(4): 561–566
- H. Paul Grice [1975]: “Logic and Conversation”. In P. Cole and J. Morgan (Eds.) *Syntax and Semantics*, vol 3. 41–58
- Mario Günther [2017]: “Disjunctive Antecedents for Causal Models”. In *21 St Amsterdam Colloquium*, 25
- Alan Hájek [2021a]: “Contra Counterfactualism”. *Synthese*, 199(1): 181–210
- [2021b]: “Counterfactual Scepticism and Antecedent-Contextualism”. *Synthese*, 199(1): 637–659
- [2023]: “Most Counterfactuals Are False”
- Roni Katzir [2007]: “Structurally-Defined Alternatives”. *Linguistics and Philosophy*, 30: 669–390
- Nathan W. Klinedinst [2007]: *Plurality and Possibility*. PhD Thesis, Los Angeles
- Angelika Kratzer [2012]: *Modals and Conditionals*. Oxford: Oxford University Press
- Angelika Kratzer and Junko Shimoyama [2002]: “Indeterminate Pronouns: The View from Japanese”. In *Proceedings of the 3rd Tokyo Conference on Psycholinguistics*, Tokyo: Hituzi Syobo, 1–25
- Daniel Lassiter [2018]: “Complex Sentential Operators Refute Unrestricted Simplification of Disjunctive Antecedents”. *Semantics and Pragmatics*, 11: 9:EA–9:EA
- Hannes Leitgeb [2012]: “A Probabilistic Semantics for Counterfactuals. Part a”. *The Review of Symbolic Logic*, 5(01): 26–84
- David Lewis [1969]: *Convention: A Philosophical Study*. Cambridge (Mass.): Harvard University Press
- [1973]: *Counterfactuals*. Oxford: Blackwell

-
- [1977]: “Possible-World Semantics for Counterfactual Logics: A Rejoinder”. *Journal of Philosophical Logic*, 6: 359–363
- [1979]: “Counterfactual Dependence and Time’s Arrow”. *Noûs*, 13: 455–476
- Karen S. Lewis [2016]: “Elusive Counterfactuals”. *Noûs*, 50(2): 286–313
- Yael Loewenstein [2021]: “Heim Sequences and Why Most Unqualified ‘Would’-Counterfactuals Are Not True”. *Australasian Journal of Philosophy*, 99(3): 597–610
- Barry Loewer [1976]: “Counterfactuals with Disjunctive Antecedents”. *The Journal of Philosophy*, 73(16): 531–537
- [2007]: “Counterfactuals and the Second Law”. In *Causation, Physics, and the Constitution of Reality: Russell’s Republic Revisited*, Oxford University Press
- Thomas McKay and Peter Van Inwagen [1977]: “Counterfactuals with Disjunctive Antecedents”. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 31(5): 353–356
- Donald Nute [1975]: “Counterfactuals and the Similarity of Words”. *The Journal of Philosophy*, 72(21): 773–778
- [1980a]: “Conversational Scorekeeping and Conditionals”. *Journal of Philosophical Logic*, 9(2): 153–166
- [1980b]: *Topics in Conditional Logic*, vol 20. Springer Science & Business Media
- Judea Pearl [2000]: *Causality: Models, Reasoning and Inference*. Cambridge: Cambridge University Press
- Jacopo Romoli, Paolo Santorio and Eva Wittenberg [2022]: “Alternatives in Counterfactuals: What Is Right and What Is Not”. *Journal of Semantics*, 39(2): 213–260
- Alexander Sandgren and Katie Steele [2021]: “Levelling Counterfactual Scepticism”. *Synthese*, 199(1): 927–947
- Philippe Schlenker [2004]: “Conditionals as Definite Descriptions”. *Research on Language and Computation*, 2: 417–462
- Katrin Schulz and Robert Van Rooij [2006]: “Pragmatic Meaning and Non-monotonic Reasoning: The Case of Exhaustive Interpretation”. *Linguistics and Philosophy*, 29(2): 205–250
- Wolfgang Schwarz [2021]: “Discourse, Diversity, and Free Choice”. *Australasian Journal of Philosophy*, 99(1): 48–67
- Gregory Scontras, Michael Henry Tessler and Michael Franke [2018]: “Probabilistic Language Understanding: An Introduction to the Rational Speech Act Framework”. Retrieved March 2023, 17: 2021

-
- Raj Singh, Ken Wexler, Andrea Astle-Rahim, Deepthi Kamawar and Danny Fox [2016]: “Children Interpret Disjunction as Conjunction: Consequences for Theories of Implicature and Child Development”. *Natural Language Semantics*, 24(4): 305–352
- Brian Skyrms [1980]: “The Prior Propensity Account of Subjunctive Conditionals”. In W. Harper, R. Stalnaker and G. Pearce (Eds.) *Ifs: Conditionals, Belief, Decision, Chance and Time*, Springer, 259–265
- Benjamin Spector [2007]: “Scalar Implicatures: Exhaustivity and Gricean Reasoning”. In *Questions in Dynamic Semantics*, Brill, 225–249
- Robert Stalnaker [1968]: “A Theory of Conditionals”. In N. Rescher (Ed.) *Studies in Logical Theory*, Oxford: Blackwell, 98–112
- [1980]: “A Defense of Conditional Excluded Middle”. In William Harper, Robert C. Stalnaker and Glenn Pearce (Eds.) *Ifs*, Reidel, 87–104
- [1984]: *Inquiry*. Cambridge (Mass.): MIT Press
- Robert van Rooij [2006]: “Free Choice Counterfactual Donkeys”. *Journal of Semantics*, 23(4): 383–402
- Robert van Rooij [2010]: “Conjunctive Interpretations of Disjunctions”. *Semantics and Pragmatics*, 3: 11:1–28
- Kai von Fintel and Irene Heim [2021]: *Intensional Semantics*
- Ken Warmbrod [1981]: “Counterfactuals and Substitution of Equivalent Antecedents”. *Journal of Philosophical Logic*, 10(2): 267–289
- J. Robert G. Williams [2010]: “Defending Conditional Excluded Middle”. *Noûs*, 44(4): 650–668
- Seth Yalcin [2010]: “Probability Operators”. *Philosophy Compass*, 5(11): 916–937